

# **Method and Intelligence**

Digital Approaches to Memory and Communication in Historiography

Jonah Lynch (matricola 491509)

Tesi per il conseguimento del Ph.D. in Storia, XXXVI Ciclo

Dipartimento di Scienze Politiche e Sociali

# Table of Contents

Chapter 1 – Making Models.....	4
Introduction.....	4
Method.....	7
Connections Between Data.....	13
Inner-referential Interpretation.....	16
Split Fragments, Recomposed Wholes.....	20
Chapter 2 – Elements of Digital Humanism.....	24
Natural Language Processing.....	26
Neural Correlates and Language Models.....	28
Critiques of Statistical Language Models: Where Word Vectors Fail.....	31
Time Dependency.....	33
Words Slip, Slide, and Sometimes Break.....	35
In Defense of Using Language Models.....	37
Example: Are Mesopotamian Gods Active?.....	38
Transitivity from a grammatical point of view.....	41
Example: Comparing Shuila prayers with the Psalms.....	46
Structure.....	51
Example: Representing Structural Differences.....	52
Conclusion.....	59
Chapter 3 – Intellectual Prosthetics.....	61
The ‘Extended Mind’.....	61
The ‘Battle for Synthesis’.....	63
A Short Recent History of Memory Aids.....	66
Vannevar Bush and the Memex.....	69
J.C.R. Licklider: the Birth of the Knowledge Worker.....	70
Douglas Englebart: Augmenting Human Intellect.....	72
Ted Nelson: the Hyperlink Controversy.....	74
Tim Berners-Lee: Connecting Computers.....	75
Building a ‘Memex’.....	78
Building Text Networks.....	83
What is a node?.....	86
What are relations?.....	88
Interfacing with the machine.....	96
Using the Memex.....	101
Comparing a Book to its Reviews.....	101
Expanding the experiment.....	110
Chapter 4 – Models of Historiography.....	118
Writing the Mes-Rel Website.....	119
Wholes and Parts.....	120
Digital Writing.....	124
The core narrative(s).....	125
Bibliography and notes: in support of the core narrative(s).....	126
Excerpts, reviews, sources.....	128
Themes and monographs.....	128

Indices and ‘maps’ .....	129
Automation.....	130
Validating the Mes-Rel Website.....	134
The Memex as “X-Ray”: The Multinodal Index.....	137
The Memex as “Bird’s Eye View”: Validating the Bibliography.....	150
Conclusion.....	157
Chapter 5 – Further Models of Historiography.....	158
The Memex as “Fisheye Lens”: the Author Comparison Tool.....	158
The Memex as “Filmstrip”: Time-dependent Projection of JSTOR Datasets.....	165
Representation, encoding, and intelligence.....	170
The Memex as “Microscope”: Visualizing the Atoms of Discourse.....	176
Simple Topological Features.....	182
Overview.....	186
Chapter 6 – The Struggle for Synthesis.....	193
What is Digital Humanism?.....	198
The Database Problem.....	200
From Data to Conclusions.....	202
Redefining the Humanities?.....	209
Conclusion.....	214
Appendix.....	216
Encoding divinities in Cypher.....	216
How are relations encoded?.....	218
Logical relations.....	221
Bibliography.....	223

# Chapter 1 – Making Models

## Introduction

In these pages I will present an argument for the use of computational machinery in historical research, describe the theory and practice of one such method as instantiated in a computer program that I have written, and also apply this program to several historical questions in order to show the relevance of the method.

Documenting historical details and their provenance occupied historians like Ranke and de Thou, and the history of the footnote apparatus shows both the increasingly close attention historians have paid to documenting their sources, as well as the perennial human tendency to prefer the appearance of erudition to the genuine article.<sup>1</sup> Detailed attention to sources is as often a cudgel to beat one's opponents with as it is a sincere instrument of the search for accuracy and precision in one's truth-telling. I do not think my embryonic method is immune to this problem, nor do I expect myself and my contemporaries to be naturally more truth-seeking than our ancestors were.

But to the degree that I and my contemporaries *do* care about truth, there is the serious and disturbing issue of the multiplicity of sources that in good conscience we should consider in the attempt to discern the features of the past, yet cannot, due to the limitations of our capacity to read, remember, and synthesize information. Could a computer help us do better?

It is useful to attend to the “bootstrapped” character that human knowledge has. In the adventure of knowing the world, humans cannot avail themselves of an external vantage-point. We see the world only through our own eyes, with the prejudices of our own cultures and experiences, and within the confines of the categories in our languages. This means that our knowing is necessarily incomplete and structurally biased by the viewpoint of each observer. However, through language and social structures, we have been able to continually expand the

---

1 Anthony Grafton, “The Footnote from De Thou to Ranke,” *History and Theory* 33, no. 4 (December 1994): 53, <https://doi.org/10.2307/2505502>.

sphere of our knowledge. This thesis presents a new method which could help us to expand our knowledge even further, both by using computational methods to peer more efficiently into the mass of data we already have, and by scaffolding collaboration between researchers.

The invention of writing greatly increased the ability of thought to exist without a thinker, speech without a speaker. Concepts could be reified and travel beyond the bounds of the place and the time where they were conceived. It became possible for sages to peer into the wisdom of many peoples, and compare their own local interpretation of things with experiences well beyond their ken. The thoughts of people distant in time and space could be incorporated into the web of ideas and explanations that each scholar made their own.

The central problem discussed in this dissertation comes into view at this point. The quantity of information written on tablets or on papyrus contained in great libraries grew without bound. In the early stages of this process, it was possible for great kings to aspire to collect every text that had been written. This situation remained perhaps until the time of the library of Nineveh, or perhaps someone had read everything in the great library of Alexandria. But at some point, the quantity of information available exceeded what even the most diligent scholar could collect and read in a lifetime.

What then? How could new knowledge be generated by a scholar unfamiliar with what had already been discovered? In some fields, it is possible to add to the sum total of knowledge without having an encyclopedic grasp of the subject. But many important questions link multiple domains of inquiry together.

It is easy to see that at some point, the complexity or even the sheer quantity of relevant information cannot possibly be managed by a human brain. And yet, in some sense, it must. For we will not cease from exploration. Adding incremental contributions to the galaxies of information already discovered can be the necessary preparation for true scientific advances. The patient compilers of dictionaries or tables of logarithms prepare the ground for the syntheses discerned by the greatest minds, those that discern large-scale patterns in the cloud of apparently disconnected data.

As historians, we want to understand the story told by the fragments of the past. Our close cousins the archaeologists seek not only the dimensions, composition, and emplacement of the

sherds they find, but also their use, their antecedents, and their meaning. We want to perceive the life of our ancestors, understand their reasoning, vibrate sympathetically with their experience of the world. This is the work of interpretation. But the problem mentioned about is a formidable challenge. How can our interpretations of ancient civilization claim to be more than fiction? We must take into account the data. And if the data exceeds our ability to remember, we must build tools, prosthetics of our bodies and minds.

I hope to contribute to the solution of this problem in this thesis by exploring a sub-set of the problem as I experienced it while working with a small research team on questions related to the history of Mesopotamian religion. The four researchers in our team shared bibliographical references and met for a few hours together each week. Inevitably, some of us had read articles or books that others had not. The varying points of view of the scholars were the expression of the richness of the research team, as well as a problem that had to be resolved when it came time to write. Whose opinion should be expressed? Which interpretation most fully covered the available information? Our project was in support of the publication of a large-scale synthesis written by Giorgio Buccellati, and part of our aim was to buttress or critique his arguments through our work. As we gathered hundreds of relevant sources, I began to wonder how we could map our coverage of his text. How could we direct our limited energies toward those parts of the work that remained undone? And how could we identify which topics were incomplete?

\* \* \*

The result, described in detail in this thesis, is a model of historical knowledge expressed in a computer program I wrote during my doctoral research. I began with the large-scale problem described above and attempted to build a technical solution to it. My method has been iterative, proceeding first from doing, and then to critical reflection upon what I have done, including engagement with scholarly literature, both in history and interdisciplinary sources. Several times, I returned to the start and created a new version of the “product” that I set out to build, and each time I also used the product in its current state to address historical questions. These

iterations will continue after the defense of this thesis, since the problem is far larger than a single scholar can hope to address in a few years.

This method may be more typical of engineering or product design than history or epistemology. However, I believe it is also an adequate method, precisely in the classic sense of *adequatio*: the gradual, iterative molding of the human mind to form a map of the domain it is engaged in studying. This map is never complete, and must continually be improved and critiqued. In my view, the ever-incomplete and ever-improved map of reality that single scholars and groups of scholars produce is an essential feature of the study of history today.<sup>2</sup> It is also worth mentioning at the outset that my method appears to be original. I am not aware of any analogous methods having been proposed in the field of Digital Humanities; two recent publications from Oxford University Press and Cambridge University Press describe techniques that come close to my approach, but do not arrive at it.<sup>3</sup>

## Method

Before discussing the specific method I have developed, it is useful to gesture at the broader conversation about method in general. The topic is far too broad for more than a

- 
- 2 Calls for an increase in collaboration and joint publication in the humanities have been made for decades. While it is common practice for natural scientists to work and write together, humanists tend to still prefer writing alone. The literature on the topic is vast. A starting point referring generally to humanities is Simon Moreton, “Rethinking ‘Knowledge Exchange’: New Approaches to Collaborative Work in the Arts and Humanities,” *International Journal of Cultural Policy* 22, no. 1 (January 1, 2016): 100–115, <https://doi.org/10.1080/10286632.2015.1101081>. On collaboration in the specific domain of digital humanities, see Gabriele Griffin and Matt Steven Hayler, “Collaboration in Digital Humanities Research – Persisting Silences,” n.d.Griffin and Hayler write: “collaboration has become a hallmark of Digital Humanities (DH) research. Nonetheless it remains under-discussed and for those not deeply engaged in DH a bit of a mystery. Drawing on recent DH work and publications that engage with questions of DH collaboration in different ways (e.g. [Deegan and McCarthy] [Griffin and Hayler 2016] [Hayler and Griffin 2016]), we analyze three types of DH collaboration: 1) human-human interactions; 2) human-machine/material interactions; and 3) machine/material-machine/material interactions. We argue that engagement with collaboration processes and practices enables us to think through how DH tools and practices reinforce, resist, shape, and encode material realities which both pre-exist, and are co-produced by them. We suggest that understanding these entanglements facilitates a critical DH in which academic hierarchies and disciplinary preconceptions are challenged.”
- 3 See Martin Paul Eve, “Distance and History,” in *The Digital Humanities and Literary Studies*, by Martin Paul Eve, 1st ed. (Oxford University PressOxford, 2022), 129–50, <https://doi.org/10.1093/oso/9780198850489.003.0005>. See also Peter De Bolla, ed., *Explorations in the Digital History of Ideas: New Methods and Computational Approaches* (Cambridge New York: Cambridge University Press, 2024), <https://doi.org/10.1017/9781009263610>.

summary presentation in this context. The goal of this section is to note the multiplicity of methodological approaches and to establish a few points in the existing literature to which my own project is anchored. It also aims to specify an important epistemological choice that underlies my method: in favor of a “realistic” stance with regard to knowledge about the past, neither absolute in its affirmation of “objective” knowledge, nor “relativist” in its admission of the multiplicity of narratives and interpretations. This clarification is important to keep in mind later on, when my deployment of mathematical tools and computer programs might seem to be “positivistic” in nature. I endeavor to make use of the power of mathematical methodologies, while always maintaining the critical consciousness both of the limits of these methods and of the fragmentary and biased nature of the underlying data.

First, methodological multiplicity. Galyna Bezarova notes that historians make use of multiple methodologies, and that these methods are complex and to some degree conflicting. She affirms that “the first task of the methodology of history is the establishment of principles of ‘organization’ of the material”<sup>4</sup>, its selection, and the determination of the subject and object of study. But therein lies the crucial problem: “the theoretical position of the historian is largely determined by his own choice and independent creative development, depending on the current level of knowledge and the nature of the social environment”. For Bezarova, the method of the historical science is not an anchor that the scholar can hold firm to, but is rather a chosen lens, with advantages and disadvantages, through which to view the object under study.

Methodological multiplicity may derive from generations of sophisticated scholarship, and the various schools of thought that have arisen are each in their own way serious attempts at handling the data in a respectful manner. But according to Bezarova, the historian’s choice of which school to belong to is largely a *choice*, depending on their “creative development” and social milieu. What are the main choices they might make?

Roberto Franzosi answers this question by starting from the binary distinction between narrative or “scientific/cliometric” history as expressed in Elton and Fogel’s 1983 work *Which*

---

4 Galyna Bezarova, “The Problem of the Subject and Object in the Methodology of History,” *Current Issues of Social Sciences and History of Medicine*, no. 2 (February 23, 2017): 67, <https://doi.org/10.24061/2411-6181.2.2017.41>.



*Road to the Past? Two Views of History.*<sup>5</sup> Elton had focused on the journalistic “5 Ws + H”, derived from the older septet plus how: “who, what, and where, by what helpe, and by whose, why, how and when”, to quote the sixteenth century spelling in Franzosi’s article.

Elton’s focus was on narrative, while Fogel emphasized analysis and quantification. Franzosi proposes his own method of Quantitative Narrative Analysis as a resolution of the dichotomy. He represents the 5 W’s + H in rigorous detail within a computer database, which allows him to apply the “invariant narrative structure in large-scale socio/historical projects”.<sup>6</sup> Finally, Franzosi’s approach applies data visualization tools to “reflect the underlying properties of narrative, namely, network models that graph the Who and the What, pro or against whom, and GIS models that map Where and When, ‘Who did the deede and What was done’”.<sup>7</sup>

Franzosi applies his method to the Georgia lynchings (1875-1930) to show the power of his computer-aided approach to analyzing and displaying the data. He then wonders if his approach would pass muster with Elton, who wrote in 1967 that historians

go to sociology and the like not only for inspiration but especially for method and can speak with hope of an ‘age of the historical factory’ with its cooperative and organized scholarship calling upon arithmetic for aid. I am less frightened by the thought of cooperative research or ‘quantification’ than unimpressed, so far, by its results, and I believe that these supposedly sophisticated innovators are guilty of a little naivety.<sup>8</sup>

Franzosi thinks that he escapes Elton’s judgment at least in part, in that his approach is highly narrative, and indeed can *only* be successfully applied to narrative texts. What is more, he asserts that he had “never done as much *close reading* of lynching articles as *after* [he] had developed the specialized search routines for *distant reading!*”.<sup>9</sup> He might, however, be guilty of speaking to sociologists.

---

5 Roberto Franzosi, “A Third Road to the Past? Historical Scholarship in the Age of Big Data,” *Historical Methods: A Journal of Quantitative and Interdisciplinary History* 50, no. 4 (October 2, 2017): 227–44, <https://doi.org/10.1080/01615440.2017.1361879>.

6 Franzosi, 4.

7 I developed an analogous method for representing religious beliefs, which I discuss in Chapter 2. At the time I developed it, I was unaware of Franzosi’s work. It is comforting to know that I was (unwittingly) within an existing tradition.

8 Franzosi, “A Third Road to the Past?,” 13.

9 Franzosi, 14.

The stiff disdain in Elton's comment is a warning that will accompany this thesis, since much of what follows is an exploration of the territory that Franzosi has pioneered. Cooperative research and quantification are both central to my proposal. Elton is certainly right that "calling on arithmetic for aid" can hide superficiality in thought. He is also right that this approach had not born much fruit in 1967. But is that still the case? Franzosi thinks not. He closes his article with a quick glance over the Natural Language Processing (NLP) tools he uses, and states that

if we do not wake up soon to the new realities of big data, computer scientists will leave us behind, biting the dust in this road to knowledge. The good news is that computer scientists are no less eager to become bad historians, social scientists, linguists, or literary critics than we are to become bad computer scientists. [...] Which points to the increased need of collaboration among scholars from different disciplines in the world of big data.<sup>10</sup>

I heartily agree. Scholars working alone cannot successfully manage the specificities of many domains. Collaboration is essential to the knowledge project at this time—and so is thought. Analysis and cybernetics cannot replace the complementary work of synthesis; synthesis is suspect when it does not proceed from an analysis of all the available data.

To what end should the historian tend? A naive realist would say, the truth about what really happened in the past.<sup>11</sup> The historian should discover what happened and why. Such a view has been criticized from many quarters, and it has been pointed out that historians do not describe the real world like scientists do. There is debate on the issue. Behan McCullagh emphasizes that "it is reasonable to believe some historical statements correctly describe events which have actually occurred in the past, even though these cannot now be observed. [...] it is a realist position, because it affirms the reasonableness of faith in the reality of many of the events described by historians, which a non-realist would deny."<sup>12</sup>

According to McCullagh, belief in the reality of certain events, although they are unobservable to the scholar at present, is the best explanation of the features of the world that are known to us. Of course, this explanation may be changed if more information comes to

---

10 Franzosi, 15.

11 To understate it, the word is rich with history. A good introductory essay can be found at the Stanford Encyclopedia of Philosophy, Michael Glanzberg, "Truth," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Summer 2021 (Metaphysics Research Lab, Stanford University, 2021), <https://plato.stanford.edu/archives/sum2021/entries/truth/>.

12 C. Behan McCullagh, "Historical Realism," *Philosophy and Phenomenological Research* 40, no. 3 (March 1980): 421, <https://doi.org/10.2307/2106406>.

light. McCullagh again: “If an historical narrative accounts for certain facts believed to be true of the world, and does so better than any other explanation of them we can devise, that is a reason, if not always a very strong reason, for believing that the narrative correctly describes what actually happened to bring about the facts which it is designed to explain.”<sup>13</sup>

McCullagh recently engaged with Roy Harris, Emeritus professor of linguistics at Oxford, on the topic of historical realism. Harris holds that traditional theories of meaning are unsound, because the meaning of words changes in every new context in which they appear. He derives the idea from Saussure, who held that “when new words are introduced into a language, or old words are used in new ways, the language itself adjusts to accommodate the change, so it is wrong to suppose that languages continue relatively stable over time, with small adjustments”.<sup>14</sup> If this were the case, the historian could not properly translate an ancient text, since he has no way of accessing the meaning of the words it is composed of.

Furthermore, Harris proposes an “integrationist theory of meaning”, in which the meaning of a sentence depends on the “communication situation” in which it is produced. If we accept these premises, we will also have to accept McCullagh’s summary:

If the meaning of words is so context-dependent and variable, historians who have little knowledge of the context of the texts they study could not possibly know what they mean. And what chance have they of discovering that context if they cannot rely upon the written documents they are trying to understand?<sup>15</sup>

But in McCullagh’s estimation, it is not strictly true that the entire language changes meaning when a new word is added or an old word is modified. Communication remains possible, precisely because the language is mostly stable. Translation remains possible: even if some words are so special as to defy a one-to-one correspondence in another language, most words are not. And words do have a relation with the world, notwithstanding Saussure’s extrapolation from exceptions to affirm the contrary.<sup>16</sup> McCullagh explains that

---

13 McCullagh, “Historical Realism.”

14 C. Behan McCullagh, “Language and the Truth of History,” *History and Theory* 44, no. 3 (October 2005): 441–55, <https://doi.org/10.1111/j.1468-2303.2005.00335.x>.

15 McCullagh, 443.

16 The debate about the relation between word and world has a long and illustrious history. Plato dealt with the question in the *Sophist* and in *Parmenides*, with particular attention to the knowledge of ‘forms’ or abstract ideas, called universals in the medieval tradition. Many medieval scholars, such as Duns Scotus and William of Ockham, debated the realist and nominalist understanding of universals. The issue has been greatly expanded in recent decades, in part due to the advances in linguistics and neuroscience. In 1960, Quine

Dictionaries acknowledge that words often have slightly different meanings in different contexts, but these variations are regular, which is why dictionaries are able to list them. The mistake made by Harris and others who deny the importance of rules of language is that because context is always important in determining the interpretation of a text, they assume rules play no part in the process. But rules are essential to making verbal communication possible. Just try conversing in a language of whose rules you are entirely ignorant, and you soon discover how very little you can communicate.<sup>17</sup>

The historian can use words to learn about the past, and can translate old texts and arrive at some measure of knowledge. *Contra* Harris, McCullagh affirms that he can justify the “habit of regarding well-supported historical descriptions as conveying information about the real world”: “the materials that historians use as evidence have observable features that can be explained by means of a hypothesis about events in the past.”<sup>18</sup>

The hypothesis is an “argument to the best explanation”, analogous to the way a physicist produces a line of best fit to suggest that the messiness of the observed data can be described theoretically by a mathematically simple curve. These explanations are not always correct, because the evidence is not always complete: “there might be a better one they have not considered, and there might be more evidence that will cast a different complexion upon the historical events that interest them. But if the evidence in support of an explanatory hypothesis is strong, and there is no alternative hypothesis supported nearly as well, it is reasonable to believe it is probably true, at least for the time being.”<sup>19</sup>

McCullagh explores the problems that Harris raises with some sympathy. If our perceptions and our words are imprecise, under what conditions could they be called “true”? For McCullagh, “we would think them true if, when every possible perceptual experience of the

---

proposed that the study of epistemology must be informed by neurobiology, and many more recent scholars have followed this proposal. See Quine, W. V. O. (1960) “Word and object.” Cambridge, Mass.: MIT Press. More recently, the successes of artificial ‘neural networks’ in a variety of language processing tasks lead many to surmise that the connectionist view of neural representation at least within some high level of approximation can be considered a correct theory of how the brain represents the world. A classic from within linguistics is Eric H. Lenneberg, *Biological Foundations of Language* (John Wiley and Sons, 1967). An interesting idea regarding the concept of infinity, which represents a stumbling block for many recent linguistic theories, was recently presented in Ryan M. Nefdt, “Infinity and the Foundations of Linguistics,” *Synthese* 196, no. 5 (May 2019): 1671–1711, <https://doi.org/10.1007/s11229-017-1574-x>. See also Jozsef Andor, “Cognitive Grammar,” *Acta Linguistica Hungarica* 52, no. 4 (2005): 341–66. For a philosophical overview, see Gonzalo Rodriguez-Pereyra, “Nominalism in Metaphysics,” in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Summer 2019 (Metaphysics Research Lab, Stanford University, 2019), <https://plato.stanford.edu/archives/sum2019/entries/nominalism-metaphysics/>.

17 McCullagh, “Language and the Truth of History,” 447.

18 McCullagh, 452.

19 McCullagh, 453.

world was in, and when the implications of all possible perceptions had been considered, they remained the best explanation of our perception.”

However, this idea seems to bolster the skepticism in Harris’ position: it would be properly god-like to have access to a state in which “every possible experience of the world was in”. Humans can neither experience all states, nor can they remember all the states they do experience. The problem compounds when we note (as McCullagh does) that “historians build explanations upon explanations”.<sup>20</sup> The uncertainties of each layer multiply, and epistemological skepticism is their bitter fruit.

For McCullagh, Peircean pragmatism is the way out. “In order to act in the world, in our own and other’s interests, we must decide what the world is like, or to put it another way, which descriptions of the world to believe.”<sup>21</sup> I share this position. The fact that translation is approximate, and that historical narratives are ever-changing and never fixed for all time, is an important observation. However, I would not construct a relativist theory of truth upon this basis. Instead, I propose to treat the problem of ignorance as a problem that can be partially solved in time through the collection of new data. Thus, I shift attention away from the details of epistemology as applied to a limited data set, and toward the improvement of the data set itself. In this way, I allow for a multiplicity of historical methodologies to coexist and to continue to debate, while attempting to improve the capacity of the field to produce narratives and analyses that illuminate the past.

I now outline the epistemological choices that inform the rest of this work.

## **Connections Between Data**

The information gathered by a historian generally has a fragmented form.<sup>22</sup> A scholar finds fragments and hints, and must piece together an interpretation. Many of these fragments are textual primary and secondary sources. They also can be the material objects the archaeologist,

---

20 McCullagh, 454.

21 McCullagh, 454.

22 This affirmation is particularly true of archaeology, which deals with sherds and broken traditions. It also applies to history in that the connections between narratives, points of view, and material culture are often missing and must be reconstructed.

the historian's next-of-kin, discovers. Yet we wonder about those topics, and to some extent can reconstruct an understanding of them through the patient collection of details.

Each detail constitutes an “atom” of information, and is connected to other details in a complex web of connections. The fragments and their connections form a pattern that is not known *a priori*. This fact means that the structure of the connections should not be predetermined or influenced by the establishment of categories too early on, which might falsify the future development of links between data points along the lines of an inadequate ontology. At the same time, the fragments must be organized *somehow*, and a linear narrative (an article, speech, or book) is generally the desired endpoint of scholarly work. How can the complexity of the pattern of fragments be converted into a one-dimensional structure, a linear discourse?

As fragments are gathered and their connections are discovered, they form a complex pattern. Within the pattern, clusters appear: categories of similar ideas, objects, writings, experiences. If the establishment of categories happens organically as more data is included in the system, the structural connections between data points will actually be one reflection of the order in the system, and not an external structure imposed by the scholar. It is possible that biases and distortions enter the system, and for a variety of reasons some categories of information will be over-represented, and others will be scarcer in the study than they are in reality. The problem of bias is inescapable, but as the above discussion of McCullaugh and Harris indicates, we can aspire to imperfect, but real knowledge.

The brain of the scholar too is a patterned machine. The hundred-billion neurons in their brain connect and disconnect too, and form a continuously changing pattern of the world. The pattern of the material under study is repeated in the grey matter: cognitive scientists believe that the pattern formed in the brain is a representation of the external world, like a map that the human brain forms through experience.<sup>23</sup>

---

23 It is often postulated that experience is represented in the human brain by the connections between neurons. A good overview of the older debate about the distinction between body and mind, and the different ideas about how the brain represents the world, can be found in Patricia Smith Churchland and Terrence J. Sejnowski, “Neural Representation and Neural Computation,” *Philosophical Perspectives* 4 (1990): 343, <https://doi.org/10.2307/2214198>. At the time Churchland and Sejnowski wrote, artificial neural networks had only been recently described by Hinton et al in 1986. Later developments have continued to support the view that the brain represents information in its ‘connectome’ rather than in discrete physical locations. In

The brain-representation of the world is a form of knowledge, archived information, a map of realities and their connections. The historian needs both a rich representation of the world, which is to say a *lot* of data, and an exquisite sensitivity for narrative, the meaning-filled stories that lie embedded within the data. Narratives too are models, in the sense that they are a simplification of the facts that aim nonetheless to communicate a significant aspect of the facts.

If we expand our gaze from the single scholar to the broader community, we see that the narratives told by each scholar contribute to the information available to all the other scholars. Each narrative enriches the brain-representations of others, and sparks the search for new information and the generation of new corroborating or contrasting narratives. The scholarly community is thus like a super-brain in constant growth. New connections are formed, the representation grows more nuanced and complex, and once in a while the discovery of a piece of information connects domains of thought in such a convincing way that it feels like a perfectly-fitting puzzle piece. Other times, new information blasts a scholarly field wide open, and shows the limits of all the existing interpretations. All this movement and growth in complexity is an important source of joy.<sup>24</sup>

These considerations are generally important for any domain of knowledge, but they are particularly important for the work of an historian of the ancient world, since a great deal of relevant information is missing. The adequate structure of the data is often unknown, and unknowable, except through the reconstruction of the connections between data points that the scholar collects through the methods of historical research, including archaeology and the study of texts and artifacts. There is no external reference point that can be used as an ordering principle. The entire cosmos of meaning has to be reconstructed through the data themselves.<sup>25</sup>

---

particular, the success of artificial neural networks in a large variety of tasks that seemed quintessentially ‘human’ until they were produced by computers, bolsters the view that connectionist models of brain representation are closer to the truth than competing explanations. Interesting extensions and caveats to this view are presented in Anna M. Borghi et al., “Varieties of Abstract Concepts: Development, Use and Representation in the Brain,” *Philosophical Transactions of the Royal Society B: Biological Sciences* 373, no. 1752 (August 5, 2018): 20170121, <https://doi.org/10.1098/rstb.2017.0121>.

24 See the classic *Flow* by Mihaly Csikszentmihalyi (1990).

25 This does not mean that data is considered to be purely “objective”: rather, it means that the data is all we have. Of course it is influenced by the imprecision of the methods used, the choice to attend to this rather than that, the history of the observer, and much more.

It is a fascinating, seemingly impossible project, like lifting one’s self out of the bog by one’s own hair.<sup>26</sup> Throughout these pages, I will explore aspects of this project in turn. The characteristics of knowledge I emphasized above could lead to a “relativistic” understanding, in which each point of view is inescapably unique and ultimately incommunicable. Through the study contained in these pages, I wish to pay serious attention to the multiplicity of points of view, the “micro-narratives” that proliferate around any serious interpretive question. Other aspects of my proposal might seem part of the old search for “objective truth”. As we proceed, I hope it will become clear that my attention to structure does not claim to arrive at universality, nor is it capable of the impossible task of a “view from nowhere”. Instead, I am convinced that attention to the multiplicity of interpretations can, with adequate scaffolding and perhaps a technological prosthetic, aim to improve the interpretive enterprise by taking a greater quantity of information and micro-narratives into account. And I am convinced that while human knowledge may inescapably be only the construction of models, and not the sure knowledge of immutable Truth, it is also possible to *improve* our models and more closely approximate the realities they represent. The process of improvement can be colloquially described as “moving closer to the truth” without thereby intending that it is possible to arrive at that destination once and for all—but this is not my argument. My central focus is interdisciplinary, and aims to produce better models by taking into account *more* data and by representing *invariant structures* within that data, with the help of mathematical techniques.

## Inner-referential Interpretation

A recent article by Johanna Drucker entitled “Modeling Interpretation”<sup>27</sup> showcases several projects that aim in various ways to collect data and interpret it without reference to external sources of truth. The first, which Drucker calls temporal modeling, attempts to take into account the multiple values time has for different observers, without reference to an external, common reality. “Our goal (never quite fully achieved) was to construct a system that was not based on a container model (time as a preexisting framework into which events or incidents

---

<sup>26</sup> With a nod to Baron von Munchhausen.

<sup>27</sup> In Martin Brückner, Sandy Isenstadt, and Sarah Wasserman, eds., *Modelwork: The Material Culture of Making and Knowing* (Minneapolis London: University of Minnesota Press, 2021), 227–55.



were put) but based on relationality.<sup>28</sup> The proportions and scales of the temporal model would emerge as events and references were put into relation with each other.”<sup>29</sup>

Drucker is certainly correct in challenging the idea of time as a “container”, which even in the natural sciences is a Newtonian conception that has been superseded by more nuanced depictions.<sup>30</sup> Her project supports “a meta-reflection on the work of interpretation while also offering the means to model interpretatively.”<sup>31</sup> In other words, instead of positing an external point of reference such as a “second” or an “hour”, she meant to construct the concept of time entirely through internal references between texts.

It would be interesting to know more in detail what elements in her view remain unrepresented in this model—why she says their goal was never fully achieved. I suspect that it might have to do with the fact that in order to compare two quantities, an external metric is usually necessary. In the absence of a direct, physical comparison between two objects (often impossible), we need more information. An external reference, such as inches or centimeters, does permit comparison.<sup>32</sup> To be sure, it does do so not without imposing strictures of its own, such as the tendency to express things in round numbers, whatever system one uses. It might also impose value judgments about the nature of space, which might need to be specified, and perhaps there would also be an anticolonialist critique lurking somewhere, along the lines of Drucker’s point about the Celsius scale not taking into account the experience of cultures that had never encountered freezing temperatures until recently.<sup>33</sup> But it might be the best one can do.

---

28 The point is well taken. In fact, the “quantitative” transcription of data that was originally humanistic and text based need not be merely quantitative in the sense of linear scales. It could also take into account relation, and overall “shape” or topology, as other mathematical elements that describe the content.

29 In Brückner, Isenstadt, and Wasserman, *Modelwork*, 235.

30 In physical space, time is intimately linked to gravity and acceleration in general (as Drucker points out); in biological space, a animal lifetime can be closely correlated to the total number of heart-beats. The psychological experience of time is a still more complex issue.

31 In Brückner, Isenstadt, and Wasserman, 239.

32 See also the profound critique of inner-referential systems understood as “encodings” carried out by Mark Bickhard (1995). “If representational contents are carried or constituted only by encodings, then how can we ever check the accuracy of our representations? To check their accuracy would require that we have some epistemic access to the world that is being represented against which we can then compare our encodings, but, by the encodingism assumption, the only epistemic access to the world that we have is through those encodings themselves. Thus, any attempt to check them is circularly impotent—the encodings would be being checked against themselves.” (19-20). This is an important point, which we take up again in Chapter 4. For the moment, let this note suffice.

Drucker emphasizes what she calls a “nonrepresentational approach”. “A nonrepresentational approach does not preclude the use of an image to depict or figure something; it simply suggests that an image cannot represent something in a relation of equivalence.”<sup>34</sup> Of course, “it is a truism that a map is not equivalent to a territory.” But consider a space as inhabited by a “small child, a threatened woman, a preoccupied man” and the meaning of the details of the space they inhabit change radically. That meaning is not represented within the map: it must be reached through interpretation. This is a point she insists upon: Drucker affirms that “within the digital humanities, this interpretative activity has not found explicit expression in graphical (or other) form, nor have the models (such as data models or algorithmic ones) by which interpretation is formulated been explicitly called to attention.”<sup>35</sup>

Drucker’s statement that “there is no unmediated relation to phenomena” sets the scene for her various proposals for modeling interpretation. “Interpretation is a constitutive act, one that makes the object of investigation through engagement,”<sup>36</sup> she claims. But “the digital humanities, in part by its concessions to the mechanistic and highly formalized, procedural logics of computational methods, has not integrated these theoretical approaches into its techniques. [...] What if, instead, the digital humanities were to take on the challenge of modeling interpretation instead of giving up on the possibility of formulating computational methods on a non-empirical foundation?”

The reason this project has not been accomplished probably lies in the incredible complexity of the endeavor. “The challenge is to understand how to model interpretation

---

33 The Centigrade temperature scale is constructed on the basis of an experience (boiling and freezing water) which some peoples had never had before refrigeration and travel made it available to them. This example is susceptible to further critique, since the experience described is also as “objective” an event as can be found. Emphasizing the cultural contextuality of this measurement scale might be pushing it a bit. Water does not freeze or boil at arbitrary or culturally-determined temperatures.

34 This quotation and all subsequent quotations on this page, Brückner, Isenstadt, and Wasserman, *Modelwork*, 227.

35 In Brückner, Isenstadt, and Wasserman, 228.

36 This quotation and the rest in this paragraph, in Brückner, Isenstadt, and Wasserman, 231.

computationally,”<sup>37</sup> writes Drucker. Is it possible to represent quality through quantity?<sup>38</sup> If so, how?

Another of the projects developed by Drucker’s team, called Interpret, “was intended to integrate the activities of scholarly research—bibliographical trails and links, text analysis, visualization, multiple (many) documents, and the layering of interpretative frameworks onto the underlying evidence—and vice versa.”<sup>39</sup> If only we could retain evidence of every “trail”, if only we could layer evidence and interpretation together: then! We would have the mind of God.<sup>40</sup>

Like we saw in the debate between Harris and McCullagh, god-like omnipotence lies in the background of Drucker’s attention to critical theory. She affirms that “Critical theory [...] seems to get jettisoned in the implementation and use of digital environments that provide a sense of omnipotence—or, at the very least, potency.”<sup>41</sup> It is true that the speed and ease of digital manipulation seems to conjure images of omnipotence, as if the user were a god-like figure, capable of seeing the overall structure of things from a disinterested, neutral point of view. This is of course an illusion, and Drucker does well to warn us of the danger. No matter how extensive our collection of data points, it is only a small subset of the information that could be discovered in the future. The data we see—all the patterns we have—are like a few branches illuminated by a campfire, while out beyond the flickering light of our intelligence lurk the unexplored forests and trackless prairies of the Being we have not yet encountered. In the semi-darkness, we can tell stories and make models, for we may discover invariant features as well as culturally-determined structures: -etic and -emic elements, as they can be called in shorthand.<sup>42</sup>

The project of this thesis thus lies within the furrow attempted by Drucker in her “inner-referential” modeling. I will pursue an approach that she did not explore directly, although she

---

37 In Brückner, Isenstadt, and Wasserman, 231.

38 Pirsig’s *Zen and the Art of Motorcycle Maintenance* explores this theme.

39 In Brückner, Isenstadt, and Wasserman, Modelwork, 241.

40 The oblique reference to “trails” conjures the memory of Vannevar Bush and his memex, which remains a constant dream in the minds of digital humanists.

41 In Brückner, Isenstadt, and Wasserman, Modelwork, 248.

42 See a short introduction to this topic in Buccellati, G., “On (e)-tic and -emic”, *Backdirt. Newsletter of the Cotsen Institute of Archaeology* (Winter 2006), pp. 12-13. I return to this crucial issue in Chapter 6.

partially intuited my project in her reference to “trails” mentioned above. As Buccellati put it in his *Critique of Archaeological Reason*, “The archaeological effort should be viewed in the first place as a cognitive structural whole that is self-contained and must be understood and described on its own terms, without reference to anything external to it.”<sup>43</sup> In this thesis, I aim to produce such an “inner-referential”, self-contained structural whole as a model of historiography.

## Split Fragments, Recomposed Wholes

If knowledge involves the pattern of internal connections between data, represented in the human brain as a pattern of connections between neurons, could it be mimicked by a machine that carries out the same operation? Many have wondered if a machine could be built that would contain a rich representation of the world as fragments connected together. Others have affirmed that no process of fragmentation would be able to reproduce the “wholes” that we indicate with words like “thinking” or “intuition”.<sup>44</sup> For them, the construction of narratives lies forever beyond what a fragment-containing machine could do.<sup>45</sup>

The question lies at the threshold of the spiritual. One of the oldest ways of explaining the world is mythological, in which humans postulate the existence of powers outside the physical everyday world they inhabit. This “outside”, a world above or beyond the one we inhabit in our mundane living, was seen as controlling and causing this world. But does the “outside” exist, and is there an Archimedian fulcrum from which to exert effective pressure, or is everything we

---

43 Buccellati, *Critique of Archaeological Reason*, 17.

44 The issue has been explored by Mark Bickhard. See one overview in Bickhard, Mark H. (2016). “How to operationalize a person” *New Ideas in Psychology*. More in detail, Bickhard’s *Foundational Issues in Artificial Intelligence and Cognitive Science* (available online at <https://www.lehigh.edu/~mh0/AIFull.pdf>) explores in depth the logical and ontological reasons why, in his estimation, the usual understanding of representation as “encodingism” is inadequate and must be replaced with “interactionism”.

45 The recent development of generative AI machines based on Large Language Models (LLMs) like ChatGPT has raised this issue with new urgency. The output of these machines certainly *seems* to be a coherent narrative created by the collation of fragments. Correctly interpreting what these products are doing remains a work in progress. For some relevant recent discussion, see Moro, Greco, Cappa “Large languages, impossible languages, and human brains” (2023), and Piantadosi “Modern language models refute Chomsky’s approach to language” (2023).

know just the dimly lit trees around a fire, with countless other shapes waiting to be revealed by a higher blaze?<sup>46</sup>

In the last few centuries, it has become common to think that there is actually no “outside” and no immutable “nature”, and that intrahistorical relations are determined not by an external structure of reality, but exclusively by the actions and interplay of the historical agents themselves.<sup>47</sup> And in scholarly pursuits, it is broadly accepted that it is correct to proceed *as if* there were no outside. Science does not use the hypothesis of gods or angels pushing the stars around, even if it does not have the right to negate their existence either. The question of the existence of “something” that is not a thing, that is not made of matter, by definition does not belong to the domain of physics, and cannot be posed as a properly scientific question. Furthermore, the existence of ways of action of matter and energy that we do not yet know about is a basic postulate in research. There are always more things under heaven than are contemplated in our (current) philosophy.<sup>48</sup>

Natural science proceeds as if there were no outside, and attempts to connect the data together in a coherent picture that fits the phenomena, building a model without reference to forces that cannot be queried or known.<sup>49</sup> History too attempts to understand events and their

---

46 Plato was pointing at a truly profound mystery when he suggested that knowledge is remembering, and that we cannot see something except insofar as it has the “look” of something (an *eidōs*, which we often translate as “idea”)—and so in that view, it is the “outside” that informs the “inside”.

47 We are well accustomed to the bracing excitement and the solitude of the idea that we are alone in the cosmos, solely responsible for understanding it. At the same time, the traditional perspective continues to be attractive. It does seem intuitively true that reality has a structure, that this structure can be known, and that we fail to respect it at our peril. It does seem true, as George Grant pointed out, that if we think that all the horizons that gave meaning to our actions are just a human invention, a darkness falls over our minds and we are strangely unable to act. See George Grant’s interpretation in “Time as History” (2000). For Nietzsche, this was properly the role of the super-man: the ability to dwell in a world of pure possibility without any sacred outside, and to order reality according to his own will. For others, this has been a demonic call, alluded to in Goethe’s *Faust*, Mann’s *Magic Mountain*, and Bulgakov’s *Master and Margaret*. Each of these works chronicle a period of power-drunken ecstasy, followed by disaster. The plot was followed not only in fiction, but also in the politics of the twentieth century, as well as in other periods. Johann Huizinga famously described the end of the Middle Ages in terms that present much the same dynamics—the movement from a society that knew what it was and what the rules of life were, to a deeply anguished and uncertain age, which was also the crucible from which enormous human achievements arose.

48 With apologies to Shakespeare and his Hamlet (Act I, scene 5).

49 It runs into significant difficulty in the liminal spaces, for instance when it asks what consciousness is. Is it an “emergent property” of a highly complex system? Is it an illusion? Is it a form of computation we do not yet understand? The difficulty in responding to these questions from within the scientific paradigm does not necessarily imply that there is no spirit in the human, capable of producing cause-less free actions, nor does it mean that there is no spirit in the world, somehow transcendent beyond what is physical. The hypothesis of

causes and effects, without reference to gods or *res cogitans* except insofar as they are elements of the beliefs of the people living at the time under examination. It is essentially a choice, and perhaps a humble choice, to concentrate on what presumably can be known by all humans who have access to the underlying data.

But the quantity of data that must be interpreted is astronomical. Could the endeavor be aided by computational methods to simplify the mass of data, and help find patterns within it? Our ancestors had mythological interpretations to help simplify their understanding of complex dynamics: we too need practical heuristics to reduce the dimensions of the data we have to process.<sup>50</sup>

Earlier in this chapter, I alluded to the ‘-etic’ and ‘-emic’ distinction. This alternative can provide us with a helpful framework within which to situate our topic. The two terms allude to the distinction between phonetics and phonemics in linguistics, although in current usage their reference is broader.<sup>51</sup> Buccellati affirms that “the basic underlying concept is the distinction between an open and a closed system, where -etic refers to the first, and -emic to the second.”<sup>52</sup> The open -etic information is extracultural, and can be studied from a purely structural point of view. In contrast, -emic information is intracultural, and refers to the meaning of data that from an -etic point of view can be categorized, but not interpreted. To take up the Platonic allusion of this section title, the split fragments are recomposed into provisory and partial wholes: -etic categories. Then, from within their culture, the categories can be read, interpreted as having meaning.

My primary focus will be on methods that can surface -etic structures and patterns within historiographical information. This is not at all meant as a dismissal of the value and

---

the spirit is simply not a scientific topic.

50 Today, critical readings of history emphasize the degree to which one’s point of view affects the interpretation one gives to events. This is important and laudable; it also creates an infinite web of different interpretations, literally as many as there are interpreters. Sometimes there are fashionable sets of interpretations, such as the idea that European culture was the pinnacle of human society, and should be exported to the “little black faces” in Ethiopia by fascist Italy, or the later idea that European culture is among the main oppressors of human society, and many statues of its “great” figures should be eliminated from public spaces by forcible removal. These fashionable interpretations often last a brief season before being replaced by another, sometimes opposite, dogma.

51 Buccellati, G., “On -etic and -emic”, *Backdirt* 2006, 12.

52 Buccellati, G., “On -etic and -emic”, *Backdirt* 2006, 12.

importance of intracultural interpretation. Rather, it is the choice to concentrate my contribution on the first steps of the long road from fragments to wholes. I would like to run down that road more quickly, and my research aims to provide a tool for doing so. The aim remains to characterize extracultural structures precisely, with the expectation that this will help others refine their intracultural interpretations.

It is the choice to focus on what can be said from a position exclusively “within” the data. In a sense, it is thus the choice to emphasize the certainty and poverty of mathematics, before enjoying the wild and rich ambiguity of language. Though our approach be poor, one can find some comfort in an ancient idea attributed to Augustine of Hippo: if it is truly stated that a leaf fell from a tree, not even God can disagree.

In the next chapter, I will explore some major elements of “digital humanism”, the application of computers to humanistic work. These first examples are as much an exploration of the limits of computation as applied to humanistic disciplines as they are a celebration of its power. Using the lens of computation, I will seek for structures and underlying patterns in two specific directions: augmenting human memory, and communicating the details of that ‘prosthetic memory’ across a group of scholars.

## Chapter 2 – Elements of Digital Humanism<sup>53</sup>

This chapter is about two central elements of my approach to model-making in digital humanities: attention to structures, and automated natural language processing.

The intuition that structures are a powerful way to study cultural realities has left a profound mark on linguistics (de Saussure<sup>54</sup>, Jakobson<sup>55</sup>), anthropology (Lévi-Strauss<sup>56</sup>), and several other fields, including archaeology (Buccellati<sup>57</sup>). In my usage, the word ‘structure’ refers to the invariant patterns that can be found within information, and that can be described without reference to information outside the system.

If we accept as our starting point the postulate that intrahistorical relations are to be studied exclusively on the basis of the actions and interplay of historical agents themselves, we find ourselves constrained to the inside of a structure. Even if it does have an “outside”, we have chosen to suspend judgment about such features.

Edwin Abbot’s mathematical romance *Flatland* offers a useful analogy.<sup>58</sup> The inhabitants of “Flatland” sometimes observe higher-dimensional objects passing through their world, but they see them only within the two dimensions allowed in their space. A three-dimensional sphere passing through Flatland is seen first as a point, then it expands to a circle, then contracts down

---

53 The plural “humanities” instead of “humanism” is arguably a distinct view of the nature of the topic under discussion. Giorgio Buccellati proposes that “digital humanities” be used to refer to the technical dimension proper to data and its manipulation, whereas “humanism” and “digital humanism” aim specifically to “appropriate past human experience”. (See *Critique of Archaeological Reason*, p. 193. I use the two terms interchangeably here, since I name the crux of that distinction in other ways, especially in the relationship between partial and competing narratives with a postulated (though perhaps ultimately unknowable) unified view of the truth.

54 Ferdinand de Saussure, *Course in General Linguistics*, translated by W. Baskin, New York: Philosophical Library (1916).

55 Roman Jakobson and Linda R. Waugh, *The Sound Shape of Language*, de Gruyter (2002).

56 Claude Lévi-Strauss, “Structuralism and Ecology”, in *Graduate Faculty Philosophy Journal* 7:2 (1978), pp. 153-78.

57 Giorgio Buccellati, *A Critique of Archaeological Reason*, Cambridge University Press (2017).

58 Edwin Abbott, *Flatland*, Seeley & Co. (1884).



to a point again before disappearing: the plane presents a “slice” of the sphere moving through it.

An inhabitant of Flatland would need to display special powers of intuition to grasp that the locally experienced phenomenon (the changing circle) was actually a higher-dimensional reality projected onto the plane. We do not know if the Flatlanders were actually capable of such imagination. However, and crucially for our subject, humans *have* demonstrated just such feats of intuition many times in human history. Galileo correctly intuited the spherical shape of the moon and correctly deduced the orbit of Jupiter’s moons, even though they appeared two-dimensional in the frame of his telescope.

We too can move from a low-dimensional shadow to the intuition of the high-dimensional reality of structures. My choice to follow Drucker in the search for inner-referential methods that permit the scholar to characterize the structure without reference to external realities restricts the interpretative possibilities, but is justified by the desire to create a firm foundation for hermeneutics. Once the bare possibilities of inner-referential analysis are exhausted, less certain and more fascinating approaches will have their place. Within this thesis, I will focus on ways to characterize the data and its structure, in preparation to deploying other, more complex ways to approach interpretation: I will describe several methods for characterizing and communicating structures, leading from “atoms” of information, to their connections, and finally to an overall interpretation of the meaning of the information thus gathered and ordered.

In order to make use of structural thinking, we will need a way to treat natural language text as an element of a structure. The importance of this aspect became clear to me when I realized how much more information is available than any scholar can actually read and remember. Scholars vary widely in their ability to remain up-to-date in their field, but certainly no scholar is able to read everything that is published yearly even in a narrowly defined domain. For this reason, it would be useful for scholars to have an additional tool capable of performing a first level of dimension reduction—if, of course, such a tool were actually capable of summarizing and surfacing relevant information. Such a task presupposes the invention of algorithmic ways of characterizing text as a mathematical object, in order to take advantage of the processing power of computers.

## Natural Language Processing

Much work has already been done on this issue, and much more is ongoing. Here, I concentrate on methods that represent text as a mathematical object.<sup>59</sup> On the basis of such methods, we can characterize groups of texts as structures, and carry out the inner-referential analysis we have set out to perform.

In this section, I begin by discussing a digital method for translating text into numerical data, “word embeddings”, as they are known.<sup>60</sup> This method can be described as the creation of a numerical representation of textual information. My hope is that the speed and scale advantages afforded by computer processing of text will provide a helpful addition to other, older methods. I will cover the background of this method and attend to several important critiques of it. Then, once aware of its strengths and limitations, I will employ it in several specific examples in which it is used in an attempt to decide between competing historical interpretations of Mesopotamian culture.

The underpinnings of the most powerful methods of automated natural language processing were first proposed by Rumelhart, Hinton, and Williams in a 1986 article.<sup>61</sup> The idea was simple and revolutionary: a network of abstractions in a computer program representing artificial “neurons” could be constructed to receive inputs and generate outputs with varying probabilities. A message passed through the system would produce an output, and that output could then be checked against a known correct answer and an error signal “backpropagated” into the weighting of the “neurons” in the program, thereby altering the output of the next iteration. By recursively passing messages through the system and checking for a desired output, the weightings could be successively tuned, until they form a representation of whatever process the user wishes to encode. In 1986, the idea ran aground because computers were not sufficiently powerful to carry out the computation necessary.

---

59 To wit: as a vector in the high-dimensional space of a language model.

60 The idea that “a word is characterized by the company it keeps” was proposed in 1957. See John Rupert Firth, “A synopsis of linguistic theory 1930–1955”. *Studies in Linguistic Analysis*: 1–32. A review of the different ways word embeddings may be generated is found in Lavelli, Alberto; Sebastiani, Fabrizio; Zanoli, Roberto (2004). “Distributional term representations: an experimental comparison.” *13th ACM International Conference on Information and Knowledge Management*. pp. 615–624.

61 David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams, “Learning Representations by Back-Propagating Errors,” *Nature* 323, no. 6088 (October 1986): 533–36, <https://doi.org/10.1038/323533a0>.

But fourteen years later, the average desktop computer was as powerful as the supercomputers of 1986. An important step forward became possible: Paccanaro and Hinton applied the idea of backpropagation in a neural network to “word embeddings” in 2000, and trained the network to represent natural language as a set of probabilities, and words as vectors.<sup>62</sup> At first, their approach treated language as a “bag of words”, and did not encode information about word order or context but just text frequency.<sup>63</sup>

Later, it was refined to generate text: for each word in the training set, the program would predict what the next word should be. Since the next word in the training set is known, an error signal could be generated to modify the neural network and improve the prediction quality as training proceeded. The training sets can be quite large, and over time the network would settle into a pattern capable of predicting the next word with good probability.

This work was developed and improved by a team at Google. In 2013, Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean published an article describing their invention of a new mechanism for computing vector representations of words, known as Word2Vec.<sup>64</sup> Using the powerful computers at Google, they trained a neural network on a very large corpus of English text, essentially the billions of words that comprise the entire English language internet, in order to represent all English words in a 640-dimensional vector space.<sup>65</sup> By converting words

---

62 Alberto Paccanaro and Geoffrey E. Hinton, “Learning Distributed Representations of Relational Data Using Linear Relational Embedding,” in *Neural Nets WIRN Vietri-01*, ed. Roberto Tagliaferri and Maria Marinaro, Perspectives in Neural Computing (London: Springer London, 2002), 134–43, [https://doi.org/10.1007/978-1-4471-0219-9\\_12](https://doi.org/10.1007/978-1-4471-0219-9_12).

63 Stephen Robertson, “Understanding Inverse Document Frequency: On Theoretical Arguments for IDF,” *Journal of Documentation* 60, no. 5 (October 2004): 503–20, <https://doi.org/10.1108/00220410410560582>. They used the tf-idf approach, short for term frequency–inverse document frequency. This approach characterizes the importance of single words in a text by normalizing them by their frequency in an overall corpus. It is a curious mix of obviously insufficient philosophical and linguistic theoretical underpinnings—which gives surprisingly good results.

64 Tomas Mikolov et al., “Efficient Estimation of Word Representations in Vector Space” (arXiv, September 6, 2013), <http://arxiv.org/abs/1301.3781>. See also Mikolov, T. *et al.*, “Distributed Representations of Words and Phrases and their Compositionality”, proceedings of NeurIPS 2013, [proceedings.neurips.cc/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf](http://proceedings.neurips.cc/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf)

65 Clearly, the quality of the language model will depend on the training corpus used. Several large corpora are publicly available, and tech companies are also creating proprietary text sets to finetune their models for various use cases. For instance, there are corpora comprised entirely of books; one language model has been trained on all the medical research articles contained on PubMed, others use the web articles that have received at least three upvotes on Reddit...

into vectors, they found, a “word arithmetic” became possible, which retained the semantic content of the words and permit computerizing a series of operations.

Mikolov et al. give a simple example of their results in their paper: “Using a word offset technique where simple algebraic operations are performed on the word vectors, it was shown for example that vector (“King”) minus vector (“Man”) plus vector (“Woman”) results in a vector that is closest to the vector representation of the word Queen.”<sup>66</sup> Further refinements of this idea have led to vector representations of full sentences (Sent2Vec), and even entire documents (Doc2Vec).<sup>67</sup>

Since then, many other models of language have been built. Transformer-based models are able to account for the context of words, and thus identify the disparate meanings of homophones as well as the similarity of terms like “car” and “automobile”, that share no typographical features beyond the letter ‘a’, yet mean very nearly the same thing. These models now power systems that produce automatic translations, perform part-of-speech parsing, and suggest relevant information to the users of online shopping systems. Vector representations exhibit some useful properties and can be used to generate recommendations for similar topics or articles, such as what Amazon and Netflix do when they predict what a user might prefer to see next and what scholarly sites such as PubMed do when they suggest articles similar to the ones their users have already accessed.

Even more recently, language models have been used to create programs like ChatGPT that perform many tasks like writing and summarizing articles in a way that is uncannily close to what a human can do. Clearly, language models capture something essential about language. But before entrusting our research to them we must look more closely at the details of what they can and cannot do.

## Neural Correlates and Language Models

If vector representations are an effective way to translate human speech into mathematical space, some researchers have hypothesized that analogous features might be observable in the

---

<sup>66</sup> Mikolov, T. *et al.*, p. 2.

<sup>67</sup> The literature on this topic is enormous. For an introduction, see Gidi Schperber’s article “A Gentle Introduction to Doc2Vec” at [medium.com/wisio/a-gentle-introduction-to-doc2vec-db3e8c0cce5e](https://medium.com/wisio/a-gentle-introduction-to-doc2vec-db3e8c0cce5e).

patterns of neural activation as measured by an fMRI device. If so, we would be justified in thinking that language models built upon vector representations are indeed capable of capturing part of the essence of human language, and therefore can justifiably be used in algorithmic processing of large amounts of text.

A groundbreaking study was published in 2008 in *Science* by Mitchell, *et. al.* The team built a computational model of language based on a trillion-word corpus and observed fMRI data for 60 nouns which were viewed by nine college-age participants. The fMRI patterns corresponding to the 60 nouns were averaged across the participants and used in conjunction with the language model to predict the activation patterns of other words not included in the corpus. The researchers were able to predict the fMRI images for stimulus words to within a meaningful accuracy, and based on this correspondence they believe that they establish “a direct predictive relationship between the statistics of word co-occurrence in text and the neural activation associated with thinking about word meanings.” Furthermore, their work suggests that “the neural encodings that represent concrete objects are at least partly shared across individuals.” This fact supports the hypothesis that semantic features of natural language can indeed be encoded as a vector space or as the topography of an fMRI scan.<sup>68</sup>

Other recent research has continued along these lines. For example, Caucheteux and King (2022) “compared a variety of deep language models to identify the computational principles that lead them to generate brain-like representations of sentences.” Like Mitchell *et al.*, their research supports the affirmation of a similarity between algorithms such as Word2Vec and the ability of the brain to predict words from context. According to Caucheteux and King, “modern language algorithms partially converge towards brain-like solutions.”<sup>69</sup>

Kaiser *et al.* (2022) built on Mitchell’s work to predict fMRI patterns related to abstract concepts and thereby inform theories about conceptual representations in the human brain.<sup>70</sup>

---

68 Tom M. Mitchell et al., “Predicting Human Brain Activity Associated with the Meanings of Nouns,” *Science* 320, no. 5880 (May 30, 2008): 1191–95, <https://doi.org/10.1126/science.1152876>.

69 Charlotte Caucheteux and Jean-Rémi King, “Brains and Algorithms Partially Converge in Natural Language Processing,” *Communications Biology* 5, no. 1 (February 16, 2022): 134, <https://doi.org/10.1038/s42003-022-03036-1>.

70 Daniel Kaiser, Arthur M. Jacobs, and Radoslaw M. Cichy, “Modelling Brain Representations of Abstract Concepts,” ed. Leyla Isik, *PLOS Computational Biology* 18, no. 2 (February 4, 2022): e1009837, <https://doi.org/10.1371/journal.pcbi.1009837>.

Fernandino *et al.* also investigate the information structure underlying the neural representation of concepts. They find that “lexical semantic information can be reliably decoded from a wide range of heteromodal cortical areas in the frontal, parietal and temporal cortex.”<sup>71</sup>

Schrimpf *et al.* (2021) applied state-of-the-art “transformer” models to predict variants in neural responses to language and found that such models strongly correlate with the next-word prediction task. This result further bolsters the case that essential features of human language use are indeed expressed in the language models using vector-space embeddings.<sup>72</sup>

Another recent article, published in June 2022 by Millet *et al.* expands on these results with a neural network called “Wav2Vec”.<sup>73</sup> Instead of encoding words as vectors, the team based their vector encodings on the auditory waveform of spoken text. Using fMRI techniques, they were able to trace the processing of auditory input in a fairly large cohort (n=412) of English, French and Mandarin speakers. Interestingly, they gave the Wav2Vec algorithm access to only about 600 hours of speech data, a quantity comparable to what a human child encounters while learning language. Their results were spectacular: the algorithm learns “brain-like representations” in the time allowed, its “functional hierarchy aligns with the cortical hierarchy of speech processing”, it “learns sound-generic, speech-specific, and language-specific representations similar to those of the prefrontal and temporal cortices”, and their results are robust when compared with a second large cohort (n=386) of participants. Even more

---

71 Leonardo Fernandino *et al.*, “Decoding the Information Structure Underlying the Neural Representation of Concepts,” *Proceedings of the National Academy of Sciences* 119, no. 6 (February 8, 2022): e2108091119, <https://doi.org/10.1073/pnas.2108091119>.

72 Martin Schrimpf *et al.*, “The Neural Architecture of Language: Integrative Modeling Converges on Predictive Processing,” *Proceedings of the National Academy of Sciences* 118, no. 45 (November 9, 2021): e2105646118, <https://doi.org/10.1073/pnas.2105646118>. A natural application of these principles aims to realize a long-standing dream in human-computer interfacing, namely the ability to interpret “inner speech”, and control a computer by thinking the commands one wishes to execute. Some exciting work on this topic, which has immediate application to apparently comatose patients—but who are actually awake, just completely paralyzed—has been carried out in recent years. One of these, Nieto *et al.* (2022), developed an open-access dataset by recording ten subjects with an EEG-based brain-computer interface. Nicolás Nieto *et al.*, “Thinking out Loud, an Open-Access EEG-Based BCI Dataset for Inner Speech Recognition,” *Scientific Data* 9, no. 1 (February 14, 2022): 52, <https://doi.org/10.1038/s41597-022-01147-2>.

73 Juliette Millet *et al.*, “Toward a Realistic Model of Speech Processing in the Brain with Self-Supervised Learning” (arXiv, March 20, 2023), <http://arxiv.org/abs/2206.01685>.

surprisingly, two researchers in Japan recently showed that latent diffusion models could reconstruct images seen by test subjects only using their brain activity as measured by fMRI.<sup>74</sup>

This is just a survey of some of the most recent and interesting results that support the hypothesis that language models are indeed a reasonable representation of natural language. But there are important criticisms as well.

## Critiques of Statistical Language Models: Where Word Vectors Fail

W. H. Auden presciently described, and belittled, the “word algebra” proposed by Mikolov *et al.* In his foreword to Barfield’s *History In English Words* (1967), Auden thought that Barfield’s historical approach to understanding language is “much more likely to be effective than the approach of the linguistic analysts. The latter seem to believe that, by a process of ‘demythologizing’ and disinfecting, it should be possible to create a language in which, as in algebra, meanings would be unequivocal and misunderstanding impossible.”<sup>75</sup>

Although there are similarities between vector representations of natural language and neural activity, the dissimilarities are important as well. To give one obvious example, current language models are mostly unable to distinguish negation. They usually assign very similar vectors to two phrases whose semantic meaning is diametrically opposed. A second important point is that as far as we know at present, the operations of the brain are not like the probabilistic operations of a neural network: neurons do not emit continuous values, but discrete spikes. Notwithstanding the suggestive results surveyed above, we are still far from having resolved the question of the neural correlate of language.

Orthographically identical words whose meaning can be understood only from their context are also imperfectly modeled in current systems. It is true that transformer models attempt to take context into greater regard, but distinguishing ‘saw’, the object, from the past tense of the verb ‘to see’, which any English reader would do without difficulty, is not so easy

---

74 Yu Takagi and Shinji Nishimoto, “High-Resolution Image Reconstruction with Latent Diffusion Models from Human Brain Activity,” preprint (Neuroscience, November 21, 2022), <https://doi.org/10.1101/2022.11.18.517004>.

75 Owen Barfield and Wystan H. Auden, *History in English Words* (Gt. Barrington, MA: Lindisfarne Books, 2002).

for a model. Thus, Auden's point. Human language is not unequivocal, and whatever translation into mathematics it affords will also not be unequivocal.

The reason for this, according to Auden and Barfield, is that human language is mythological and metaphorical by nature. In their view, by its very essence it refers to an 'outside', and cannot be entirely understood in terms of 'inner-referentiality'.

Barfield writes that "language has preserved for us the inner, living history of man's soul. It reveals the evolution of consciousness."<sup>76</sup> Barfield documents this affirmation over several books, including *Poetic Diction*, *History In English Words*, and *Speaker's Meaning*. In his theory of poetics, he affirms that language continually changes and presents a poetic character only at the moment of its changing. Thus, poets must constantly reinvent the language. Barfield's examples make his reader perceive that the richness of language indeed lies far beyond what can be translated into vectors or mathematical spaces.

Perhaps, as Auden pointed out, the ground-level of language is like a code, empty variables that stand for needs common to all human beings, which can easily be translated between all languages, like in a tourist's phrasebook. But poetry, he believed, cannot even be approximately translated. "A poet, one might say, is someone who tries to give an experience its Proper Name."<sup>77</sup> As historians, we can agree that if we are to give past civilizations an honest and sensitive reading, we must strive to describe them with Proper Names, and not only as instances of abstractions like civilization, the city, law, religion, and language.

We do well to listen to Auden and Barfield even though I have announced my intention to suspend, for the moment, the specificities of Proper Names and the richness of poetry. I must also announce my intention to distinguish between a "word algebra" that pretends to establish univocal meanings for words, and a "word algebra" such as Word2Vec's statistical interpolation. In the former case, the failure of formalization experienced by Hilbert and Russell would await.<sup>78</sup> But in the second case, humbler but useful results are yet possible.

---

<sup>76</sup> Barfield and Auden, 18.

<sup>77</sup> Barfield and Auden, *History in English Words*, 5.

<sup>78</sup> At the turn of the twentieth century, David Hilbert attempted to carry out a complete formalization of mathematics. He was particularly interested in the independence of axioms in geometry and the consistency of the resulting theorems. At first, Russell's work on the logical foundations of mathematics seemed to provide the necessary underpinnings. But a series of irresolvable contradictions were discovered, culminating in Kurt Gödel's proof of two incompleteness theorems. See



## Time Dependency

A second critique of language models is the fact that texts written at different times use the same terms in ways that can be markedly different. Language models are usually built by “training” a neural network with a large amount of text, and using the next word as the error signal to backpropagate and improve the network. Thus, the network is trained by predicting what word should follow the current word it is “reading”. However, such a procedure does not take into account the date each text was written. Words are implicitly treated as objects outside of history, with stable meanings, even though in actual fact words change dramatically over time.

This is particularly true of words relating to changing religious and philosophical concepts and technological developments. For instance, Barfield begins *History In English Words* with the example of the word ‘electric’. If we were to find the phrase ‘the atmosphere was electric’ in a letter purported to be written by Dr. Johnson, the word’s context would prove that the document had not in fact been written by Dr. Johnson—not because the thing meant by ‘electric’ did not exist, but because at that time the word simply meant an ill-defined property ‘in some bodies, whereby when rubbed so as to grow warm, they draw little bits of paper, or such-like substances, to them’.<sup>79</sup>

The burgeoning of the understanding of the phenomenon we call electricity lie in the ancient observation that amber and other similar substances exhibit a mysterious property when rubbed. Only after Johnson’s time would humans realize that this phenomenon, so-named because the Greek term for amber is ‘elektron’, is connected also to the production of lightning in storm-clouds and related to magnetism and light and many other phenomena besides. This example suffices to show that a word can change its meaning dramatically over time, through technical and scientific development as well as through the employment of poetic metaphor.

Linguists have been aware of this fact since the beginning of the discipline, and a variety of ‘laws’ of semantic change have been proposed. For instance, the Law of Differentiation

---

<https://plato.stanford.edu/entries/hilbert-program/>

79 Owen Barfield, *History in English Words*, Eerdmans (1967), pp. 15-16.

proposes that near-synonyms tend to diverge in their meaning over time, and the Law of Parallel Change proposes that related words tend to change in similar ways over time.<sup>80</sup>

A more recent survey of computational approaches to semantic change, Tahmasebi *et al.* (2019), observes that “a variety of techniques for lexical semantic change and detection” have been developed over the last few years. “The state of the art is represented by methods based on word embedding techniques. However, most of these approaches are sense-agnostic, effectively focusing on the mixture of word senses expressed by a lexeme. Although some claim that their methods utilize the dominant word sense, they use each occurrence of the lexeme or word form without detecting if it is indeed representing the dominant sense or not.”<sup>81</sup>

The point is well-taken. All forms of word embedding, even sophisticated transformer models, produce vectors to represent word meanings that are an average of several meanings. A map containing all discrete meanings of a lexeme, like an automatic dictionary, has proven too unwieldy for use as a model of human speech and writing.<sup>82</sup> And yet, if language models only provide an approximation of meaning—and worse still, average across quite unrelated meanings—their usefulness in creating a map of invariant structures would diminish notably.

Another point made by Tahmasebi *et al.* is that “only a few approaches propose techniques capable of analyzing semantic change in words with relatively few occurrences.”<sup>83</sup> In scholarship relating to remote historical periods, for which only small quantities of text are available, this point is highly relevant. A small number of scholars are attempting to face these issues by constructing language models that take historical development into account. For instance, Hosseini *et al.* (2021) present four types of neural language models trained on an English-language dataset published between 1760 and 1900.<sup>84</sup> Their dataset and models are

---

80 Kevin J. Peterson and Hongfang Liu, “An Examination of the Statistical Laws of Semantic Change in Clinical Notes,” *AMIA Joint Summits on Translational Science Proceedings. AMIA Joint Summits on Translational Science 2021* (2021): 515–24.

81 Nina Tahmasebi, Lars Borin, and Adam Jatowt, “Survey of Computational Approaches to Lexical Semantic Change” (arXiv, March 13, 2019), <http://arxiv.org/abs/1811.06278>.

82 But see the Cyc project of the recently deceased Doug Lenat, which has had moderate success in modeling language and logical relations using rules instead of probabilities.

83 Nina Tahmasebi, Lars Borin, and Adam Jatowt, “Survey of Computational Approaches to Lexical Semantic Change” (arXiv, March 13, 2019), <http://arxiv.org/abs/1811.06278>.

84 Kasra Hosseini et al., “Neural Language Models for Nineteenth-Century English” (arXiv, May 24, 2021), <http://arxiv.org/abs/2105.11321>.

publicly available and have been shown to improve several NLP tasks relating to word meaning change over time.<sup>85</sup>

Future work in this field could involve mapping semantic changes in a large number of words as a way to validate language models and attempt to account for change over time. I expect that the large amount of resources dedicated to improving the state of the art in language models will continue to produce improved versions of the existing tools, and further breakthroughs may well solve the outstanding problems discussed above.

## Words Slip, Slide, and Sometimes Break

A central mystery in linguistics is: how can human children learn rich word meanings from very few examples, and how is it that at the origin of language, these meanings were developed in the first place?<sup>86</sup> We also wonder what this fact implies about the structure of grammar and the human brain.<sup>87</sup>

---

85 Others are attempting to solve the time-dependency problem of language models. Zampieri et. al. (2016) modeled text change in a Portuguese corpus containing texts from the 16th century to the early 20th century containing over 5 million tokens. Marcos Zampieri, Shervin Malmasi, and Mark Dras, “Modeling Language Change in Historical Corpora: The Case of Portuguese” (arXiv, September 30, 2016), <http://arxiv.org/abs/1610.00030>. Zhu and Bhat (2021) present a way of improving the BERT model to detect drug-related euphemisms in social media communication. Their results are between 20 and 50 percent better than baseline static language models. It is to be expected that this field of research will grow rapidly as automated content-moderation systems are employed increasingly by social media, business, and government across the world. Wanzheng Zhu and Suma Bhat, “Euphemistic Phrase Detection by Masked Language Model” (arXiv, September 10, 2021), <http://arxiv.org/abs/2109.04666>. Yixiao Wang et al., “Deriving Word Vectors from Contextualized Language Models Using Topic-Aware Mention Selection” (arXiv, June 15, 2021), <http://arxiv.org/abs/2106.07947>. attempt to improve existing language models by using contextualized language models rather than ‘bag of words’ vectors to encode contexts, and they use a topic model to partition the contexts in which words appear in order to learn different topic-specific vectors for each word. Their approach improves existing models and leads to high-quality word vectors. A recent conference, the Eighth Estonian Digital Humanities Conference (October 2022) was entirely dedicated to studying computational approaches to variation and change in language and culture. Qiu and Xu (2022) have built the most sophisticated historically-sensitive BERT-based language model yet, which they call HistBERT. See Wenjun Qiu and Yang Xu, “HistBERT: A Pre-Trained Language Model for Diachronic Lexical Semantic Analysis,” 2022, <https://doi.org/10.13140/RG.2.2.14905.44649>. They appropriately challenge the assumptions involved in applying pre-trained BERT models to historical texts: these models were trained on corpora that are almost entirely composed of contemporary English text. Qiu and Xu present results that suggest that their method improves BERT’s capacity to detect semantic shifts, but also say that more studies are necessary.

86 Important contributions to this theme are found in Noam Chomsky, *Syntactic Structures* (Mouton, 1957). See also Lenneberg, *Biological Foundations of Language*.

87 Noam Chomsky was one of the most important scholars to propose a mathematical theory of syntax. He has emphasized the structures of languages and grammars, including artificial grammars that have no ‘meaning’ but that model the mathematical characteristics of natural languages. Some of his students, for instance

These fundamental problems may or may not admit solution at some point in the future. In the meantime, they should cause us to pause and consider both the strengths and the limitations of our existing language models. It does seem that, as Barfield points out: “We think by means of words and we have to use the same ones for so many different thoughts that, as soon as new meanings have entered into one set, they creep into all our theories and begin to mould our whole cosmos; and from the theories they pass into more words, and so into our lives and institutions.”<sup>88</sup>

We should not expect that our presently-available language models be more than snapshots of the state of a language at a time. Perhaps, if enough text can be found to span some number of decades, they can become a partial history of the development of a language. The imprint left by many human hands may thus be mapped into vector space. But the continual activity of those hands and the mysterious source of their activity remain outside the map, a mystery which we should not underestimate or ignore.

A few more examples will make this point sufficiently clear. Consider the English language as spoken by the Irish. They were oppressed for hundreds of years and forced to speak the language of their oppressors.<sup>89</sup> Their own language was forbidden, as was their music and their religious worship. It is reasonable to expect that such a people would interpret the language of their oppressors in a highly original and layered manner. The meaning expressed in a language also exceeds its lexical content through moods like irony, sarcasm, and hyperbole, as well as the use of euphemisms as briefly discussed above. One can damn with faint praise. One can also emphasize by understatement.

This observation helps us to avoid the simplistic way of understanding words as ciphers in a code that could be exchanged with other ciphers to translate into another language. In this view, the things represented by the ciphers would remain unchanged, while they were assigned to new variable, a word in another language. But as anyone who speaks more than one

---

Andrea Moro, have continued his line of research and by resisting the siren call of brute-force connectivist neural networks have made important contributions to a richer understanding of grammar.

88 *History in English Words*, 189.

89 Laws forbidding the use of Irish Gaelic date to 1367. See Pádraig Óg Ó Ruairc, “‘To extinguish their sinister traditions and customs’ – the historic bans on the legal use of the Irish and Welsh languages”, *The Irish Story* (2018). At <https://www.theirishstory.com/2018/10/11/to-extinguish-their-sinister-traditions-and-customs-the-historic-bans-on-the-legal-use-of-the-irish-and-welsh-languages/>, accessed 29 March 2024.

language knows, this is not how translation works except at its most basic level. What makes up the character of a language and a culture is somehow between the lines, expressed in the juxtapositions and in the silences as well as in the explicit lexemes.<sup>90</sup>

## In Defense of Using Language Models

And yet, part of the central argument of this thesis is that even our current language models, with all their shortcomings, *are* a useful tool for the historian. We return to the point made early on: even in a narrowly-defined field of study, there is more information than a single scholar is able to absorb and interpret even in many lifetimes. If scholars are not to despair of improving their interpretations, they do need the help of computational techniques. They need ways to access a broader range of existing information and interpretations, and ways to remember more precisely the contents of the material they access.

So even if my proposal is not yet a mature system, since it is built with shaky components like the language models I have been criticizing, it is worthwhile that it be built. And since my main point is methodological, I may hold the hope that as its constituent parts are improved, so will its overall results. At present, as we have seen, the accuracy of language models are quite imperfect, but the subfields involved in this technique are advancing rapidly and it is likely that their results will improve too.

Statistical models of language will probably never arrive at what we mean by the word ‘understanding’ for many reasons<sup>91</sup>; yet even a few-percent improvement in the ability to manage massive quantities of text – from the zero percent of those texts that lie unread, to some

---

90 T.S. Eliot’s comparison of Shakespeare and Dante casts light on this issue. He points out that Dante is much easier to translate than Shakespeare because the Italian poet speaks in images. His words are mostly ciphers for concrete objects and actions and thus can be translated to a much greater degree than can Shakespeare, whose language depends greatly on atmosphere and wordplay to achieve his effect. It is probably also relevant that the English dictionary is far larger than the Italian, and so there are more words that can convey a specific atmosphere in English than in Italian. Dag Hammarskjöld wrote in *Markings*: “Respect for the word is the first commandment in the discipline by which a man can be educated to maturity – intellectual, emotional, and moral. Respect for the word – to employ it with scrupulous care and an incorruptible heartfelt love of truth – is essential if there is to be any growth in a society or in the human race. To misuse the word is to show contempt for man. It undermines the bridges and poisons the wells. It causes Man to regress down the long path of his evolution.”

91 See a brilliant exposition of some of these reasons in Andrea Moro, *The Boundaries of Babel*, 2nd ed. (MIT Press, 2015).

small percentage of ‘understanding’ offered by the use of automated systems, could represent a useful advancement for the field.<sup>92</sup> In order to begin to develop and detail this hypothesis, I now give three specific examples of what statistical representation of language can do for a historian.

## **Example: Are Mesopotamian Gods Active?**

While studying the comparison of Mesopotamian divinities with the Biblical God, I came upon an intriguing metaphor. The metaphor was the concept of “transitivity” applied to the gods. Thorkild Jacobsen used the word in several articles of the 1950s through the 1970s, first in the Haskill lectures he gave on the figure of Dumuzi/Tammuz at Oberlin College in the early 1950s. Then in 1961, he published *Formative Tendencies in Sumerian Religion*, where he wrote that

It is characteristic for Sumerian religion, especially in its older phases, that the human reaction to the experience of the numinous remained singularly bound by the situation in which the numinous was encountered, and by some central phenomenon or group of phenomena in it particularly. The numinous appears to be immediately and unreflectingly apprehended as a power in, underlying, and willing the phenomenon, as a power within it for it to come into being, to unfold in this its particular and distinctive form. In consequence the phenomenon largely circumscribes the power, for the numinous will and direction appear as fulfilled in the phenomenon and do not significantly transgress it. This boundedness to a phenomenon one might describe with a grammatical metaphor as intransitivity.<sup>93</sup>

The article “Towards the Image of Tammuz” (1962) contains a rewritten version of part of the Haskill lectures. It is not, Jacobsen clarifies, that Tammuz as power in the milk dies when the churn and the cup are empty, or that Tammuz as the power in grain dies when the grain is crushed between millstones. Rather, it is that the Tammuz-power “does not, either in action or as will and direction, ever transcend the phenomenon in which it dwells.”<sup>94</sup> Tammuz does not act on anything or anybody, his activity dissolves into the contrast between being and not being, living and dying. In the spring, the power “is there”. In the winter, the power is absent. The power inheres completely in the phenomena that manifest it, which makes Tammuz a poor helper to needy humans. In fact, Jacobsen points out that very few prayers are addressed to

---

92 If the introduction of errors does not exceed the increase in information, of course.

93 Thorkild Jacobsen, *Formative Tendencies in Sumerian Religion* (Doubleday, 1961).

94 Thorkild Jacobsen, “Toward the Image of Tammuz,” *History of Religions* 1, no. 2 (n.d.). p. 191.

him.<sup>95</sup> At most, he seems a “delightfully self-centered youth”, the “intransitive” object of others’ attention.

Jacobsen contrasts this view of things to the wider range of activity and power in other members of the Mesopotamian pantheon, such as Enlil, Enki, and Ninurta. These three gods are also connected to specific phenomena—storms, sweet waters, and thunderclouds, respectively—but their activities are broader. They can be of help or remain silent; create or destroy; they make demands and enforce them. They can be prayed to, and Jacobsen describes them as “transitive active” powers.

The metaphor appears again at the beginning of Jacobsen’s *Treasures of Darkness* (1976) to refer to the characteristic of deities who “made no demands, did not act, merely came into being, [were], and ceased being in and with [their] characteristic phenomenon.”<sup>96</sup> In this context, Jacobsen further points out that this characteristic was common to the older strata in the Mesopotamian pantheon, and contrasts with the younger gods who had power and interests beyond their defining characteristic phenomenon. Again, Jacobsen presents Tammuz as an “intransitive” figure of the power of fertility and new life in the spring, and writes that “there is no instance in which the god acts, orders, or demands; he merely is or is not.”<sup>97</sup> According to Jacobsen, later gods such as Marduk present elements of personal activity, or “transitivity”, that can be contrasted with the passivity he emphasizes in the figure of Tammuz.

Jacobsen’s idea sounds convincing, especially within the context of his long-running interest in Tammuz. But could it be supported with reference to a broader swath of Mesopotamian historical texts? Is it really true that the other gods never “act, order, or demand”?

Giorgio Buccellati proposes that one of the decisive characteristics that distinguish Mesopotamian spirituality from the adjacent Biblical spirituality is the transitivity-intransitivity metaphor. He elevates it to a category that describes not only a moment in the process of development of Mesopotamian religiosity, but that defines the gods in general. In *When on*

---

95 Jacobsen., pp. 74-78.

96 Thorkild Jacobsen, *The Treasures of Darkness: A History of Mesopotamian Religion* (New Haven London: Yale University Press, 1976).

97 Jacobsen.p. 10

*High the Heavens*, Buccellati presents an alternative between “whether one sees the absolute as the subject or as the object of the relationship that is perceived by individuals and by human groups in their intuition of the absolute.” He writes that

the two terms (subject and object) are to be understood in the syntactic sense of what governs a predicate or is governed by it, with the consequence that the nature of the absolute radically changes the nature of the predicate itself, as follows. In Mesopotamia, we can say that the predicate is intransitive, because the absolute (subject or object) is not conceived in a personal way. As such, it is not really the term of a relationship that is based on a face-to-face encounter. The Mesopotamian absolute does not look us in the eye, nor can human beings look him in the eye. Because, in fact, it has neither face nor eyes. This is how the absolute, seen as a profound immanence, qualifies the nature of the relationship. ... In the biblical perception, exactly the opposite occurs. The relationship is transitive in the specific sense that there is a face-to-face posture that is symmetrical with, if not on the same level as, the absolute conceived as God. This profoundly colors the nature of the relationship, and indeed changes its deep structure, even when there are strong similarities in external appearances.<sup>98</sup>

Note Buccellati’s enrichment of the distinction that Jacobsen indicated with the word “transitivity”. Not only is there an axis regarding action being done to a subject, but there is also an axis regarding the *personal* nature of the action done, the “face-to-face” posture mentioned above. Buccellati considers the difference between transitive and intransitive versions of the divinity as indicative of a radically different religious structure in Mesopotamia and in the Biblical world. When I encountered these ideas, I began to wonder if direct evidence of this difference could be found within the sacred texts themselves.<sup>99</sup>

The difference described by Jacobsen could be a moment in development from intransitivity to transitivity, as he seems to imply. The later gods (who are also closer in time to the formation of the biblical texts) exhibit more personal activity than the earlier gods like Tammuz. Following Jacobsen’s logic, one could think of the biblical God as a rarefied and simplified version of what came before, but not as exhibiting a difference in kind.

However, such a conclusion would not explain two notable differences between the Mesopotamian and Biblical religions. First, the concept of fate in the Mesopotamian religion does not have a personal character, and cannot be said to ‘change its mind’ even when the gods argue over what should happen in the future. This is a clearly different understanding from

---

98 Giorgio Buccellati, “*When on High the Heavens...*”: *Mesopotamian Religion and Spirituality with Reference to the Biblical World*, trans. Jonah Lynch (London New York: Routledge, 2024)., §7.3

99 A similar intuition was proposed by Friedrich Max Müller in his *Introduction to the Science of Religion* (London, 1871): 366-70. With thanks to prof. Andrea Moro, who suggested I look in this direction.



what is presented in the Bible, where the One God makes decisions, changes his mind, and acts with sovereign power.

Secondly, we note that even if Mesopotamia moved toward a personal conception of the absolute over its millennia of existence, as Jacobsen implies, it was not sufficiently convinced of the power of that absolute to survive the fall of the social and political structures that held its culture together. Buccellati points out that the people of the Bible, on the other hand, notwithstanding utterly destructive events, continued to read their history through the lens of a personal relationship with a personal—and historically active—absolute. This fact supports the claim that together with deep similarities between the two cultures and religions, there is also a radical structural difference.

### **Transitivity from a grammatical point of view**

Could these alternative interpretations be tested and verified by looking for transitive and intransitive forms in the Mesopotamian and Biblical sacred texts?

The grammatical term “transitivity” can be understood as referring to a verb that acts through to a subject, often the direct object of the phrase. For instance, “The pigs splashed mud on the wall” can be contrasted with “Mud splashed on the wall (when the pigs ran past).”<sup>100</sup> Some linguists have proposed more complex understandings of the structure indicated with the term. For instance, Michael Halliday proposes a further distinction between transitivity and ergativity, and places the focus more on the process and its initiator than on the type of verb used. Halliday points out that a great number of verbs can be used in both senses, transitive and intransitive, on the basis of their subordinate clauses. He notes that “these concepts relate more appropriately to the clause than to the verb. Transitivity is a system of the clause, affecting not only the verb serving as Process but also participants and circumstances.”<sup>101</sup>

Thus, in Halliday’s view, it is not so important to look at the verb as it is to examine the clause. He suggests that “(i) generalization across process types and (ii) transitivity model are independently variable. In English and in many other languages, it is the transitive model that

---

<sup>100</sup>See Kenneth L. Hale and Samuel Jay Keyser, *Prolegomenon to a Theory of Argument Structure*, Linguistic Inquiry Monographs 39 (Cambridge, Mass: MIT Press, 2002), 31.

<sup>101</sup>M. A. K. Halliday and Christian M. I. M. Matthiessen, *Halliday’s Introduction to Functional Grammar*, Fourth Edition (Milton Park, Abingdon, Oxon: Routledge, 2014), p. 226

differentiates the different process types and it is the ergative model that generalizes across these different process types.”<sup>102</sup> Halliday’s transitive model construes the actor as bringing about the unfolding of a process in time, thus supporting Buccellati’s distinction between transitive and intransitive versions of divinity as a way to differentiate between an absolute who acts as a "person" (or "affecting presence") with an unpredictable will, and an absolute which, like fate, simply “is” without acting or willing.

If divine action can be construed as an “occurrence”, fate bringing about an event without will or finality, or instead as the action of a willing person, I wondered if it would be possible to find grammatical features of the religious texts of the two traditions that support one or the other interpretation.

Apparent counterexamples to Buccellati’s claim exist in the Mesopotamian corpus. See for instance Enki and Ninmah: “Enki answered Ninmah: I decreed a fate for the first man with the weak hands, I gave him bread. I decreed a fate for the man who turned back the light, I gave him bread. I decreed a fate for the man with broken, paralyzed feet, I gave him bread. I decreed a fate for the man who could not hold back his urine, I gave him bread. I decreed a fate for the woman who could not give birth, I gave her bread.”<sup>103</sup> The god seems “transitive” in this case: active, present and speaking in the first person.<sup>104</sup> Certainly, we would expect large corpora to contain a variety of formulations, and not exhibit strict formal precision regarding the way the text speaks of the gods. But taken as a whole, I reasoned, a statistical analysis of the texts (translated into English) might reveal that the Mesopotamian corpus tends toward “intransitivity”, and the Bible to “transitivity”.

Until recently, answering such a question would have required many months of patient work collating texts manually. But in our time, automated natural language processing makes it possible to write a program to quickly parse the texts of the Mesopotamian and Biblical religious writings into their dependencies. On the basis of this parsing, it is possible to read

---

<sup>102</sup>Halliday and Matthiessen., p. 334.

<sup>103</sup>1.1.2 ETCSL

<sup>104</sup>According to Buccellati however, these are not really counterexamples, because "to destine a destiny" does not imply a control over fate.

through thousands of lines of text and filter and count several grammatical forms that might indicate the “transitivity” in discussion.

In the first iteration of the program, I used the SpaCy library as the basis of a parsing system that identifies the universal dependency labels of words in the English language translation of the religious texts in both Mesopotamian and Biblical traditions. SpaCy<sup>105</sup>, which released its third version in 2020, has the Word2Vec approach at its core. SpaCy is oriented toward the development of applications more than research, unlike other popular language processing packages such as NLTK<sup>106</sup> and CoreNLP<sup>107</sup>. It has been rapidly updated to take advantage of technological developments, including transformers such as BERT, and is supported with easy to follow tutorials and example programs.<sup>108</sup> As such, it seemed well suited to my task: I wanted a system that would help extract keywords, group them by similarity, and help me do so quickly. Improvements to precision and accuracy could wait until after I had validated the underlying logic.

I chose to work with translations in order to simplify my first-order procedure. It is important to recognize that this choice is fraught with consequences: the translations are vastly different, and comparison between them is affected by their differences. One may hope that translation allows us to find a match between two corpora that are originally in two different languages, by locating the words in the language of translation which renders the same concept in the two different words in the original. This can be seen as a way to resolve a fundamental

---

105See [spacy.io](https://spacy.io)

106NLTK, the Natural Language Tool Kit, was developed at the University of Pennsylvania and made available for free public use with a first release in 2001. It has been updated frequently over the course of the last two decades. See Bird, Steven, Edward Loper and Ewan Klein (2009), *Natural Language Processing with Python*. O’Reilly Media Inc. Over 16,000 articles have been published about this system. See [scholar.google.com.au/scholar?q=%22natural+language+toolkit%22](https://scholar.google.com.au/scholar?q=%22natural+language+toolkit%22).

107CoreNLP was developed by a team at Stanford that began publishing in 1999. It currently supports eight languages (Arabic, Chinese, English, French, German, Hungarian, Italian, and Spanish). By contrast, SpaCy currently supports 64 languages, 19 of which have pretrained pipelines available. See [spacy.io/usage/facts-figures](https://spacy.io/usage/facts-figures). For more information about CoreNLP, see Manning, Christopher D., Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. *The Stanford CoreNLP Natural Language Processing Toolkit In Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pp. 55-60. See also [nlp.stanford.edu/](https://nlp.stanford.edu/) for CoreNLP generalities, and [github.com/stanfordnlp/CoreNLP](https://github.com/stanfordnlp/CoreNLP) for source code and updates.

108Transformers are a new way to make neural network architectures that have led to dramatic improvements in accuracy and predictive capacity, most famously in GPT-3, developed by OpenAI, which at the time of writing holds the crown as the most impressive predictive natural language processing system.

problem with comparing texts in two different languages. To some extent, translation does this, but not in a systematic way.<sup>109</sup>

In my approach, syntax is more important than single words, and lends itself more readily to this procedure. Translations, especially literal ones, tend to represent fairly closely the structure of the constituents in a phrase or sentence. Since I am looking for transitive and intransitive structures, it is less important to capture the shades of meaning attached to words that are different in the original but translate to the same English word. It is more important to reveal the structure of the phrases, which is visible in translation.

Finally, I can justify my approach because this is a functional test which could be used on different corpora, including on the original languages, given an adequate language model. My use of translation limits the value of the specific conclusions this experiment can support, but it still permits me to demonstrate the overall method which is the proper object of this thesis.

The program<sup>110</sup> used a logical structure as follows: a list of names of gods was manually prepared. If the name of a god is found in a sentence, the program considers several possibilities. If the name of a god is the subject of the sentence, and the head of the phrase is a verb, then it is considered a case of an “acting god”, and the direct object of the action is identified if possible. If the divine name is the subject of the sentence, or if it is the direct object of the phrase, but does not act, then the phrase is counted as a case of a “passive god”, where the divine appears to be an object acted upon by some other actor.

The English translations of Sumerian religious texts in the ETCSL library, and the NRSV English translation of the Bible (excluding deuterocanonical and New Testament texts), were run through the program. The outputs include lists of subject, verb, object and complete phrase for error checking, and counts of how many times the divine name appears as active or passive. By normalizing (dividing by the total number of phrases containing the divine name), the two corpora can be compared.

---

109For instance, in the gospel both *λόγος* and *ῥῆμα* are translated as "word". In the first case, one might distinguish it by capitalizing it, *Λόγος* or *Word*.

110The source code of the program is in the appendix. It is also publicly available on my GitHub account. See [github.com/sibelius/transitivity](https://github.com/sibelius/transitivity). This process currently has about 87% accuracy using the models available on SpaCy.

I found that there is not a great statistical difference in the presence of the two verb forms as I defined them in the two corpora: by the definition of activity and passivity used, the Mesopotamian gods are 15.2 percent active, and 3.3 percent passive.<sup>111</sup> By contrast, the Biblical God is 16.4 percent active and 3.0 percent passive. These figures marginally support the hypothesis that the Biblical God is more active than the Mesopotamian gods, but the difference probably not significant.

However, by searching for occurrences of the verb forms “to say” and “to speak” in phrases including divine names, a more significant difference appears. Biblical texts of this form are nearly eight times more frequent: Mesopotamian gods speak 1.8 percent of the time, whereas the Biblical God speaks 14.2 percent of the time. This result tends to support our hypothesis, insofar as speaking can be construed as a “personal” trait. Words do not merely appear; they are spoken by a person who decides to say them.

This initial program has several weaknesses, which professors Mauro Giorgieri and Maria Freddi from the University of Pavia and prof. Marco Passarotti of the Catholic University of Milan helped me notice. First, as mentioned, it operates on English translations instead of on original texts.<sup>112</sup>

Another serious problem is that the binary opposition between phrases in which the god is the subject of a verbal clause, versus phrases in which the god is the object of the verbal clause, does not entirely capture the difference Jacobsen called “transitivity”. Some direct forms of speech would be correctly classified: “God speaks to man” while other equivalent phrases such as “man is spoken to by God” would not. Further, some Mesopotamian texts recount conversations among gods: “Ninurta spoke to Enlil”. In this case, should the god be classified as active or passive? Clearly, the program was not able to generate robust conclusions without further development. I had to look deeper for evidence of a structural difference between Mesopotamian and Biblical spiritualities.

---

111 This proportion reflects a widely noted feature of language: active forms are more common than passive forms in general. See David Banks, “The Extent to Which the Passive Voice Is Used in the Scientific Journal Article, 1985–2015,” *Functional Linguistics* 4, no. 1 (December 2017): 12, <https://doi.org/10.1186/s40554-017-0045-5>. for research on the active vs. passive voice in scientific articles.

112 Prof. Giorgieri suggested that we could justify using translations by assuming that whatever imprecision is introduced in this manner will affect both corpora equally.

## Example: Comparing Shuila prayers with the Psalms

In order to reduce the complexity of the text and give the algorithmic analysis a better chance of providing high quality information, I proceeded to work with two smaller corpora: the Shuila prayers<sup>113</sup> from Mesopotamia, and the book of Psalms from the Bible.<sup>114</sup> As above, my goal was to find evidence for a structural difference between Mesopotamian and Biblical religions.

My choice to restrict the corpora to these texts was determined by the fact that several scholars affirm that the two corpora are similar in intent and use, as well as in length and to some degree in style.<sup>115</sup> I hypothesized that these facts would reduce the complexity of analysis: instead of the enormous variety of styles, authors, and periods of composition in the overall corpus I attempted to analyze in my first example, two roughly comparable sets of prayers might prove homogeneous enough to show a clear difference.

Again I used the NRSV translation of the Psalms, and compared it to the entire catalog of Shuila prayers collected in transliterated and translated versions by Alan Lenzi.<sup>116</sup> As with many Mesopotamian texts, the Shuila prayers are quite fragmentary in many cases, but they do contain enough text to make statistical comparison possible.

I examined the typical language of each corpus, irrespective of speaker. Using a program, all verbs and adjectives were extracted, lemmatized, counted, normalized on the basis of corpus length (number of each verb / total number of verbs extracted), and graphed for the 60 verbs and adjectives that showed the greatest difference between corpora, which I call “prevalence”. Words for which either corpus gave zero results were excluded from the calculation of

---

113 Alan Lenzi, who has studied these prayers extensively and maintains a website archive of all the extant shuila prayers, writes that they are “liturgical ritual-prayers that were directed to the high deities of the Mesopotamian pantheon such as Marduk, Shamash, and Ishtar, among others. A ritual official (i.e., an exorcist) recited these prayers to assist a troubled client, often a Babylonian or Assyrian king. The exorcist would read the prayer aloud and the (presumably illiterate) client would repeat the words after him.” See <http://www.shuilas.org/>

114 This program is also available on GitHub: see [github.com/sibeliu/transitivity/blob/main/shuila%20vs%20psalms%20v2.ipynb](https://github.com/sibeliu/transitivity/blob/main/shuila%20vs%20psalms%20v2.ipynb)


115 In the introduction to his site [shuilas.org](http://www.shuilas.org), Alan Lenzi writes: “Comparatively inclined scholars in various fields of Religious Studies have used shuila-prayers to shed light on other religious corpora. For example, biblical scholars have been making comparative observations between the biblical psalms and the shuila-prayers in terms of both form and content since the late nineteenth century.”

116 See his [shuilas.org](http://www.shuilas.org) website.

prevalence. “Lift”, which is the meaning of the term “shuila” and appears in nearly every shuila prayer. In the Psalms, “be” is quite frequent, but the meaning of this fact is obscure. “To be” is a common helping verb, as well as an indicator of state. “Say” and “speak” appear at opposite extremes, which might indicate a stylistic difference in translation, or it might indicate a more profound difference, but it would be necessary to examine each occurrence to reach a robust conclusion.

After removing these words from the list, we see that in the Psalms, the next most prevalent verbs are *bless, let, deliver, praise, go, remember*. The correspondingly prevalent words in the Shuila prayers are *set, establish, accept, stand, look, forgive*. Results for 60 verbs are displayed in the figure below as a keyness graph. The list of verbs that display the greatest difference between the two corpora are listed vertically, and each one has a blue bar representing the degree to which it is prevalent. Bars that extend to the right of the vertical line indicate verbs that are prevalent in the Shuila corpus, to the percentage degree indicated on the x-axis. Thus, ‘lift’ is 17% more prevalent in the Shuila corpus than in the psalms. Verbs whose blue bars extend to the left are correspondingly more prevalent in the Psalms:

## Shuila vs. Psalms verb prevalence




One tantalizing feature of the graph is the distinct difference in the verb “say” (Biblical prevalence) and “speak” (Mesopotamian prevalence). Could this be a feature indicating the “transitive” biblical preference for a God who communicates directly, as a free “person”?



Could the use of the term “speak” in the Mesopotamian corpus indicate a more factual, fatalistic understanding of the communication of the gods? While the distinction is attractive, it is too tightly bound to the vagaries of translation and its proper interpretation too dependent on a line-by-line analysis of each occurrence, to firmly support one or the other conclusion. A case could be made for some difference in the relationship to the god on this basis, but it would not be strong. One would have to account for stylistic variety and differences in the translation, and it does not appear possible to arrive at any conclusion regarding “transitivity” on these grounds. The fact that the program counts lemmatized versions of the verbs eliminates finer-grained syntactic information about tense, mood, voice. Since my aim was to arrive at robust conclusions by means of algorithmic analysis, this result was still too weak.

My program also counted adjectives in the two corpora. Relative adjective frequency is shown below:

## Shuila vs. Psalms adjective prevalence



In this case, some evidence might be deduced from the importance of moral terms in the psalms. "Wicked" and "righteous" set up an alternative typical of the psalms, and emphasize the concept of a moral choice, while the shuila prayers emphasize the condition of the person praying as being in need, but not as having offended the gods with his actions. The

Mesopotamian image of god, at least as shown by the prevalence of the adjective, emphasizes “greatness”, which is not connected to the will. It is often used as a *captatio benevolentiae* in order to obtain the desired favor.

These explorations gave some tentative confirmation in favor of Buccellati’s affirmation that for the people of the Bible, God is a person who has a will and is active, while in Mesopotamia, man speaks to the god as to a technician, and expects a specific response that is not freely given or withheld, but is rather the result of a correctly executed prayer and ritual. However, it also became clear that basic statistical techniques involving counting words are too coarse-grained to characterize texts with a high enough degree of certainty to be useful to the historian. My hope to use language models as a tool in the service of history had not yet reached a definitive result. In order to proceed further, I needed to engage with the second feature announced at the beginning of this chapter: not only the fragments taken as isolated elements, the ancient texts algorithmically disintegrated into a cloud of words, but also their structure: the web of connections between the fragments.

## Structure

In his essay about the theoretical underpinnings of archaeology, *Critique of Archaeological Reason*, Buccellati argues for the “digital” character of archaeology, in the sense that the basic information in this field is first of all fragmentary. It may of course be postulated that the fragments had a unitary origin (many sherds derive from one vessel, for instance), and the scholar intends to recover that origin and interpret it. But the information itself, as it presents itself to the scholar at the starting point of his study of a “broken tradition”, is a (large) collection of fragments.<sup>117</sup> A “grammar” of the fragments, in the sense of a reasonable heuristic for categorizing the information coming out of the ground, and in the sense of the slow build-up of reasonable inferences about the way the fragments relate to each other, is the foundation of interpretation.

Buccellati further develops this grammatical metaphor as regards the development and communication of interpretations of the fragments. It should be possible, he affirms, to move

---

<sup>117</sup> See Buccellati’s article on this topic at <https://critique-of-ar.net/themes/broken.htm>.

from a high-level interpretative narrative back down to the lowest level “atomistic” detail of the source material that informs and justifies the narrative. Given that archaeologists today enjoy the use of a wide variety of digital tools like computer databases and web sites, it should further be possible for the connections between fragments and narrative to be communicated in a highly transparent manner.

Indeed, Buccellati and his collaborators have created several websites that do just that. For instance, [urkesh.org](http://urkesh.org) presents an enormous variety of source material regarding the excavation of Urkesh in Syria, all the way down to the descriptive notes regarding individual sherds. These notes reveal all the levels of interaction with the single sherds: who first observed them, what their first impressions and interpretations were, and what later interactions (even years later) further revealed.<sup>118</sup>

The emphasis on structure led me to hypothesize that it might be possible to digitally represent the structure of the religious texts I was studying, and in this way arrive at a robust characterization of the differences between the two corpora.

I therefore returned to the Shuila – Psalms comparison and created a model of the relations between entities in these texts. My aim was to demonstrate that a network, or graph, representation of the relations present within the texts offers a way to compare the two corpora and uncover the underlying structural presuppositions of their authors.

## **Example: Representing Structural Differences**

As we saw above, merely counting words in the Shuila prayers and the Psalms did not reveal any strong correlations. It did not give access to properly structural characteristics, which was the central affirmation I was hoping to corroborate or critique. In an attempt to overcome these difficulties, I developed the beginnings of a methodology for comparing two religious experiences through graph representation and network analysis.

Instead of applying automatic techniques to extract syntactic and semantic features of the text, I decided to begin by manually extracting structural features of a small group of texts. These features include the relations between gods, humans, and the natural world. For instance,

---

<sup>118</sup>See for instance the entries under “Record” on the [urkesh.org](http://urkesh.org) site.

one Shuila prayer contains the affirmation “Shamash-shum-ukin fears Marduk”. I coded this relationship as a relation (“fears”) between two entities, a human and a god.

These structures were then represented using a computer language called Cypher, which is specialized in the representation and analysis of graph databases. A network representation of the relations between entities in the religious texts might corroborate Buccellati’s affirmation of a structural difference between Mesopotamian gods and the Biblical God by helping to describe the structure of the interplay between agents in the texts in a rigorous manner.<sup>119</sup> In this example, the computer was used for its ability to render complex visual structures automatically and interactively. If indeed this method proved robust, I could then extend its use to algorithmic analysis of the text.

---

<sup>119</sup>This section was first presented at the EASR conference in Pisa (2021). A video version of that presentation can be found at <https://www.youtube.com/watch?v=JCLU2CoN3NM>.

In order to compare anything, it is necessary to find an element that can be compared.<sup>120</sup> Could a graph representation of the experiences described in ancient texts be used as an invariant element to permit comparison?<sup>121</sup> By shifting the emphasis from the definitions of the “essences” that are represented by nodes (gods, God, etc.), and also away from the linguistic relations between the words in the sacred texts (such as the transitive character of verb forms), to the topological relations between nodes, perhaps new light could be shed on some of the essential characteristics of the two religions. If so, perhaps those characteristics would permit

---

<sup>120</sup>It has been pointed out that comparison takes place using the categories of the observer—and can misrepresent what an observed group actually thinks about itself. (The longstanding debate about the word “religion” is a salient example.) Some scholars therefore interpret religious and cultural structures fundamentally as an expression of power, and emphasize considerations about colonialism, imperialism, and “platonizing” categories of thought that can distort and misrepresent the cultures under examination. This critique sometimes implies that the groups under examination see each other as competitors in a zero-sum power struggle. Perhaps in reaction to the ugliness of such a struggle, other scholars emphasize an irenic approach in which each culture is seen only in its own terms, without comparison to other cultures, in order to avoid suggesting that one be considered “better” than another. In my view, both of these positions fall short of the goals of understanding, respecting, and esteeming one another. In a recent book, Corinne Dempsey summarizes the complaint against a comparative approach. It “conjures and imposes abstracted categories that too often erase culturally embedded distinctions and realities” Corinne G. Dempsey, *Bringing the Sacred down to Earth: Adventures in Comparative Religion* (New York: Oxford University Press, 2012), 5. She continues: “Likewise, critics of religion often note how religious systems impose on adherents spiritualizing abstractions that deflect and neglect material needs and realities. As both sets of critics have it, scholarly comparison and religion, imposed from above, easily lend themselves to imperialistic structures of oppression.” On page 7, she writes that “a core critique of comparativism today is of intellectual imperialism, a process that imposes universal categories that distort or disregard locally embedded meanings and differences.” She claims that her book was not intended as a defense of comparativism, and rather that it was the very dynamic of encountering new realities and attempting to “translate” the contents of other people’s experience into her own language that led her to affirm the comparative approach as an “interactive adventure that is never complete, perpetually enticing those who engage into further exploration and discovery.” Dempsey, 19. She recasts the problem: perhaps it is “not whether we should compare but, in comparing, how we can do so transparently and responsibly. How do we choose among infinitely available points of comparison such that the process advances rather than predetermines or undermines our knowledge of religious phenomena?” Dempsey, 4. Dempsey maintains that “the key to maintaining productive exchanges between scholar and subject, helping to deflect comparativism’s imperialistic potential, is the resolve to take seriously the religious experiences and expressions of those we study. Although this may seem an obvious suggestion, it appears that the current emphasis religion scholarship tends to place on power dynamics is precisely capable of diminishing this resolve, creating barriers between students of religion and the communities they study.” Dempsey, 10. I agree with Dempsey that it is possible and necessary to carry out a comparative study of religions. But it must be a continuous “adventure”, to use Dempsey’s word, a continuous hermeneutic circle, in which new information updates the information that is already available, and provisional conclusions are ever improved. Such comparison cannot escape categorization, because human experience forms patterns in the brain and is expressed in patterns of language. It would be interesting to explore the relevance of these issues in a further study.

<sup>121</sup>A first version of this section was presented at EASR 2021 in Pisa. A video version of that presentation can be found at <https://www.youtube.com/watch?v=JCLU2CoN3NM>.

comparison. Graph representations offer three particularly useful features to this end: they encode relations, not just “essences”; they are natively capable of continuously incorporating new data; and some of their characteristics are encoded in their very shape. It is this topology that principally interested me.

I chose a very small subset to begin with: four Shuila prayers dealing with lunar eclipses, and four Psalms.<sup>122</sup> The use of a small number of texts also permits easy comprehension of the images in the printed text of this thesis: a larger number of examples would necessitate mathematical analysis, and would obscure the fundamentally methodological point I wish to make here.

I manually encoded the relations expressed in those texts as a directed graph, using the Cypher language and Neo4j database management software. The Cypher language uses an ASCII-art representation of nodes and edges, as well as a multilayered tagging and properties structure to encode each node and relation with rich data. The first few lines of the encoding is below<sup>123</sup>:

```
1
2 //create gods
3 CREATE (enlil:God {name: 'Enlil', source: 'p393721'}),
4 (fate:God {name: 'Fate'}),
5
6 //create persons
7 (generic:Person {name: 'generic persons'}),
8
9 //create events
10 (eclipse:NaturalObject {name: 'lunar eclipse'}),
11
12 //create relations
13 (generic)-[:FEARS]→(eclipse),
14 (generic)-[:REQUESTS_LIFE]→(enlil),
15 (generic)-[:TRUSTS]→(enlil),
16 (generic)-[:REQUESTS_PROTECTION]→(enlil),
17 (fate)-[:CAUSES]→(eclipse),
18
```

*Figure 1: Cypher encoding of Shuila prayers and Psalms*

---

<sup>122</sup>Specifically: prayers with CDLI numbers 259033, 393721, 393771, and 393796 on Lenzi’s site [shuilas.org](http://shuilas.org).

The prayers, with original cuneiform text as a photo, complete with transcription and translation, can be found by replacing the number in the URL: <http://shuilas.org/P398671.html> The Psalms do not deal with lunar eclipses, so for a random first test corpus I simply used the last two digits of the Shuila prayers as references. Thus the psalms encoded are numbers 33, 21, 71, and 96 in the numbering of the Revised Standard Version (RSV). All prayers and psalms were encoded according to their English translations as available on [shuilas.org](http://shuilas.org) and in the RSV.

<sup>123</sup>The complete code is in the appendix. It is also publicly available on GitHub: see <https://github.com/sibelius/Creator-of-gods/blob/main/EASR%20psalms%20shuila%20graph>

By encoding the texts in this way, they could be visualized using the Neo4j software. The four Mesopotamian texts result in the following image:




Figure 2: Four Shuila prayers, represented as a directed graph


The lunar eclipse, represented in green as part of the natural world, is at the center of three “fear” relationships with human persons (in purple) who are praying using the text of the prayers. Some of them are named, some are generic clients of the clerical structures in Mesopotamia. The persons are also connected to several named divinities (in gold) through relations such as PRAISES, REQUESTS\_LIFE, and OFFERS\_INCENSE. Some divinities are named in the prayers, but not specifically addressed: these float unconnected to the graph.

In a similar fashion, the four Psalms can be encoded and represented as follows:





At the current stage of encoding, the two graphs can be compared intuitively by collecting the three types of nodes (divinities, humans, and natural world) in three locations.<sup>124</sup> The gods (colored in gold) are in the top of the image, humans (purple) are on the bottom right, and the natural world (green) on the bottom left. Below, the two graphs presented above have been manually rearranged in order to permit easy comparison between their structures, with Mesopotamia represented on the right and the Bible on the left:



*Figure 4: Side by side comparison of four Psalms (left) and four Shuilas (right), manually rearranged for clarity*

With only four randomly-chosen prayers to represent each religion, the conclusions that can be drawn should be taken with more than a grain of salt. Yet even at this low resolution, we can already see a few significant differences. While there is a large pantheon forming on the

---

<sup>124</sup>The two graphs presented above were encoded in a single database, and do not touch each other because none of their current nodes are identical. It would be expected that if a large number of texts were encoded, this might no longer be the case.

right (Mesopotamia), only one God is on the left (Bible). More importantly, and beyond mere numerical differences, the relations between humans and gods is *unidirectional* (from man to god) in the case of Mesopotamia, and *bidirectional* in the case of the Bible. Furthermore, measuring centrality shows that in the case of Mesopotamia, humans are more central, whereas YHWH is more central in the Bible. With a larger graph based on more texts, network analysis would permit more numerous and robust conclusions to be drawn.<sup>125</sup>

Within the limitations of the data I have encoded here, one can already see distinctly different structures in polytheism and monotheism as expressed in the collections of the sacred texts of Mesopotamia and the Biblical people. This fact tends to support the conclusions reached by Buccellati in his comparative study<sup>126</sup> and to critique the conclusions of Mark Smith, who proposed that there is not a difference of kind but rather of degree between Mesopotamian polytheism and Biblical monotheism.<sup>127</sup> Having been encoded as a graph, two religions or spiritualities can be compared by comparing their structures: the topology of two networks can be compared independently of the content represented by the network.<sup>128</sup> This fact is an essential element of my proposal of an “inner-referential” representation of historical information. Even without reference to an “outside”, we are able to characterize the topology of a network of information, and its shape offers a robust tool for understanding.

## Conclusion

As I stated at the outset, the goal is not to replace human intelligence by machines, but to augment it. I seek not fully automated processes of text analysis, but a symbiotic relationship between scholar and machine that could improve the properly humanistic, and human, work of the scholar. As we have begun to see in this chapter, language models are powerful but fragile,

<sup>125</sup>It seems that Fitz John Porter Poole dreamed of something like this comparison in his seminal article of 1986, “Metaphors and Maps”, but it was not yet possible with the technology of the day. See Fitz John Porter Poole, “Metaphors and Maps: Toward Comparison in the Anthropology of Religion,” *Journal of the American Academy of Religion* LIV, no. 3 (1986): 411–98, <https://doi.org/10.1093/jaarel/LIV.3.411>.

<sup>126</sup>Buccellati, *When on High the Heavens...*

<sup>127</sup>Mark S. Smith, *The Origins of Biblical Monotheism: Israel’s Polytheistic Background and the Ugaritic Texts*, 1. issued as an Oxford Univ. Press paperback (Oxford: Oxford Univ. Press, 2003). With more data, it might be possible to reach even more interesting conclusions. For example, this method could be applied to study the origins and goals of political structures and the institution of the temple.

<sup>128</sup>This method points toward a way to compare and contrast even those human experiences that are fraught with histories of struggle and misunderstanding.

and at present cannot answer all the detailed questions a historian wishes to ask. However, they can be a crucial part of our tools.

## Chapter 3 – Intellectual Prosthetics

### The ‘Extended Mind’

The technological augmentation of human abilities like memory have led some cognitive scientists to think that our prosthetics<sup>129</sup> are part of our mind itself. Lev Vygotsky postulated that some aspects of human cognition are carried out in the cultural environment through social and technological structures. Later, Edwin Hutchins proposed an approach to cognitive science that he called “distributed cognition”, in which he argued that cognition involves both external artifacts and cultural systems for interpreting reality.<sup>130</sup>

Hutchins’ theory influenced Andy Clark, who together with David Chalmers proposed in *The Extended Mind*<sup>131</sup> (1998) that the support structures we use to aid our memories should actually be considered as extensions of the mind itself.

The argument is intriguing. Clark and Chalmers describe a thought experiment about Otto, an Alzheimer’s patient who writes a complicated series of actions in a notebook that he carries in order to be able to successfully complete a long sequence of actions without forgetting, and Inga, a normally-gifted woman. Otto and Inga both go to a museum, which in Otto’s case requires the use of the notebook. The two authors claim that the main difference between these two cases is that Inga’s memory is internally processed in her brain, whereas Otto’s memory is externally processed in his notebook—and thus his brain has been “extended” beyond the boundaries of his body.<sup>132</sup>

---

129For one early (and rather tongue-in-cheek) use of this word in a context similar to my own, see Neal Stephenson, “Communication Prosthetics: Threat, or Menace?,” *Whole Earth*, Summer 2001.

130Some of his work focused on tools for navigation in traditional sea-faring cultures who used the mattang. See Edwin Hutchins, *Cognition in the Wild*, 8. pr, A Bradford Book (Cambridge, Mass.: MIT Press, 2006)..

131A. Clark and D. Chalmers, “The Extended Mind,” *Analysis* 58, no. 1 (January 1, 1998): 7–19, <https://doi.org/10.1093/analys/58.1.7>.

132It is possible to poke fun at it, such as what Adams and Aizawa wrote in 2010: “Question: Why did the pencil think that  $2+2 = 4$ ? Because it was coupled to the mathematician.” But such insistence on the concept of causality, to downplay the importance of external prosthetics, seems mostly to avoid engaging with the crucial point. Frederick Ray Adams and Kenneth Aizawa, *The Bounds of Cognition* (Malden, Ma: Blackwell Publishing, 2008).

This is not the place for an extended engagement with this idea insofar as it bears on the definition of “mind” and the concept that the mind and cognition could occur outside of the brain. Here, we accept the premise that external support structures can aid human memory, as they have since the invention of writing, libraries, moveable type, microfilm, and electronic storage and retrieval systems. For the aims of this thesis, it is not necessary to specify whether or not these external devices are to be considered part and parcel of the human “mind”. What is important first of all in our context is finding ways to overcome or at least expand the forgetfulness and restricted field of view that affects not only Alzheimer’s patients like Otto, but any human who attempts to manage a large amount of information.

Like natural scientists and engineers, historians now have access to too much information and must find a way to “tabulate the census faster than the next census will arrive”, as Herman Hollerith was tasked with doing in 1890.<sup>133</sup> Like Hollerith, contemporary historians must find a way to retain more information, to interpret the contents of large databases (say, archaeological excavations and large bodies of text), to publish their results more quickly—in order to improve the quality of their inferences and conclusions. The alternatives are unsatisfying: a historian might decide to focus only on a certain small domain, but he does so at his peril, for the dynamics in the broader world certainly affect whatever he decides lies within his domain of study as well. How can those “external” dynamics be taken into account? The human mind is not able, unaided, to progress much further than it had in the 19<sup>th</sup> century. Machine-aided research is a necessity.

But historians pose questions which are different from the questions a natural scientist faces. This is probably a major reason that history is late to the mechanized game, and is also resistant to it. For it must be possible to pose properly *historical* questions, and receive a

---

133 “The 1890 census was the point in history where the processing of *data* as well as the calculation of mathematical equations became the object of automation. As it turned out, Hollerith was neither a mathematician nor a logician, but a data processor. He was grappling, not with numerical calculation, but with the complexity of collecting, sorting, storing, and retrieving a large number of small items in a collection of information. Hollerith and his colleagues were unwitting forerunners of twentieth-century information workers, because their task had to do with finding a mechanical method to keep track of what their organization knew.” Howard Rheingold, *Tools for Thought: The History and Future of Mind-Expanding Technology*, 1st MIT Press ed (Cambridge, Mass: MIT Press, 2000).

properly historical aid from the machine. I believe that at least one part of the process can be improved at this time.

## The ‘Battle for Synthesis’

In his classic portrait of Mesopotamian civilization, A. Leo Oppenheim notes that some superficial scholars may be tempted to accept “glib popularizations” of the central words that define the field, but

this attitude on the part of an Assyriologist would border on cowardice. The battle for synthesis is the battle he is to fight, and this battle should be considered his *raison d’être*, even though it is a battle that can know no victorious outcome. The battle as such must be the task of the Assyriologist. Typically, however, we tend to escape into peripheral skirmishes. The field of Assyriology has grown so wide and so complex that not more than a handful of scholars can claim to be at home in its manifold domains. Most Assyriologists restrict their interest to apparently well-documented subdivisions and often select, in premature specialization, a specific area as their field of research.<sup>134</sup>

One of the reviewers of this work, Hildegard Lewy, also pointed out that “fifty years ago, the great Assyriologist Friedrich Delitzsch knew every cuneiform text which was available at that time. Since then, the number of texts has increased so much that nowadays no single individual can master every one of them. The sheer volume of the material as well as its diversity make it necessary for today’s Assyriologist to specialize in the texts of one or several regions and periods” (Lewy 138). But how can specialization contribute to the ‘battle for synthesis’ that Oppenheim sees as the essential task of the Assyriologist? Specialization alone will lead instead to further fragmentation and incommunicability between the specialties: coordination of specialties is also necessary.<sup>135</sup>

Oppenheim did not believe “that the diagnosis of our malaise allows for simple medication”<sup>136</sup>, but he did point out a few important directions for research that were bearing fruit in his day. One in particular was the “spectacular success” of interpreting cuneiform astronomical texts, thanks to collaboration between Assyriologists and mathematicians and astronomers. On the topic of collaboration, Oppenheim wrote that

---

134Adolf Leo Oppenheim and Erica Reiner, *Ancient Mesopotamia: Portrait of a Dead Civilization*, Rev. ed., 11. impr (Chicago: Univ. of Chicago Press, 1998), 28.

135Make discussion of micro vs. macro narratives. Grand narratives are problematic. Show awareness of discussion.

136*Ibid.*, 29.

This may be, at last, the solution of many problems that beset Assyriology. Perhaps the descriptive linguist will help us throw off the fetters that are hampering our progress [...]; the historian of medicine may well contribute essentially toward the understanding of numerous medical texts in cuneiform [...]; and the historian of technology will show us the way we should investigate, for example, the tablets describing the manufacture of colored glass...<sup>137</sup>

Collaboration is a crucial aspect of contemporary natural science. It is an approach that historians need help in initiating, since their discipline has been mostly dominated by solitary work instead of group collaboration. In this sense, the methodological issues are important: how can historians coordinate their information and render it retrievable at the right time? How can an “intellectual prosthetic” aid the short-sighted human memory to collect, manage, and synthesize ever-larger amounts of information?

The problem is as old as intellectual work itself: in order to correctly interpret new data, one needs a repository of existing knowledge, and an efficient way to find relevant information. The larger the archive, the more complete the research that can be carried out.

The Library of Congress in the US attempted to address this issue by collecting a copy of every book in print, until this no longer proved workable because of the sheer number of books published.<sup>138</sup> Search engines also attempt to collect a “capture” of every page on the internet, multiple times a year—which amounts to astronomical volumes of data stored on their servers. Archival and retrieval is a basic problem for anyone trying to study and write detailed works about any issue. Throughout history, we have developed systems that help us to organize resources and notes, and find them again when it comes time to write.

Each member of a scholarly community is engaged in collecting information and creating “models” of the reality they study.<sup>139</sup> These models are expressed as articles, books,

---

137The scholar from the University of Chicago did in fact follow this course of action with his study dedicated to the production of glass, in collaboration with Robert Brill of the Corning Museum of Glass. See Adolf Leo Oppenheim et al., eds., *Glass and Glassmaking in Ancient Mesopotamia: An Edition of the Cuneiform Texts with Contain Instructions for Glassmakers with a Catalogue of Surviving Objects*, Repr, The Corning Museum of Glass Monographs 3 (Corning: Corning Museum of Glass Pr. [u.a.], 1988).

138In 2013, more than 300,000 books were published in the United States. See <https://www.ingenta.com/blog-article/ipa-report-says-global-publishing-productivity-is-up-but-growth-is-down-2/>. For an estimate of the space required, one library planning guide (<http://mnl.mclinc.org/wp-content/uploads/2013/05/spaceplanningguide.pdf>) suggests that the necessary square footage can be estimated as a factor between one-tenth to one-twenty-fifth the number of books to be archived. That would mean that each year the Library of Congress, under the most conservative estimate, would need to acquire a new 12,000 square foot space.

139Talk about structuralism and my stance. Show that I know the objections.



conferences, and other forms of communication to a broader community of scholars who have the formation to understand and critique the models presented. This community also needs mechanisms through which to share information, recall it, give each other feedback, and as a group move toward greater understanding. This iterative process has taken the form of conferences and print publications for several generations; new tools seem to be needed today in order to process the ever-growing quantity of specialized information and form new syntheses.

Historians face the immense complexities involved in interpreting civilizations, the fragmentary nature of much of their data, and the constant stream of new discoveries that force them to re-think their interpretations. In the past, most historical research was published by single authors, and represented the conclusions of a single mind that had explored the sources and reached a synthesis. But archaeological research, which rose to prominence in the 19<sup>th</sup> century, began to change this, since it was not possible to excavate a site without the coordination of large groups of people. The cybernetic aspect of managing people, and the data management aspect of cataloging and archiving immense quantities of artifacts, opened the door to collaborative forms of historical research.

The research behind this thesis began within the context of a group of historians and archaeologists working on a large-scale interpretation of Mesopotamian civilization. While working on this project, I tried to manage several competing priorities. Our various opinions and points of expertise had to coalesce into a choral unity. It was not sufficient that our various voices sound independently, nor yet did we aim for simple uniformity. We also faced a problem due to our various levels of expertise. Some members of the team were highly trained and widely-read in the field; others arrived from outside the domain of studies on the Ancient Near East and needed to get up to speed on basic data, as well as contribute to the interpretation of new information. Finally, the very work we were collaborating to carry out, the 4banks.net group of websites, explicitly aims to be a forum for conversation and debate between different voices: how can it maintain its identity, and also host a variety of voices?

As I worked with these scholars, I began to see the need for a tool that would permit our collaboration to bear the complex and rich fruit we hoped for. Every new piece of information,

every reference to relevant articles and books, every interpretation, was valuable. And our team had already developed several effective ways to archive and retrieve the wealth of information we generated.

Even so, as our work grew it became impossible for all the scholars to take the mass of data into full account. In a certain sense, it is a simple question of combinatorics: the connections between data points increase as the square of their number, which quickly grows without bounds.<sup>140</sup> Now, it is true that not all data points are connected, so the actual number of relevant connections is much smaller. But the devil lies precisely in the detail: *which* data points are connected? Which connections are relevant? How can information be remembered, and how can it be found again in the archive?

The all-encompassing project of a total history is obviously impossible within the space of a single human brain and a single human life-time. After Hegel, few have thought themselves capable of such a project. But it is not necessary for historians to simply accept to become specialists. Such an outcome would be self-contradictory, because the interpretation of a detail depends upon its broader context. If the context is unknowable because it is too broad, then we could only hope to connect small domains of knowledge together, or produce scholarship which underlines details, heretofore unexplored correspondences perhaps. But they would be splinters of the whole, irremediably disconnected and unverifiable in their accuracy. In my view, it is not necessary to accept these outcomes, because as our data has grown, so have our technical capabilities.

## A Short Recent History of Memory Aids

We can begin to detail the “intellectual prosthetic” by reviewing some pertinent examples of past solutions to the same problem. Besides the more ancient versions of data archival and

---

<sup>140</sup>The number of calculations increases as the square of the documents—each document must be compared with all the others. The number of paths, however, must also take into account the fact that paths could travel through more than two nodes. Counting paths involves the use of the factorial. We recall that the factorial of a number is represented as  $n!$ , and is calculated by multiplying together the integers from 1 through  $n$ . Thus  $4!$  equals  $1 \times 2 \times 3 \times 4$ , which equals 24. The first five positive integers to the factorial are as follows:  $1! = 1$ ,  $2! = 2$ ,  $3! = 6$ ,  $4! = 24$ ,  $5! = 120$ . As the number of data points increases, their factorial increases extremely rapidly.

retrieval such as libraries, some more “technological” thinkers have proposed solutions that converge toward what I will propose here.

We can first refer to a curious antecedent, the “book wheel” invented in 1588 by Agostino Ramelli.<sup>141</sup> He never made a working version of the device, but it has been subsequently built by a few people. It allowed a reader to move between several open volumes displayed on a rotating wheel. It was like a primitive form of “browser tabs” that allowed the reader to consider several sources in rapid succession, instead of working through one volume at a time.<sup>142</sup>

Of course, in the contemporary context, the more important issue is not merely how multiple texts can be accessed in rapid succession, but rather how information that exceeds the ten or twelve volumes that could fit on such a wheel should be managed. What does one do with the thirteenth book, or the thirteen-hundredth?

One solution is to record information on notecards, collected in small boxes, which Niklas Luhmann called “zettelkasten”. Many scholars have found that paper-based zettelkasten served to extend their memory and improve their ability to manage a large amount of information over a long period of time.<sup>143</sup> For a note to be useful, according to Luhmann, it should contain several features. It should contain some text, usually copied from a source, along with bibliographical information to find that source again. It should also contain a heading, or a caption, which encapsulates in a few words what it is about. The action of writing a short summary is not only useful for recalling notes, but it is also a first and crucial part of the act of saving a note.<sup>144</sup> Tags could also be useful, or perhaps it could be better to consider tags as part of the unity of a caption; or perhaps multiple captions would be better: there is much debate about such matters.

---

141 <https://en.m.wikipedia.org/wiki/Bookwheel>

142 Older, somewhat analogous examples exist. For instance, Thomas Aquinas was famous for dictating several books almost simultaneously to a group of scribes, not unlike chess grandmasters who almost simultaneously face multiple opponents. I use the phrase “almost simultaneously” because the action occurs in sequence, not in parallel. This of course is more about performance than receptivity, but it bears a strong analogy to the closely serial processing allowed by a “book wheel”.

143 See Willard McCarty, “Making and Studying Notes: Towards a Cognitive Ecology of Annotation,” in *Annotating Scholarly Editions and Research: Functions, Differentiation, Systematization* (Berlin: Walter de Gruyter, 2020), 271–97.

144 These features are all present in the DABI files used on the Mesopotamian history project led by Buccellati.

In his paper *zettelkasten*, Luhmann suggested that it is useful to always code a link between the information in a new card and something in the existing archive. This linking creates an organic branching structure within the archive that in time becomes a rough copy of the memory of the human who created the map.<sup>145</sup> Notecards, linked together, slowly form a web of interconnected “atoms”, like a gigantic molecule slowly being synthesized.

The information that historians have regarding the past is mostly written. We speak of “prehistory” as the time before the invention of writing. Ancient history deals with the fascinating period of the beginning of writing, wherein we must make difficult inferences from data that is very incomplete, fragmentary, and also deals mostly with administrative issues that hide and falsify our view of the complete society we wish to understand. Perhaps this difficulty is also the attraction, as we push into the darkness at the origin of our human adventure.

The texts that we deal with are primary and secondary. We have material that comes directly from the time we wish to understand, and we have layers of interpretation that later generations have made. Inferences and patient reconstructions, new information and new inferences, running after each other down the centuries. Historical thought has grown exponentially in recent times.

The texts we use are dated, both in their origin and in the moment in which we discover them and begin to use them in our own research. The timing is often unplanned, yet it has an important effect on the formation of our conclusions. Some insights order a large amount of other information, and form the ontology that gives structure to an entire article or book. Other insights, while attractive at first, are proven false by further data. It is sometimes important for us to remember not only the information, but also the order in which we encountered it. From the collection of atoms, a large pattern emerges in time: a model of reality.<sup>146</sup>

---

145If more than one person contributes to the archive, multiple minds working on the branching structure will generate alternative structures that in some cases coincide and strengthen each other, and in other cases will be at odds. Computer based versions can solve some of those problems and lead once again to a fuller representation of the data and its interconnections. Many such systems have been invented, such as Evernote, Margin Notes, Dendron, Roam, The Archive, and more. See [zettelkasten.de](http://zettelkasten.de) for a helpful overview of the technique, as well as a community of aficionados and a software tool designed to help take notes with this method.

146Today there is a great deal of interest in automated intelligence, and the specter of systems that take over control from humans is a constant trope in popular fiction and film. The questions about the existence of freedom, and whether consciousness is an “emergent property” of complex information, are at the center of a

Feedback is a crucial part of the process. A scholar proposes an interpretation; another scholar proposes an alternative, and criticizes the first. This interplay is a crucial part of the endeavor to understand. Different models are built, played against each other, and in the interplay and conflict of interpretations, the models are improved. But when the quantity of information at play is large beyond some threshold, and when the languages involved are multiple and only imperfectly translatable, the risk grows exponentially that instead of being an interplay that reaches toward a common understanding, the conversation remains only a collection of skew lines that never meet, and do not add up to anything coherent even after a great deal of time and energy has been expended.

How can the interplay of intellectual model-making, -testing, and -improvement be aided? One of the first to pose the question and answer it in a fundamentally new way was Vannevar Bush.

### **Vannevar Bush and the Memex**

Vannevar Bush published a famous paper entitled “As We May Think” in *The Atlantic Monthly* in 1945. The short and surprisingly prescient article describes a “memex” capable of retaining not only a large library of documents (which Bush imagined as microfilmed texts) but also the “trails” connecting documents, encoded as the user read topics in sequence and annotated them with further comments.

At the time he wrote, Bush’s idea was imaginable using existing technology—in fact, a large part of the success of his article is probably due to this fact. He imagined how photo sensors could detect text, how speech could be automatically converted and saved to notes, how levers and punch-cards could reference the relevant information, and even estimated the

---

passionate debate. My interest in this dissertation does not follow this path. Instead, I am interested in the older path, as old as the invention of writing: the path of inventing intellectual technologies that expand the human ability to manage large amounts of information, but not to replace the properly human with a mechanism. In this, I am squarely in the tradition of the grandparents of computation: Babbage and Lovelace, Turing and von Neumann.

time (“between 3 and 5 seconds”) necessary for a search to produce information. But the device was never built with Bush’s methods. It remained a tantalizing mirage.

### **J.C.R. Licklider: the Birth of the Knowledge Worker**

According to Howard Rheingold, J.C.R. Licklider was impressed with Bush’s work and interested in optimizing his own work, and timed himself for a few weeks in order to understand where he was using his time, and came to the surprising conclusion that about 85% of his time was spent in clerical activities that could probably be done better and faster by a machine.<sup>147</sup> He then set out to build such a machine, in order to dedicate a greater part of his time to the task of interpreting the information he had gathered rather than in the mechanics of gathering and presenting that same information.

A similar dynamic is at work in the historian’s study: much time is spent gathering information, forgetting it, re-finding it, and it is natural to desire a tool that would offload some of this activity in such a way that the historian could focus on his central competency, interpreting the information gathered. Licklider was pragmatic, and wanted a machine to help him as a clerk. Howard Rheingold, a historian of the development of computing, wrote that

To those who were wild enough to make such a suggestion — especially the young MIT computer mavericks who were founding the field of artificial intelligence around that time — the idea might have seemed too obvious and too trivial to pursue. In any case, the AI founders were more interested in replacing the scientist than the scientist's file clerk. Licklider, however, was neither a respectable computer scientist nor a computer maverick, but a psychologist with some expertise in electronics. And like any other competent investigator, he followed where the data led him. [...] Then he learned that although the computer was the right *kind* of machine he needed to build his models, even the PDP-1 was hopelessly crude for the phenomena he wanted to study. Nature was far too complicated for 1960-style computers. He needed more memory components and faster processing of large amounts of calculations. As he began to think about the respective strengths and deficiencies of computers and brains, it occurred to him that what he was seeking was an alternative to the human-computer relationship as it then existed.<sup>148</sup>

What Licklider dreamed of in 1960 has now become reality, but it is worth noting that the amount of time that has passed is barely enough for such a revolution to have born mature fruit.

---

147 See Howard Rheingold, *Tools for Thought*, MIT Press (2000), chapter 7. I did not have access to a printed copy of this book, and read it in the online version at <https://www.rheingold.com/texts/tft/07.html#Chap07>

148 Howard Rheingold, *Tools for Thought*, MIT Press (2000), chapter 7. Available online at <https://www.rheingold.com/texts/tft/07.html#Chap07>

I therefore think it natural that what I currently am proposing is but another step forward. A way to handle the historian's information, an "intelligent assistant", and a way to extract greater value from it, is a further step on the path that began with the intellectual prosthetics developed in Mesopotamia six thousand years ago.

Licklider said in 1983 that the PDP-1 opened him up to ideas about how people and machines might operate in the future, but he did not think it could ever become economically feasible to give everyone a computer. However, "it did occur to him that these new computers were excellent candidates for the super-mechanized libraries that Vannevar Bush had prophesied. In 1959, he wrote a book entitled *Libraries of the Future*, describing how a computer-based system might create a new kind of 'thinking center.'"<sup>149</sup>

It is interesting to see that Licklider's dreams, like Bush's before him, bear a strong resemblance to the dreams of modern-day digital humanists. He wished he could develop a kind of computation that was more of a dialog between human and machine, an aid both in plotting and formulating models, an aid in managing complexity:

The information processing equipment, for its part, will convert hypotheses into testable models and then test the models against data (which the human operator may designate roughly and identify as relevant when the computer presents them for his approval). The equipment will answer questions. It will simulate the mechanisms and models, carry out procedures, and display the results to the operator. It will transform data, plot graphs, ("cutting the cake" in whatever way the human operator specifies, or in several alternative ways if the human operator is not sure what he wants). The equipment will interpolate, extrapolate, and transform. It will convert static equations or logical statements into dynamic models so the human operator can examine their behavior. In general, it will carry out the routinizable, clerical operations that fill the intervals between decisions. In addition, the computer will serve as a statistical-inference, decision-theory, or game-theory machine to make elementary evaluations of suggested courses of action whenever there is enough basis to support a formal statistical analysis. Finally, it will do as much diagnosis, pattern matching, and relevance recognizing as it profitably can, but it will accept a clearly secondary status in those areas.<sup>150</sup>

Licklider imagined a partnership between man and machine, not the master-slave relationship often present in science fiction. He spoke of "man-computer symbiosis" on the model of biological symbioses that see two species existing in mutually supporting harmony.

---

<sup>149</sup>Howard Rheingold, *Tools for Thought*, MIT Press (2000), chapter 7. Available online at <https://www.rheingold.com/texts/tft/07.html#Chap07>

<sup>150</sup>Quoted in Howard Rheingold, *Tools for Thought*, MIT Press (2000), chapter 7. Available online at <https://www.rheingold.com/texts/tft/07.html#Chap07>.

He quoted the fig wasp of the family Agaonidae that lives in the ovaries of the fig tree: neither the fig nor the wasp can exist without the other. Perhaps a similar future awaits humanity, a symbiosis between men and machines, beneficial to both.

## **Douglas Englebart: Augmenting Human Intellect**

Another of Bush's readers was the young Douglas Englebart, who read "As We May Think" while serving in the Philippines as an engineer in the US Navy. The article had a profound impact on him, as can be seen in his subsequent work. In October of 1962, Englebart published an internal report at the Stanford Research Institute under contract with the Air Force, entitled "Augmenting Human Intellect: A Conceptual Framework".<sup>151</sup> The article explicitly quotes long passages from Bush's article, and extends the mechanism to the improved possibilities of the electronic computer.

In his 1963 paper at SRI, Englebart mused that

we might imagine some relatively straightforward means of increasing our external symbol-manipulation capability and try to picture the consequent changes that could evolve in our language and methods of thinking. For instance, imagine that our budding technology of a few generations ago had developed an artifact that was essentially a high-speed, semiautomatic table-lookup device, cheap enough for almost everyone to afford and small enough to be carried on the person. Assume that the individual cartridges sold by manufacturers (publishers) contained the lookup information, that one cartridge could hold the equivalent of an unabridged dictionary, and that a one-paragraph definition could always be located by the average practiced individual in less than three seconds. What changes in language and methodology might not result? If it were so easy to look things up, how would our vocabulary develop, how would our habits of exploring the intellectual domains of others shift, how might the sophistication of practical organization mature (if each person could so quickly and easily look up applicable rules), how would our education system take advantage of this new external symbol-manipulation capability of students and teachers and administrators?

The example is cogent, since just such a device now exists—with a host of associated problems that Englebart had not dreamed of. The "publishers" of the "cartridges" he imagined are now multinational corporations who make much more money by manipulating knowledge and access, and through advertising, than they do from the specific intellectual product their users use them to find. Information, by becoming so quantitatively overwhelming, has become hard to navigate.<sup>152</sup>

---

<sup>151</sup>[https://dougengelbart.org/pubs/papers/scanned/Doug\\_Engelbart-AugmentingHumanIntellect.pdf](https://dougengelbart.org/pubs/papers/scanned/Doug_Engelbart-AugmentingHumanIntellect.pdf)

<sup>152</sup>A technical antidote seems necessary for this too: which is perhaps but the latest version of original sin playing out in the human field of action.



Englebart, something of an idealist, dedicated his life to helping humanity deal with complexity. His dream was to create tools to augment human intelligence, and he spent his entire life following this program. A lot of what he wanted to do was to help people collaborate, and many of the tools he dreamed up in the 1960s are common features of the latest productivity suites today—shared files, collaborative editing, remote video conferencing, links between texts...

Englebart also paid close attention to the gestural mechanisms for interaction between humans and machines, and is the inventor of the computer mouse as well as several other aspects that are familiar to computer users today—pixelated bitmap screens, real-time video communication, and the beginnings of graphical user interfaces, which he showcased in what must have been a truly remarkable demonstration in San Francisco in 1968, and which later became affectionately known as “the mother of all demos”.<sup>153</sup> To a contemporary viewer, the technology looks crude—but it is remarkable that years before these technologies began to be widespread, they were already in existence, born from Englebart’s fertile mind.

The demonstration laid the groundwork for technologies that Xerox developed in the early 1970s at the Palo Alto Research Center, and which predate the more famous products produced by Apple and later by Microsoft. Xerox was unable to market its machines, which were still too expensive for mass market use, and management pulled the plug from their research group. Apple licensed some of their inventions in their first personal computers to use a graphical user interface (GUI) and a mouse. Later, Microsoft Windows made the GUI technology ubiquitous.

Englebart and Bush’s real dream did not lie merely in the slick and intuitive interface that large tech companies later capitalized upon. Their more profound contribution lay in their understanding of memory, archival, and retrieval. In a 1986 presentation, Englebart described a two-part augmentation system that contains a human system and a tool system. He states that every skill a human has is composite: and parts of the skill can be outsourced to a tool. “We have to work on boosting people’s capability—but the technology side alone won’t make all the

---

<sup>153</sup>The video is available at <https://www.youtube.com/watch?v=yJDv-zdHzMY>. It is well worth watching, to see the future present in the past, and for a taste of Englebart’s gently humorous presentation of world-changing inventions.

difference”.<sup>154</sup> Englebart’s hope was to structure the relationship between humans and tools and permit a “co-evolution” between the two. As machines are able to help humans work faster and better, humans will make better machines, and so on, “bootstrapping” an accelerating evolutionary process.

## **Ted Nelson: the Hyperlink Controversy**

Ted Nelson was a life-long close friend of Englebart’s, and gave the eulogy for him after his death in 2013. On that occasion, he called his friend a “luminous innocent”, and insisted that the real point behind augmentation technologies was structure and collaboration, not the bells and whistles of typesetting and WYSIWYG.<sup>155</sup> “For Doug, that great demo was only the beginning... square one of the great workplace of sharing understanding that he sought to create, and that only he could imagine.”<sup>156</sup> Computers promised more than the mere simulation of paper. They could help structure discourse and improve collaboration, they could help keep track of the connections between what necessarily exists separately in the world of paper. “Documents do not exist in isolation, but we present them in isolation.”<sup>157</sup> Computers make possible the visible representation of parallel content, which according to Nelson we should be able to see not only in successive chapters, but also on the margins of a text. Comments, outlines, sources, and summaries should be presented graphically on the side of the system, in order to allow the reader to connect several streams of relevant information.

Nelson developed several demonstration versions of a system called Xanadu that instantiates these ideas.<sup>158</sup> The term “hypertext” was first coined by Nelson in the 1950s, but he meant something different from what we now know by the term.<sup>159</sup> The essential point for

---

154Minute 55 of <https://www.youtube.com/watch?v=sG3PWet8fDk>

155See 1.30 of <https://youtu.be/yMjPqr1s-cg>. WYSIWYG is an acronym for “what you see is what you get”, as in modern word processors which show a typeset version of text, not only the text itself.

156Ibid.

157See 3.40 of <https://youtu.be/DRGMFd7ue8k>

158The project never made it to production. A demo and some background information exists at <https://www.xanadu.net/>. Nelson explains his project in detail in a series of YouTube videos. The first is at <https://www.youtube.com/watch?v=hMKy52Intac>.

159There are important distinctions between Nelson’s idea of hyperlinks and Berners-Lee’s: the former insists that links be encoded *outside* of the source files they refer to, while the hyperlinks used by Berners-Lee, ubiquitous in today’s internet, are “jump links” and point from within a document to another document outside. The distinction is important, but would take us too far afield here. One author believes that both Englebart and Nelson arrived at the idea independently around the same time. See Rheingold, *Tools for*

Nelson (and Englebart) is that documents be represented as a list of pieces to bring in, instead of as a complete document containing links. Nelson's links were published separate from their underlying source, so that the source can be reused in other documents without requiring any modification.

It may seem like a minor difference, but the difference between creating a protocol which interlinks pieces of the web not only by way of blind, unidirectional hyperlinks as we know them today, and a global unified system of archival, reference, and typesetting is very great. Nelson is probably right in his insistence that in the choice for "jump links", we missed the opportunity to create a truly electronic form of publishing. Instead, according to Nelson, we merely copied paper, be it in .doc or .pdf form, or any other format.

But Nelson appears to not notice that the same effect that his "xana-links" aim to reach, can be realized in other ways too. Tim Berners-Lee et al. recently proposed (2008) a logical framework built upon URIs, which are references to name anything from a document to a part of a document to a concept or an object. With such a referencing system in place, bi-directional linking becomes an easy step. And over the last few decades, some of the other points Nelson insists upon have become part of popular programming and writing environments. For instance, GitHub has a collaboration tool that bears many points of similarity with Nelson's philosophy of collaboration and document "forking".

## **Tim Berners-Lee: Connecting Computers**

At the same time as Englebart was building his human-augmentation devices and Nelson was developing his thought about hyperlinks, Tim Berners-Lee was working as a contractor at CERN in Switzerland. In 1980, he wrote a program for his own personal use that had some of the characteristics of the memex, in that it encoded information in such a way that links could be created between "snippets" of information.<sup>160</sup>

In 1989, again at CERN, Berners-Lee wrote to his bosses and proposed the creation of a network of the computers at the research site, using his system of hyperlinked texts, and a

---

*Thought*, MIT Press (2000), chapter 14.  
160<https://www.w3.org/History/1989/proposal.html>

shared transmission protocol that would allow different machines to communicate.<sup>161</sup> His proposal was accepted and gave birth to what we now know as the Internet.

In his 1989 proposal, Berners-Lee listed a series of practical requirements necessary for their needs: the system needed to allow “remote access across networks” and also to support access to the same data from a variety of systems. It also needed to be non-centralized, since “information systems start small and grow. They also start isolated and then merge. A new system must allow existing systems to be linked together without requiring any central control or coordination.”<sup>162</sup> The proposal suggests that “if we provide access to existing databases as though they were in hypertext form, the system will get off the ground quicker”. While this emphasis on open access leads to security issues, Berners-Lee suggested that security was not the primary issue for CERN, and that information access was more important. This tension foreshadows future developments in the web.

Berners-Lee’s proposal continued by proposing that users be able to add private links to public information, along the lines of what we call “bookmarks”, and also be able to annotate public links privately. Annotation is a less-developed feature of the current world wide web, although it has been repeatedly proposed and many systems exist to implement it.<sup>163</sup> For Berners-Lee, writing in 1989, the possibility to use graphics was an “optional extra”, and display on 24x80 character screens was sufficient. It is remarkable how quickly those specifications would seem inadequate to users all over the world!

The most interesting part of Berners-Lee’s proposal, perhaps also the most nebulous in his time and in ours, but which promises to offer important insights in the future, is the “data analysis” possibility he tacked on as an appendix to the main specifications his system needed to have. He wrote that

An intriguing possibility, given a large hypertext database with typed links, is that it allows some degree of automatic analysis. It is possible to search, for example, for anomalies such as undocumented software or divisions which contain no people. It is possible to generate lists of people or devices for other purposes,

---

161 See an overview at <https://webfoundation.org/about/vision/history-of-the-web/>.

162 Ibid.

163 Perhaps the best-implemented system is called Hypothes.is, developed in large part by Jon Udell, which creates an annotation layer on top of the entire web and supports a variety of useful actions. See [www.hypothes.is](http://www.hypothes.is) for an overview, and Udell’s Github page for source code for many Hypothes.is extensions: <https://github.com/judell>

such as mailing lists of people to be informed of changes. It is also possible to look at the topology of an organisation or a project, and draw conclusions about how it should be managed, and how it could evolve. This is particularly useful when the database becomes very large, and groups of projects, for example, so interwoven as to make it difficult to see the wood for the trees.<sup>164</sup>

The structure of a network contains useful information about the contents of the network itself, and can be leveraged to do useful tasks like create automated email lists and organigrams of how people and institutions are connected. This is increasingly useful as structures become so large and dense as to be impervious to intuitive understanding by a single human being. The intuition is the same as what I presented in the last chapter, regarding the comparison between Shuila prayers and Psalms.

The next affirmation by Berners-Lee is even more interesting:

In a complex place like CERN, it's not always obvious how to divide people into groups. Imagine making a large three-dimensional model, with people represented by little spheres, and strings between people who have something in common at work. Now imagine picking up the structure and shaking it, until you make some sense of the tangle: perhaps, you see tightly knit groups in some places, and in some places weak areas of communication spanned by only a few people. Perhaps a linked information system will allow us to see the real structure of the organisation in which we work.<sup>165</sup>

He was describing what mathematicians call a “graph”, an abstract form of data representation that encodes both “nodes”, which each represent some entity, and “edges” which encode the relations between entities. Like “little spheres” connected by “strings”, the institutions, computers and users across the world could be connected. The world wide web was born. Analyzing this web could reveal further information: “shaking the graph” reveals that some nodes are weakly connected; others form locally dense regions.

Within the domain of history, a variety of techniques are applied to social network graphs that encode the relations between people, institutions, and trade routes to study the movement of technologies and materials, and derive considerations about the cultures encoded in those networks. Within the financial industry, these networks are studied to discover fraud rings, which are evident by the special topology of relations between a relatively small group of actors in the network.<sup>166</sup> There are many mathematical techniques of varying degrees of

---

<sup>164</sup><https://www.w3.org/History/1989/proposal.html>

<sup>165</sup>Ibid.

<sup>166</sup>Tahereh Pourhabibi, Kok-Leong Ong, Booi H. Kam, Yee Ling Boo, “Fraud detection: A systematic literature review of graph-based anomaly detection approaches”, *Decision Support Systems*, Volume 133, (2020).

complexity that aid in other analyses of complex networks to which we will return. For now, the main point is that by connecting data in a web, and by representing that web as a “graph”, a new feature of the information is encoded which can be queried and interpreted: the topology of the graph itself.

Berners-Lee has had an abiding interest in “knowledge engineering” that would augment human capacities to handle exponentially larger amounts of information. Recently, he worked in particular on the development of logical ontologies such as RDF and OWL, which are the basis for the semantic web, a not yet fully realized technology that encodes not only text, but also logical connections between ideas, in a computer readable format.<sup>167</sup>

## Building a ‘Memex’

With this background in place, we can now approach the definition of the “intellectual prosthetic” that is the subject of this thesis. It bears some similarity to the ‘memex’ described by Vannevar Bush, so I will use that name as shorthand to refer to my device.

Reading about archaeology introduced me to a basic problem: the archaeologist must patiently sift, interpret, and archive data long before an explanatory structure has been validated. The interim time, which can last years or decades, requires the creation of a “grammar” for cataloging and storing the data, both physical artifacts and mental constructions.<sup>168</sup> During that time, the scholar’s interpretation grows, is enriched and changed, perhaps many times before expression in a publication. And even then, more data is continuously coming out of the ground, and the scholar must revise, improve, and sometimes scrap their version of the facts as new information comes to light.

The mass of information must therefore be queried at a much later date as well as interpreted *in corso d’opera*, day by day. These facts introduce a specific challenge for the human memory, and a specific advantage that Buccellati develops at length in his study: the

---

<sup>167</sup>See Sinclair Target, “Whatever happened to the semantic web?”, *Two-bit History*, 2018. Accessed on 29 March 2024, <https://twobithistory.org/2018/05/27/semantic-web.html#fn:10>.

<sup>168</sup>My main reference point in this field is Buccellati, *Critique of Archaeological Reason*, Cambridge University Press (2017).

fact that archaeology is “born digital”.<sup>169</sup> I will not develop this theme here, but I take his conclusion as my starting place.

Fragmented information, re-discovered and brought to light in a partially haphazard manner, must be patiently re-connected before it can be interpreted. A pile of sherds looks like rubbish at first, but once recomposed into the shape of an amphora or a bowl, the sherds regain their original meaning, and the place they were found can be interpreted more fully. Archaeologists have dedicated much work to the problem of how to catalog the sherds as they come out of the ground, and how to archive them in such a way that they can be recovered and used later on in the work of interpretation.<sup>170</sup>

I wondered if a similar process could be applied not only to the physical fragments in a dig, but also the intellectual fragments of thought—sentences and paragraphs—that a scholar encounters, ponders, and files away daily. There is a similarity in both settings. One important task of the archeologist is to recompose wholes that had been fragmented, and which are unknown to any living carrier of the tradition.<sup>171</sup> The historian recomposes the story of what happened in the past, what causes and effects events had, and sometimes attempts to understand what lessons can be learned from the past and applied to the present. The problem is that the information is myriad, and historians live but for a brief moment in time, subject to all the usual pressures and restrictions of daily life. Worse, they cannot be sure of knowing the most relevant information for the problems posed: they have access to great libraries and stand upon the shoulders of the giants who preceded them, yet even so there is no lack of overturning the paradigms, critiquing the “giants” and finding them in error due to one or many blindnesses. While we may congratulate ourselves on each blindness successfully overcome and unfolded into the light of reason, we must also wonder what are the blindnesses of our own that the next generation will have to uncover, acknowledge, and heal.

Human consciousness has developed the ability to reduce the overwhelming multiplicity of data in ways that remain mysterious to cognitive scientists. Scholars can be quite effective at

---

169 See Buccellati, *Ibid.*, pp. 232-237. See also more recent related articles at <https://d-discourse.net/>.

170 See the overview in Buccellati, *Ibid.*, chapters 2 and 3.

171 See the interesting article by Buccellati and Bezzera “Broken Traditions” at <https://critique-of-ar.net/themes/broken.htm>.

intuiting a good path through the data.<sup>172</sup> No doubt, current and near-future AI systems will be even more flexible and capable of solving hard problems. But even if machines were to become able to fully replace human intellectual processes at some point in the future, and navigate the difficult and uncertain territory of historical interpretation better than humans (let us say), at present we are in the intermediate territory of having too much information to successfully ingest, understand, interpret, and use. Today, we need ways to augment human intelligence.

The use of computers to aid the historian is not a new idea. Ignace Gelb was one of the pioneers of this methodology in his computer-aided study of the Amorite language, published at the University of Chicago in 1980. In the intervening years, other tools have been developed, and many experiences have shown that digital methods enhance the work of the historian. Buccellati has extensive experience both in the creation and deployment of database structures for archaeology, and in the development of publication platforms that make use of digital methods to present a “multi-linear discourse”.<sup>173</sup> Many other publications explore further aspects of the contribution that digital and statistical methods can offer the historian. *Pars pro toto*: Furet (1971)<sup>174</sup>, Burdick et al (2012)<sup>175</sup>, and Kim (2018)<sup>176</sup>.

In recent years, archaeologists have begun using network methods at an increasing rate. While until 2006 less than five papers a year were published on this issue; in 2014 more than 25 papers were published on network methods in archaeology.<sup>177</sup> Interest has continued to grow since then. As Anna Collar put it,

A key question for archaeologists interested in using networks is whether and how their data can be represented as nodes and connections between them, or edges. But why would we want to represent our archaeological data as networks anyway, and why should using a network science approach tell us something about the past that other approaches could not? Underlying these questions is the idea that using

---

172 These brain heuristics are being studied by artificial intelligence researchers and cognitive scientists. Several game-playing devices, culminating recently in Alpha Go, have proved able to deal with the combinatorial explosion on time scales that make a game-playing computer feasible.

173 See the 4banks.net group of websites.

174 Francois Furet, “Quantitative History,” *Daedalus: Journal of the American Academy of Arts and Sciences* 100 (1971): 151–67.

175 Anne Burdick, ed., *Digital Humanities* (Cambridge, MA: MIT Press, 2012).

176 Dorothy Kim, *Disrupting the Digital Humanities*, 1st edition (Santa Barbara, CA: Punctum Books, 2018).

177 Anna Collar et al., “Networks in Archaeology: Phenomena, Abstraction, Representation,” *Journal of Archaeological Method and Theory* 22, no. 1 (March 2015): 3, <https://doi.org/10.1007/s10816-014-9235-6>.



network methods allows us to do something we could not do before, something different from “standard” archaeological practice, which will reveal new information about our data.<sup>178</sup>

Collar et al. point out that a major reason for using a network model is to be able to represent the relationships between the actors in a network.<sup>179</sup> While it may seem that relationships have always been central in the study of history, attempting to model a numerically large amount of information in a rigorous way such as Franzosi describes in his work (2017) is a new development. It is within the tradition of structuralism as used in anthropology—the logical formulation of relations between actors is reminiscent of Kingsley Davis and Lloyd Warner’s work on kinship structures (1937), and Levi-Strauss’s later work on the same theme (1949). Collar et al. note that “network science methods incorporate techniques that are already frequently used by archaeologists, or that are an element of commonly used methods. For example, a Harris matrix can be considered a network representation of the theoretical assumptions known as the laws of stratigraphy.”<sup>180</sup> They also point out that

Network science is not a single, monolithic entity, but denotes a diverse set of methods, models, and approaches concerning the study of the management, representation, and analysis of network data which represent our hypotheses about how and why relationships matter. It is not limited to the analysis of networks or the study of social networks, nor is it limited to the representation of data, nor to the fact that it offers researchers new ways to phrase research questions. The central potential of network science for archaeology is that it places relationships at the heart of our analytical techniques.

One major difference between what was possible for Levi-Strauss to carry out in 1949 and Collar in 2015 is the quantity of information that can be managed and analyzed using computers. In his day, Levi-Strauss was criticized for “threatening people with mathematics”.<sup>181</sup> Perhaps today we are used to the idea of mathematical tools being applied to humanistic data, and it is no longer felt as threatening. Instead, many hope to be able to apply rigorous transformations to our information and thereby derive new insights.

In archaeological data, “sites or assemblages of material culture form natural nodes”.<sup>182</sup> She goes on to explain that

---

178Collar et al., 4.

179*Ibid.*

180Collar et al., 16.

181Stanislav Andreski, *Social Sciences as Sorcery*, 3. impr (London: Deutsch, 1974), 72.

182Collar et al., “Networks in Archaeology,” 10.

A very common approach [...] is to use sites as nodes. Sites form natural nodes because of their relative boundedness, discreteness, and stability and persistence over archaeologically observable timescales, as well as their common use by archaeologists as analytical concepts. They offer the opportunity for mesoscale analysis of interactions: probably the level at which archaeologists most often work, due to the diachronic nature of the archaeological record and a historic interest in systemic level processes.<sup>183</sup>

These understandings of what data to represent as a node in the network are being extensively explored on the OCHRE platform at the University of Chicago.<sup>184</sup> Physical objects and persons are easily understood as nodes in relation with other objects, places, and persons. Ego networks can also be constructed by attending to individual actions and interactions, where it is possible to identify such persons.

In Assyriology, some explorations of applying social network analysis to cuneiform texts have already been carried out. One such project, realized by Tero Alstola et al. of the University of Helsinki (2019), analyzes a corpus of 1532 texts from the Neo-Assyrian period obtained from the Open Richly Annotated Cuneiform Corpus (Oracc). The scholars study the role of the god Assur in the Mesopotamian pantheon, using computational methods. Their paper concludes, on the basis of the network analysis, that “Aššur is not a very central god in our corpus despite his importance in Assyrian royal theology, but he rather joins the existing networks of gods without altering them.”<sup>185</sup>

Their method is somewhat brittle, since it defines a connection between gods simply by co-occurrence within a ten-word window of text.<sup>186</sup> Such a method can indeed churn quickly through a large corpus, but the metric misses important features of the relationship between gods. Literary texts often create powerful oppositions between characters by dealing with them in turn—at more than ten-word distances—and their rare meetings are thus freighted with the many previous pages which have given each character personality and depth. Co-occurrence might sometimes work as a representation of relation between players in a theatrical dialog (if

---

183Collar et al., 13.

184See <https://voices.uchicago.edu/ochre/project-gallery/>

185Tero Alstola et al., “Aššur and His Friends: A Statistical Analysis of Neo-Assyrian Texts,” *Journal of Cuneiform Studies* 71 (January 2019): 159, <https://doi.org/10.1086/703859>.

186“First, we produce networks of gods by looking at their co-occurrences in a text. The resulting network shows which gods were relevant enough for the ancient author to mention them in one and the same text. Second, we analyze co-occurrences in a ten-word window, meaning that two gods are connected if they are separated by no more than eight other words in the text. This allows us to map co-occurrences in more lexical semantic contexts and produce networks that show which gods are actually mentioned together in our corpus.” (160)

they do not co-occur, they do not speak), but it would certainly miss the incandescent tension between Jean Valjean and Javert in Victor Hugo's *Les Misérables*.<sup>187</sup>

For historians, a strong predilection for material culture, or an overly rigid arithmetic definition of the salient feature under examination, might not be adequate as a representation of the information they wish to analyze. Instead, it might be more important to put complementary or contradictory statements from several persons in relation, as when combining eyewitness accounts with newspapers and judicial reports, as Franzosi did when studying the Georgia lynchings.

My own experience in a research team convinced me of this point. Nearly all the information we dealt with was secondary sources: a large body of text, running in the hundreds of volumes and the millions of words. An obviously significant, but much smaller, body of information is the primary sources in textual form. Finally, the research team itself generated a sizable body of text, from brief marginal comments to short articles in length. What are the “sherds” in this pile of words? Where are the jagged edges that could be matched to their corresponding pieces, to reconstitute wholes?

## Building Text Networks

If the ‘battle for synthesis’ is a central task for the historian, as Oppenheim exhorts, then new tools are needed.<sup>188</sup> Perhaps—and this is the burden of the rest of this chapter—the idea of connecting text information into a graph structure is a viable and useful approach, but the algorithmic detail about how the graph is created must be improved.<sup>189</sup> One way to construct a

---

187The example is salient in Digital Humanities: a model of the relations between characters in *Les Misérables* is one of the first force-directed graphs I encountered. See <https://observablehq.com/@d3/force-directed-graph>.

188This aim is not unproblematic, in the sense that it points to “grand narratives” and a traditional view of the goal of knowledge-acquisition as ‘truth’. Both of these aims have been criticized, and in my first chapter I situated my own goals within the existing debate.

189In the next chapter, I describe the ‘multinodal index’ as one possible way to augment the scholar’s ability to access text information from an alternative point of view. This technique is admittedly crude, since it is based on keywords which contain no information about their context. Furthermore, the branching network structure of the multinodal index is created by a naive threshold value for shared keywords: more than five shared words and a connection was created. Less than five, and a node remains isolated. This heuristic provides an easy way to create a network representation, and it does return useful data in some cases. However, it is easy to see that many situations will not meet the threshold criterion and nevertheless be significant.

graph from text information has a significant history already: citation networks. Some researchers in the digital humanities construct networks that represent the authors who cite other authors, and attempt to derive interesting conclusions from the resulting structures.<sup>190</sup> One robust result from the clusters shown by these studies is the fact that a relatively small number of researchers are cited by most of the other authors. The distribution of high-index authors and low-index authors roughly follows a Pareto distribution: about 20% of the authors receive 80% of the citations.<sup>191</sup> To those who have, more will be given.

Citation networks are a powerful way to represent the social graph associated with a scholarly field and identify the central nodes that control the traffic of the rest of the network. Citation networks can be created algorithmically, since it is possible to parse the wide variety of bibliographic information in published articles. However, they do not contain reliable information about the actual contents of the articles cited. It is possible that authors cite articles they have not completely read, cite secondary affirmations of an article, or even simply pad their bibliographies with famous articles that are expected in a given field, even if they are not particularly useful in the context. For these reasons, we should view citation networks more as a representation of the social graph related to a field of study, and less as a representation of the content of that field of study. It has happened more than once in the history of science that the most highly-cited scholars were wrong, and a new discovery had to break through the layers of ossified social networks before being properly understood and accepted.<sup>192</sup> Social networks can both help and hinder the discovery of truth.

Another way to create a network representation of knowledge depends on manually created links between ideas. This is the case for instance on the world wide web, where authors create

---

190 See for instance Jason Portenoy, Jessica Hullman, and Jevin D. West, “Leveraging Citation Networks to Visualize Scholarly Influence Over Time,” *Frontiers in Research Metrics and Analytics* 2 (November 27, 2017): 8, <https://doi.org/10.3389/frma.2017.00008>.

191 Nunhes, Thaís & Oliveira, Otávio. (2018). “Analysis of Integrated Management Systems research: identifying core themes and trends for future studies.” *Total Quality Management & Business Excellence*. 31, p. 12.

192 Think for instance of the tragic case of Ignaz Semmelweis, who realized the importance of handwashing to prevent sepsis in obstetrics, but was ridiculed and ostracized by the medical establishment, which thought that “a gentleman’s hands could not transmit disease”. See a popular account by Francisco Domenech at <https://www.bbvaopenmind.com/en/science/leading-figures/the-man-who-discovered-that-handwashing-saves-lives/>.

robust connections to other ideas on their own site, and on other sites, across the world. Sometimes these links act like classic footnotes; in other cases they take on new characteristics of ‘digital writing’, such as the innovative project underway at d-discourse.net.<sup>193</sup> The success of the Google search engine depends on the intuition that the network structure of references between websites is a powerful heuristic for identifying valuable information: roughly speaking, the more people link to a site, the more likely it is that that site contains valuable information, and the higher it will be ranked on the search engine.<sup>194</sup> This ‘popularity contest’ has proved without doubt to be a powerful and useful way to organize information.

But what about domains where there is little or no broad public interest, and therefore no manually coded hyperlinks? What about old articles, written before the web was a reality, or articles behind paywalls? The Google approach cannot work well in this situation. And for the domain of historical research, most of the information available has the above-mentioned weaknesses.

I propose that there is another approach to the algorithmic creation and analysis of a network representation of historical information, specifically the secondary sources that make up the bulk of historiography. This network tool serves the purpose of automatically connecting the existing information in such a way that scholars can efficiently find relevant information and map the historiographical development of thought in their domain. The network tool also serves the purpose of broadening the view they have, while ordering the information thus surfaced, permitting them to take more source material into account in the ‘battle for synthesis’. Finally, the network tool permits several layers of semantic information to be generated, visualized, and accessed.

In order to create such a tool, it will be necessary to define kinds of text information on which we will principally focus. I also will need to specify how these texts will be represented as nodes in a network, and on what basis nodes will be connected. Then, I will need to test a variety of ways to query the network and interpret the results. In the rest of this chapter, I will

---

193Wikipedia is another example of a manually coded form of digital writing that contains useful information in the structure of hyperlinks.

194See the patent US7058628B1 – Method for node ranking in a linked database, and the original article by Larry Page and Sergei Brin, “The anatomy of a large-scale hypertextual Web search engine”, *Computer Networks and ISDN Systems*, 30(1-7): 107-117.

focus on the technical methods for creating the network. Testing and interpretation will be the focus of the next chapters.

## What is a node?

One very simple definition of nodes, potentially more adequate to the work of a historian than objects and persons, can be found by representing keywords in a text—lemmatized nouns—as nodes.<sup>195</sup> These words give a rough heuristic (a “bag of words”) for the content of various texts, and can be used to connect subdomains of a “thought space” that describes a historical problem. However, they also have an important limitation in that they do not account for synonyms or near-synonyms, nor do they contain logical information regarding negation or any other relation between the words in a text. Keywords and lemmas are a poor language model.<sup>196</sup>

Another more sophisticated way to think of nodes is to take a larger piece of text as the “atom”, which are composed of words, as in the physical world atoms are composed of subatomic particles. This is a common choice: digital humanists often work with “triples” formed of three adjacent words.<sup>197</sup> Sometimes these words are algorithmically selected in order to create structures such as subject-verb-object. In this case, the parts of speech could be likened to subatomic particles which form meaningful structures (“atoms”) when taken together. As we saw in Chapter 2, the current performance of language models makes this approach somewhat unreliable. It is not necessarily possible to automatically parse text into the syntactic structures that express its content.

---

195The number of such nodes rises rapidly with the length of the text, since nouns are the most common part of speech in English, and typically make up between 15 and 30 percent of the words in a text. In our case, since the keywords in the multinodal index are taken from the short summaries at the 4Banks website, their number is not excessive. By extrapolating from the *Longman Grammar of Spoken and Written English* we find that nouns are 15% of spoken English, and 30% of academic English prose. See the discussion at <https://english.stackexchange.com/questions/55486/what-are-the-percentages-of-the-parts-of-speech-in-english>

196Still, it must be noted that TF-IDF, which is based on a simple arithmetic of keywords in a text, remains one of the most powerful and accurate ways to represent a text without using the much more computationally expensive word vectors (such as Word2Vec) we use here. It is curious that TF-IDF works so well, and several explanations of this fact have been advanced. See one discussion in Robertson, “Understanding Inverse Document Frequency.” *Journal of Documentation*, 60:5 (2004), pp. 503-520.

197For instance, Tero Alstola et al., “Aššur and His Friends: A Statistical Analysis of Neo-Assyrian Texts,” *Journal of Cuneiform Studies* 71 (January 2019).

The interplay of closeness and distance in the meaning of sentences provides us with a clue as to how this information could become useful. The closeness of two sentences might signal a cluster of affirmations by different scholars who say the same thing. An algorithm capable of surfacing similar phrases and clusters of similar ideas would be useful for searching for supporting evidence. Or, the closeness of the two phrases might indicate branching, a place where two large-scale interpretations differ at a fundamental point that affects everything downstream. It would be useful to have a way of remembering that such an opposition exists and at a future time to be able to traverse the connections that lead from the data to the affirmation, and from the affirmation to its consequences in interpretation.

It is important to grasp the difference between this approach and the search engines we are accustomed to using. Older ‘search engines’, library card catalogs, provided a rough way to find relevant information based on a variety of categories. Indices and footnotes provided a more granular approach, where scholars could indicate topics in their indices, and point the reader to supporting page-level information in footnotes. Internet search engines and regular expression keyword indices for electronic databases provided a further level of granularity, and enormously expanded the domain of searchable information. However, this expansion came with the cost that the source information was returned to the user with a priority ordering that was not necessarily an aid in study. Furthermore, the user of the systems cannot interact with them and improve them with use. There is no way for a scholar to improve the search system itself, although many do create their own personal methods of archival and note-taking.

My approach aims to improve both the way information is located, and the way its quality and ordering is improved over time through use. In my approach, sentences are the basic node of the network, the “atoms” of thought.<sup>198</sup> Instead of considering entire books or articles as the fundamental division of the information, search is conducted on the far more granular affirmations made sentence by sentence. It is therefore more likely that the algorithmic procedure surface information that is relevant, both because of its similarity to the search topic, and because it might use similar wording but have an opposite meaning. In the second case, the

---


<sup>198</sup>Larger nodes could also be contemplated, such as paragraphs or entire articles and books. However, word vectors lack the ability to specify the relation between two longer works in detail.

researcher can add a layer of semantic interpretation to the network in order to point future researchers toward this decision point where two scholars differ in their interpretation. Later in this chapter I will show how this can happen.

First, I need to detail in what way the network can be created from the sentence-nodes.

## What are relations?

The most basic form of relations between sentences, or nodes, in my network is the linear sequence given by breaking a complete text into a series of sentences. While my approach emphasizes the importance of the individual sentence as an “atom” of discourse, it must also be possible to consider these atoms in their existing context: the linear sequence that starts at the beginning of a text and progresses through to its end. Therefore, my program encodes sentences as nodes that are connected by “PRECEDES” relations. A paragraph of three sentences numbered 0, 1, and 2 could be represented by a string of nodes as follows, connected by “PRECEDES” relations. Note that these relations also have a direction, which encodes the order in which the sentences form the text:




This structure preserves the overall unity of the paragraph, but also permits each sentence to be related to other sentences.

The second type of relation between nodes in my system is their ‘similarity’. This is a numerical value that can be calculated in a variety of ways. Two sentences may share some or many keywords. Based on this fact, we could connect sentences with a “strength factor” (as we did in the multinodal index) to express the number of words shared between them. This creates a type of internal relation in the graph which could help orient the researcher in her attempt to



understand which information is most relevant for the project at hand within a large collection of data.

We can now imagine another kind of relation in the graph:



Sentences 870 and 871 belong to the same overall text, and these two affirmations have been calculated as similar to sentence 879. A calculation of similarity based on keywords would show relations only between sentences that share exact words, however. This limitation is significant, and it is time to consider the possibilities offered by language models. Although they were not sufficiently precise to answer detailed questions about the syntactical differences between two religious traditions, they can provide a measure of semantic similarity that could extend and improve our memex.

Since words and phrases are represented as high-dimensional vectors in language models, they can be manipulated using linear algebra. Word vectors can be summed to form a representation of a sentence, and two sentences can be compared by calculating the “cosine similarity”, which roughly speaking can be described as the cosine of the angle between the two vectors.<sup>199</sup> The closer an angle goes to zero, the more the two vectors overlap, and the closer the cosine of the angle approaches 1.

We can see how significant this fact is through a simple example using the mid-size model provided with SpaCy and used in our previous experiment parsing ancient texts. For instance,

---

<sup>199</sup>To speak precisely, since in general we are considering high-dimensional vectors, the definition of the “cosine of the angle” is the dot product of the vectors divided by the product of their lengths.

let us take the phrases “Mesopotamian divinities”, “Babylonian gods” and “the God of Israel”. The first two phrases refer to a closely similar reality without using the same words.<sup>200</sup> The third phrase refers to a reality that is radically different from the first two, yet it shares a keyword with the second phrase. We might expect that calculating their similarity based on keywords would give a zero result for the first two phrases, and a false positive for the second two phrases because (if lemmatized) they both contain the word “god”. Clearly, a simple keyword solution would not be sufficient.

If instead we use the SpaCy cosine similarity method, we can write a simple program as follows:

```
phrase1 = "Mesopotamian divinities"  
phrase2 = "Babylonian gods"  
phrase3 = "the God of Israel"  
n1 = nlp(phrase1)  
n2 = nlp(phrase2)  
n3 = nlp(phrase3)  
print (n1, "<-->", n2, n1.similarity(n2))  
print (n1, "<-->", n3, n1.similarity(n3))  
print (n2, "<-->", n3, n2.similarity(n3))
```

The output from this program is:

```
Mesopotamian divinities <--> Babylonian gods 0.9306257467747806  
Mesopotamian divinities <--> the God of Israel  
0.5341943005758424  
Babylonian gods <--> the God of Israel 0.6212027056639037
```

Cosine similarity is expressed as a value between 0 and 1, and higher numbers mean greater similarity. A value of 1 would be obtained if the two phrases were identical. As we can see, the program successfully identified the first two phrases as being highly similar: their value

---

<sup>200</sup>There is an important caveat even in this contrived example: Mesopotamia and Babylonia are not the same, and “divinity” and “god” do not have the same gradation of gender-neutrality. Important distinctions are lost if the terms are treated as identical. The point of the example is to show that something about their similarity can be captured algorithmically, even if it does not perfectly reflect the fine-grained information contained in the words.

is .93, which is a very high similarity, even though they do not share any keywords. Likewise, the program recognizes that “the God of Israel” is significantly different from both of the previous phrases. As we might expect, it also picks up that the shared keyword makes phrase 3 more similar to phrase 2 than to phrase 1, but in both cases it is quite a lot less similar than phrases 1 and 2.

This simple example is encouraging. One important issue we should be aware of is that the SpaCy results do not render negation with the importance that it should have.<sup>201</sup> If we re-run the program above with different phrases, we can see that its results are not trustworthy in this case:

```
Mesopotamian divinities do not exist <--> Babylonian gods exist  
0.8964916722294544
```

The score is lower than before, but it is still very high—too high to capture the stark difference in meaning introduced by the negation. This is not surprising, since the calculation of similarity treats each word as a vector, and the sentence as the sum of the vectors. One word cannot flip the whole phrase around vector space, as far as the algorithm is concerned, even though it should from a semantic point of view. The creation of a system that is truly capable of understanding natural language remains imperfect, notwithstanding decades of intense work and amazing discoveries made by the academy and industry. Even so, the recent advances in representing text as vectors offers a useful tool to the researcher, as long as we remain conscious of what its results mean and what they do not mean.<sup>202</sup>

The similarity calculation allows us to recognize phrases that contain similar content, even if they do not contain the same words. This is possible because the trained model represents each word in a high-dimensional vector space that captures semantic information by mapping each word according to its context in billions of words of English text. The similarity

---

201 There are extensions available that attempt to deal with this issue. One, written by Jenő Pizarro is available at <https://github.com/jenobjp/negspacy>. However, this solution is not yet able to modify the similarity score for the overall phrase, and instead modifies the logical value of the single words the negation is attached to. It could be included in calculating similarity, but to do so effectively would require sophisticated linguistic and programming knowledge. It remains a goal for a future version of this software.

202 Some benchmark data for performance of the SpaCy system as compared with other leading NLP packages can be found at <https://spacy.io/usage/facts-figures#benchmarks>

calculation allows us to create relations between sentences automatically. These relations will then serve the dual purpose mentioned above: first, they will provide the basis for a search function capable of surfacing semantically relevant information even where there are no shared keywords. Second, they will be able to be vetted by a human user who can judge more carefully the logical relation between them.

A problem which quickly becomes relevant in this calculation is the fact that every new sentence must be compared to every existing sentence in the database. This is a process that becomes exponentially slower as phrases are added and the database grows.<sup>203</sup> The performance issues associated with calculating cosine similarity can be a major problem in systems such as the one I am presenting, since users expect fast results from their search. However, since the calculation occurs at input time, speed is not a particularly important issue in my case. It happens in the background. While reading and inputting text, a researcher is not affected by the time the computer takes, since it usually takes more than a few minutes to find another relevant text. And in any case, new texts can be saved by the user even while the computer is still calculating the previous input. These performance issues would instead be prohibitive in a commercial search engine like Google. Where the incentives align to prefer speed to accuracy, other approaches are necessary. But in our case, where historians wish to describe and interpret the world of the past as accurately as possible, the incentive toward precision and total control of the available information can and should take precedence over the commercial priorities that determine other technologies.

The basic relations that we have seen so far encode two essential characteristics of discourse. First, they encode the linear sequence of sentences that makes up a complete thought, be it a paragraph, an article, or something longer. Second, they encode the relation

---

<sup>203</sup>After inputting about 1000 paragraphs into the system, on a 2021 MacBook Pro it already takes about 70 seconds for an average length new paragraph to be processed and calculated. This time will only increase as more information has to be processed at each input cycle. In order to run more quickly, my script makes use of “threading” to create a separate process for each new sentence. In this way a paragraph of several sentences is processed in parallel, and thus in a much shorter time than would be the case if it waited for each phrase to finish before proceeding to the next. Other optimizations are also used in my code, such as a “just in time” compiling in native machine language of the routine that actually calculates similarity, which is the most processor-intensive feature of the program. The details of computational optimization are an extremely fast-moving subject in computer science. For instance, recent developments in creating a superset language of Python called Mojo promises to dramatically improve computational efficiency.

between nodes that are semantically similar. A third kind of relation that can be created between nodes is the logical connection: “AGREES” and “DISAGREES”.<sup>204</sup> At present, it is not possible to create these relations algorithmically, although research is ongoing and may well arrive at this capability sooner rather than later.


Having logical relations saved means that a researcher will be able to recover the main features of their thought processes months or years after encoding a detail. They may not remember the specific point of disagreement between two authors, and only generally recall that one held one position and the other a different one. A tool for recalling the specific point of disagreement along with all its supporting data and consequences, would be helpful. Furthermore, a research group could share information readily and efficiently in this way. All members of the group can be given access to add to and query the graph: they can all gain from one member who encodes a new text or logical relation. New members of the group who are onboarded after work began can also benefit from such a tool, since they have direct access to the state of the art reached by the team. Not only will they have access to the bibliography that has been taken into account up to the present, which is already a great advantage, they will also have access to the logical patterns the team has already discovered.

But why stop there: another possible use of the system is faculty-wide, or even inter-faculty, databases to help students begin their research. The system could be prepared with a significant sampling of the classic works in a field, ready to be queried by students writing papers and planning dissertations. It would speed up their initial research by providing vetted, widely accepted research as their first line of inquiry, in a quickly accessible format. It would also grow with their work and keep pace with developments in the field, to the extent that the students themselves continued to input new information into the system.

We can intuit the different levels of information encoded by focusing on node 879 (seen above) and show all the other nodes it is currently related to:

---

<sup>204</sup>This relation can also be expressed as a percentage of agreement, although this representation appears more meaningful than it is. What would it mean for a researcher to assign a score of 35% agreement? In my software, I have preferred to allow only a binary choice for each user of the network. At query time, the scores made by different researchers could be averaged, and thus produce a percentage agreement.



Node 879 has disagreement relations with three nodes, and it also has agreement relations with nodes 667 and 849. At this stage, we can see that something interesting lies within the structure of this small graph. Even without knowing the content of the affirmations, we can see that statement 879 is important<sup>205</sup>: and not because it is cited by others, but because its content has semantic similarity to statements made by others, whatever its “authoritativeness” in the world of academic publishing. It mediates several different arguments contained in other documents. It could be seen as a point that directs traffic through the network: because of the logical connections, node 879 initiates a branch where certain statements become alternatives and not merely apposite affirmations. Crucially, it is not a gatekeeper in a social network, a human who controls the traffic. Rather, it is a logical node that is important because of its semantic and logical connections to other affirmations made by other authors in the network. But our gaze is directed not at the relative authority of these authors, but rather at the ‘inner-referential’ system of relations that connect the affirmations made, irrespective of who said them.

The reader can intuit that there are many distinctions that must be made at this point. The branching of the graph, if the nodes correspond to natural language sentences, will not have the

<sup>205</sup>Clarify the distinction between social graph, and text relations

exclusive character of a branching logical tree. There may very well be later statements in one of the documents that agree where previously it disagreed, or vice versa.<sup>206</sup> This means that nodes like node 879 are not strictly speaking decision points, and their branchings could be recomposed at a later point, as far as the structure of the graph itself is concerned. Logical connections encoded by a team of researchers will probably not all be the same: well-intentioned scholars can disagree.

However, my hypothesis is that such nodes will still have importance for the historian, since they will serve as a dimension reduction of the number of overall nodes he has to examine and will direct attention to pivotal points in the argument. Thus, instead of having to look at thousands of sentences, the historian could start by examining the sentences that mediate the logical branching of the graph. From there, by proceeding stepwise through the connected relations, the historian's attention could be focused on the most closely relevant data for the question at hand.

In closing this section, I also admit that there is a certain naivety in the definition of nodes and relations. The former are sentences, the latter are one of four simple types of relations. Many further distinctions can and should be made, and these categories should be expanded to include smaller and larger nodes (words, graphemes, paragraphs,...) and relations should be expanded as well. However, for the methodological viewpoint that informs this thesis, such complexity would not help us reach a clear vision of the overall method and interface. Therefore, I accept the limitations of this simple model of text fragmentation and

---

206Berners-Lee et al. (2008) state that “The nonmonotonicity of many existing systems follows from a form of negation as failure (NAF) in which a sentence is deemed false if it is not held within (or derivable from) the current knowledge base. It is this concept of current knowledge base, which is a variable quantity, and the ability to indirectly make reference to it which causes the nonmonotonicity. In N3Logic, while a current knowledge base is a fine concept, there is no ability to make references to it implicitly in the negative. The negation provided is called scoped negation as failure (SNAF) and is the ability for a specific given document (or, essentially, some abstract formula) to objectively determine whether or not it holds, or allows one to derive, a given fact. However, negated forms of many of the built-in functions are available.” Tim Berners-Lee et al., “N3Logic: A Logical Framework for the World Wide Web,” *Theory and Practice of Logic Programming* 8, no. 3 (May 2008): 249–69, <https://doi.org/10.1017/S1471068407003213>. Cycorp reached a similar conclusion, that its Cyc system must be flexible enough to allow internal contradiction, to be a viable model of reality. See interview with Cycorp founder Doug Lenat on the Lex Fridman Podcast, episode 221, 15 September 2021. Lenat's approach might be a useful correction to the weaknesses of LLMs. See Doug Lenat and Gary Marcus, “Getting from Generative AI to Trustworthy AI: What LLMs Might Learn from Cyc” (arXiv, July 31, 2023), <http://arxiv.org/abs/2308.04445>.

interconnection, and leave for the future a more sophisticated approach that builds upon the same fundamental method described here.

We now face the important issue of the interface: how can the historian query the graph effectively? How can the intricacies of the network of thousands or millions of sentences become insight?

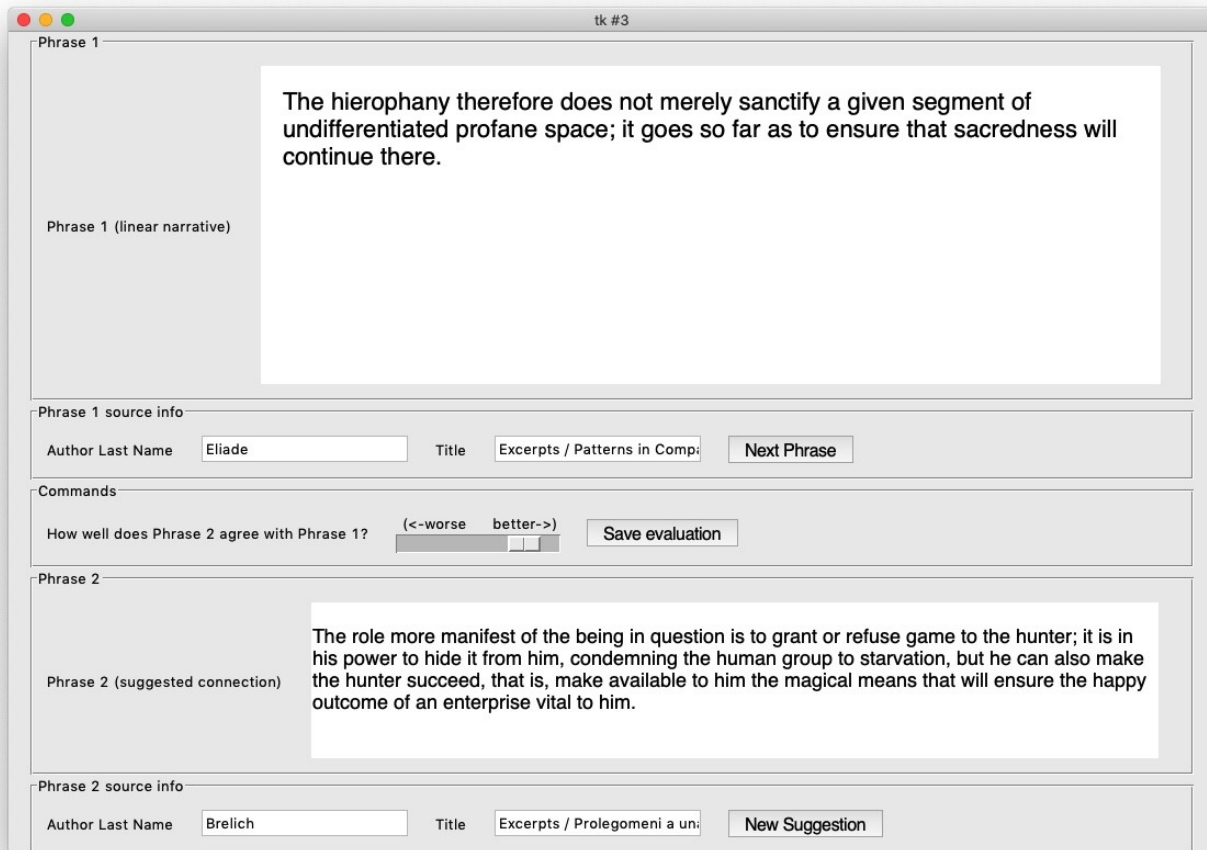
## Interfacing with the machine

The first version of my software used a rudimentary graphical interface that presented the user with two phrases and a few buttons.<sup>207</sup> The user was invited to read the two phrases and evaluate their logical relationship. After doing so, the user could set a slider to whatever approximation of agreement or disagreement they desired, and save the relation and move on to a new phrase. This interface looked like this:

---

<sup>207</sup>In this version of this software, I used a sliding scale to represent gradations of agreement. The details of such a scale would present interesting alternatives: a scale from 1 to 100 (which was my first attempt) would represent the vagueness of judgment, but it might also imply a greater degree of granularity than actually obtains in such a judgment. A low number of options, say a scale from 1 to 5, would improve this aspect and also present the possibility of a neutral judgment—a “3”—which would be difficult to interpret in a meaningful way. A scale from 1 to 4 would force the user to take sides: the relation could be weak, but it would have to lean either toward agreement or toward disagreement. The topic has been extensively studied in psychometry and sociology by Likert, Bogardus, and Thurstone, among many others. See Allen, Elaine; Seaman, Christopher (2007). "Likert Scales and Data Analyses". *Quality Progress*. pp. 64–65. See also Robert L. Armstrong, "The Midpoint on a Five-Point Likert-Type Scale," *Perceptual and Motor Skills* 64, no. 2 (April 1987): 359–62, <https://doi.org/10.2466/pms.1987.64.2.359>.





There are two ways to move through the phrases. On the top, Phrase 1 is the next phrase in the list of sentences, and clicking “Next Phrase” will move the counter forward by one, and loop through all the sentences contained in the system. Phrase 2 is a suggestion made by the computer based on the similarity calculation.

The simple graphical interface shown above has several limitations which were improved in a subsequent version of the program. First of all, the slider to represent agreement was an unnecessary complication that did not contain much useful information. Second, the visualization was complex and contained information that was not directly relevant to the task at hand: establishing logical connections between phrases.

In the second version, I opted to use two simple buttons: “agree” and “disagree”. What is needed most of all at this intermediate stage of data input is a fast and simple method for comparing sentences and inputting a logical connection between them. Moving a slider seemed

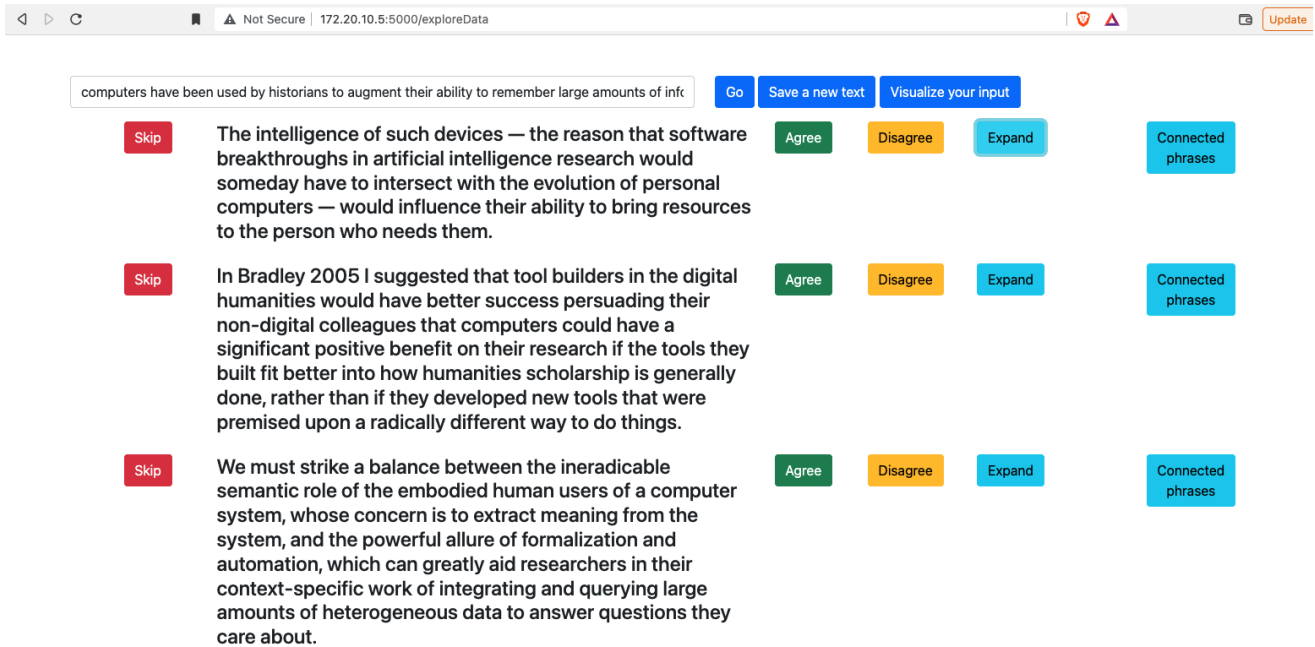
to introduce a layer of judgment that reflects uncertainty more than clarity, and slows the process down.<sup>208</sup>

The updated interface also contains a simple button that allows the user to see the context of a given phrase: the entire document could be displayed, or just the surrounding few sentences by traversing the graph a few steps in either direction along PRECEDING relation lines. In the current version of the program, I reference the document code in order to display the entire input that contains a given sentence.

The updated interface also contains a feature to see “connected phrases”, which returns a list of all the sentences that are connected to the query sentence with a similarity relation. Here is that interface, running in a web browser:

---

<sup>208</sup>In a research group there could be discord between two or more scholars. The existence of multiple relations of agreement or disagreement between two nodes could be resolved by averaging the values of the different judgments, since each relation is saved with the name of the user who saved it. The average could also be weighted according to the seniority of the scholar who made the link, or by some other metric. Multiple contradictory judgments could also simply be left as they are, a signpost of a different sort, pointing to a contended and potentially important point. The analysis of patterns of disagreement could be quite interesting in its own right. For instance, if we were using a scale with several values, the “votes” of multiple users could cluster in a normal distribution, but they could also cluster on the extremes of the scale, as in highly divisive topics that are hotly contested.

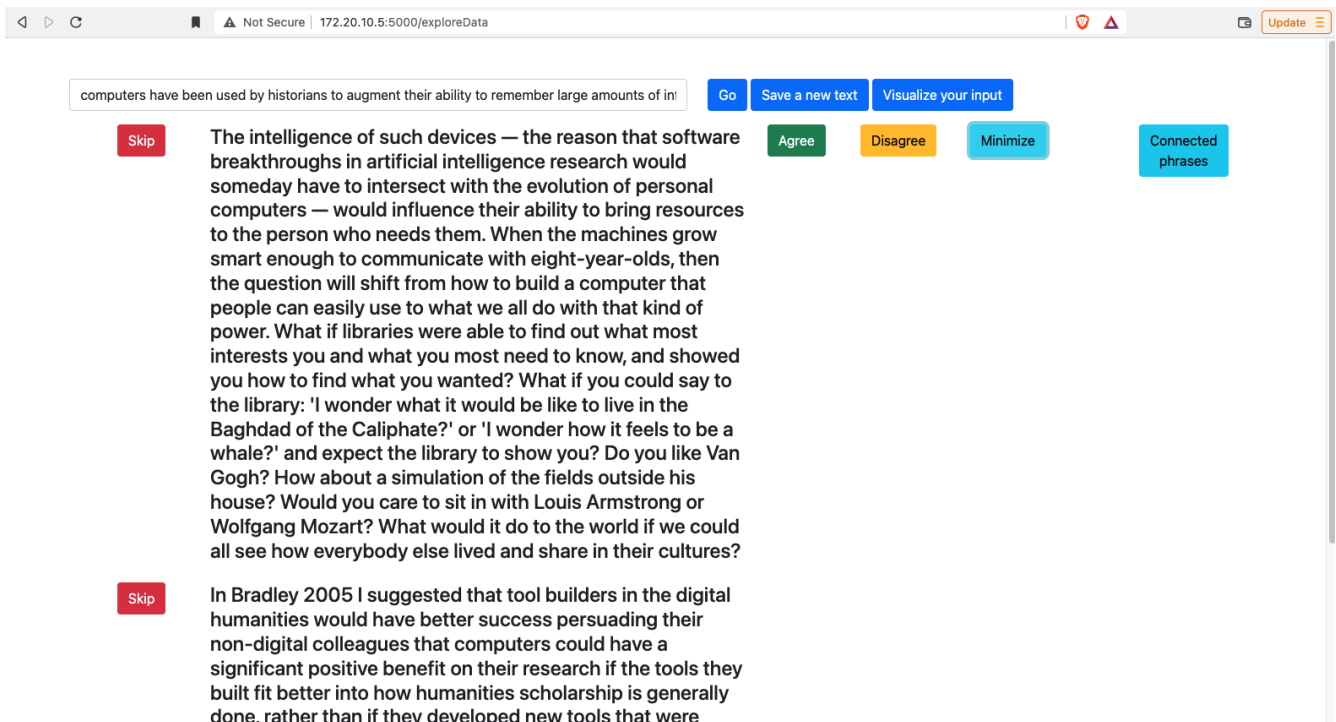


The input query, partially visible in the search line at the top, reads “computers have been used by historians to augment their ability to remember large amounts of information”. The system calculates the similarity between the query phrase and all the phrases in the database, and returns them in descending order, the most relevant first.<sup>209</sup>

By clicking on the “Expand” button, the context of the first phrase can be seen:

---

<sup>209</sup>This calculation requires a few seconds, depending on the processor of the computer. In order to provide “near instantaneous” response to the user, the current program first returns three results that are calculated by using a “keyword” search approach, which can be carried out much faster than a vector similarity search. The tradeoff is of course that keywords are less helpful in many cases. My rationale is that three responses are sufficient to keep the scholar occupied while the computer finishes calculating the similarity table, and they may be relevant enough to spark a logical connection. If not, he can simply skip through the three first results and see what the similarity table offers.



Bibliographic information regarding each phrase is not shown in this view, but it is saved in the system and can be retrieved. I decided not to include author name and title here in order to present a simple view to the user, and in order to emphasize the importance of creating logical connections between two phrases which represent two ideas, rather than potentially bias-ridden comparison between two author names or titles. The “expand” feature is useful at this stage as a way to be more sure of the meaning of the phrase.

This visualization should be considered a part of the input sequence more than data analysis: the visualizations described in the next chapter will be more useful to the researcher looking for insights from the existing data.<sup>210</sup>

---

<sup>210</sup>The two remaining buttons contain two supplementary functionalities. “Skip” removes the relative line from the visualization, and a new line is added to the bottom of the list. It is a fast way for the researcher to remove irrelevant data that the computer calculated to be similar, but which is not, and focus on other phrases. “Connected Phrases” acts by replacing the query line with the contents of the related result, and generates a new list of similar phrases. It essentially speeds up the search process by not requiring the user to copy or type the query if they see a phrase they wish to search for. It also speeds up the search process, because of the way the connected phrases are located and returned to the visualization. Since the database contains all the similarity relations between nodes that exceed a given threshold (currently a similarity value of 0.9), many nodes already have connections to other nodes, and can be queried directly through the Cypher query language instead of having to calculate similarities for all the nodes in the database. When the user pushes the “connected phrases” button, data is therefore returned very quickly. Sometimes, a node is not

## Using the Memex

One important value of my digital memex is the fact that the saved information has a structure of its own. Because the text has been expressed mathematically, it can also be studied in ways that are not usually considered in a historical context, but which can offer several advantages. I will proceed with a series of examples of increasing complexity. In the remainder of this chapter, first we will see what the memex can do by comparing a book to several reviews written about it. The small dimension of this example permits us to intuit and assess the quality of the results given by the system, since we can easily read the book and all the reviews, and compare that experience to the information returned by the computer.

Second, I will compare a series of books written on a single topic—in this case, Mesopotamian religion. Some of the questions we can address in this way are: Do various classic works on the subject cover the same territory, so to speak, or do some contain topics that are not treated by other authors? Are there authors who cover the topics with an even focus, and others who are more episodic in their treatment? Which of these features can be discovered automatically?

## Comparing a Book to its Reviews

For my first test, I chose A. Leo Oppenheim's classic *Ancient Mesopotamia: A portrait of a dead civilization*, and eight English language reviews of his volume. This was a simple test, a proof of concept of my system. Since the corpus is so small, it is easy to simply read the texts as well as analyze them with the system, and thus compare the results of the two processes.

To begin, I input the text of Oppenheim's *Introduction*.<sup>211</sup> Before proceeding to examine the whole book, I wanted to test a shorter text. Then, I input all the English language reviews I was able to find in digital format ( $n=8$ ) into the system. As detailed earlier in this chapter, the input

---

connected to any other nodes, and will return a blank page. This happens if there are no similarity relations above the threshold value. In that case, a search is still possible by pressing the "Go" button, which will perform a new calculation of similarity between the queried phrase and all phrases in the database, and return the list in descending order, even if all the relations are below the threshold value.

<sup>211</sup>Recall from Chapter 3: the input process involves splitting the text into sentences, saving each sentence as a 'node' in the Graph Object, and creating relations between nodes. First, sequential nodes from a single text are ordered by 'PRECEDES' relations. Then similar nodes are found by using cosine similarity, and saved with 'SIMILAR' relations when the similarity calculation exceeds a threshold value.

stage consisted in the creation of one ‘node’ in a graph database for each sentence in each text. Nodes are connected in sequential order and given a unique document code so that each document can be distinguished at a later time.


After inputting the sentences, they were connected by ‘similarity’ relations. These relations are the result of converting the text of each node to a vector representation, and then calculating the cosine of the angle between the vectors that represent two nodes. We can say that this number gives a measure of the semantic similarity between the two sentences: as an angle between two lines goes to zero, its cosine goes to one, and the lines overlap. We can interpret this to mean that if two vectors nearly coincide, and therefore have a very small angle between them, their high cosine score represents a measure of their similarity. In numbers, a low angle could give cosine of .98, which can be interpreted to mean high similarity between the two vectors.<sup>212</sup>

At this stage, I did not want to manage the extra complexity of links created laterally between reviews, so relations were only created between each review and Oppenheim’s text. Furthermore, links were only created if the semantic similarity score was over .90.

The first metric I examined was the overall level of connectivity between reviews and Oppenheim’s text. The following graph shows the frequency of connection as a function of their location in the introduction to *Portrait of a Dead Civilization*. The *Introduction* contains about 270 sentences, which are represented on the X-axis, and on the Y-axis we find the number of connections to the review articles. A line connects the connection number for each sentence:

---

<sup>212</sup>Of course, this description is simplified: the vectors we are discussing are not one- or two-dimensional—they have hundreds or thousands of dimensions, and they do not coincide as simply as two vectors on a cartesian grid might coincide. But for the purposes of an intuitive explanation, the above description is correct.



At first glance, the graph is promising. It is not a random pattern. It contains a few spikes that suggest that some specific sentences in the *Introduction* connect to several reviews. Which sentences are these, and what do they say?

Sentences 96, 102, 160, and 271 are highly connected, each with more than 30 links to the review articles. These sentences read as follows:

96: Because the library was not that of an individual scribe or even a school or family, but, rather, was brought together upon a royal fiat from all over Mesopotamia, we are entitled to assume that the topical range of Assurbanipal's collection is representative of the main body, if not the entire content, of the scribal tradition.

102: In spite of the fact that less than one-fourth of the body of traditional texts has been preserved, and only too often in rather poor condition, and in spite of the selection that is produced by the accidents of survival, of discovery, and - not to be underestimated - the accidents of publication, the unified picture that results from the observation of these well-distributed collections entitles us to speak of the literary tablets of Mesopotamia as belonging to a coherent and continuous stream.

160: Often this is done not out of methodological considerations or because of the range of the scholar's interest, but for reasons which seem rather to originate in a quest for a *raison d'être* for the entire field of Assyriology, not only in the eyes of other disciplines but also in the eyes of the scholars themselves.

271: All this is not meant to be a 'programme,' but neither should it be simply called wishful thinking-it is a way, well worth considering, out of the stagnation from which we suffer, a stagnation of which the most salient symptoms are the shrinkage of topics selected for research, the 'flight into specialization,' and the scarcity of students who once used to stray from theology into the perhaps greener pastures of a new and venturesome discipline.

According to our method, these are the phrases in the *Introduction* most highly connected to the reviews. These phrases could be interpreted as a summary of the points in Oppenheim's book that seem most important to his eight reviewers. For comparison, we can choose a completely unconnected node, sentence number 2, which reads: "It became obvious to me that no amount of painstaking atomizing, no endless inventories under the pretense of objectivity, and no application of any of the accepted, over-all patterns were capable of presenting the data in a way that would convey the whole as well as its integral constituents."

Sentence 2 was not picked up by the reviewers as an important idea; whereas for Oppenheim it expresses a core motivation behind the way he wrote his *Portrait*. It is also an important idea for me and for this thesis, and one of the reasons Oppenheim is relevant to my study. He reaches beyond the merely bibliographical, the 'painstaking atomizing', toward the 'whole and its integral constituents.'

This observation points us to notice that the system, prepared as it was in this first example, does not return a list of the most important phrases according to the *author*, but according to the *reviewers* taken in aggregate. We could also query the data in a different way: we could look for particularly important affirmations by reviewers or by Oppenheim, in our own opinion, and examine the sentences that are calculated as having a high degree of similarity to them. In this way, we could reverse the process, and ask the system to reveal subterranean connections to ideas that we consider important.

To take another example, note Oppenheim's explanation in the *Preface* that he does not intend to go into detail on many issues, because he aims to write a book that could "communicate with non-Assyriologists" (sentence 33). According to my program, the sentence is most highly related to an affirmation by Joan Oates, who writes in her review approvingly of this decision: "At the same time one cannot help but be aware of the enormous scope and depth



of the subject and the fact that any attempt to offer such explanations would have grossly increased the size of what is at it stands an eminently lucid and readable volume.”<sup>213</sup> In this case, a significant affirmation by Oppenheim is echoed and completed by a reviewer.

A second example of a relevant connection can be found in Kramer’s review. He writes that “among people in general there exists a diversity of types, temperaments and dispositions among scholars, and each makes a useful contribution to human knowledge”.<sup>214</sup> In the context of the review, this phrase seems to mean that Kramer would like to downplay the diatribes and divisions between methodologies.

His affirmation connects with high similarity to Oppenheim’s sentence 295: “Of course, the classical scholar may also have to face new and surprising data, but only exceptionally can such data be compared in scope and relevance to what the Assyriologist has every right to expect.” We must be careful not to read too much into these connections, but they do suggest a sort of internal dialog between the ideas, in which the issue of different approaches to the material, and different levels of difficulty among the specializations in history, take two forms: conciliatory in Kramer, and almost bitterly defensive in Oppenheim. According to the latter, the Assyriologist must face the great difficulties of fragmentary data and almost daily discoveries that upend previous scholarship. He seems to mean that his style is not only the merely obvious result of variation among humans, and the idiosyncrasy of his own personality, but is also inextricably linked to the nature of the particular study he is engaged in.

These first results suggest that we might be on the path of something, but they do nothing that a careful reader would be unable to do by simply reading the texts directly and paying attention to the salient affirmations. It is true, however, that even at this early stage we see some usefulness in using the similarity calculation as a heuristic to guide our attention, and attend first of all to those affirmations that connect two documents together. Such attention does produce significant results. Once again, the purpose of algorithmic analysis is not to replace human judgment, nor can it necessarily help one to see more deeply than one can by

---

213In order to create the graph of citation frequency, I queried the database to retrieve the sentences that are linked to other sentences by similarity relations, and plotted the sentence ID against the count of relations:

```
MATCH (n:Sentence{title:"Portrait – Preface"})-[r](m) RETURN id(n), COUNT(r)
```

214Samuel Noah Kramer, “Review of Ancient Mesopotamia: Portrait of a Dead Civilization”, *Archaeology*, 19:2 (1966), 140.

directly reading a text. In many cases, the main usefulness of the algorithm is to help a researcher quickly to survey a broad field, and focus on central issues from the outset. Other times, it helps to uncover links between ideas and texts that might be lost in the sheer quantity that a scholar must peruse.

After successfully visualizing the coverage of the *Introduction*, I proceeded to compare the entirety of Oppenheim's *Portrait* to the eight English language reviews of his work I had previously chosen. As before, I hypothesized that my method would map the coverage of Oppenheim's work by his reviewers: sections of the book that were amply discussed in the reviews should appear as highly-connected, and sections that were not discussed by the reviewers should exhibit lower connectivity.<sup>215</sup>

First, I created a graph comprising the full text of Oppenheim's work. *Portrait* is available in high-quality PDF from the University of Chicago Oriental Institute.<sup>216</sup> I extracted the text of the preface and six chapters, excluding appendices and indices. A pre-processing step involved the removal of page headings and numberings, and footnote numbers, in order to present my program with a text as cleanly formatted as possible. This text, using the method I have described above, was input line by line into a Neo4j graph database.<sup>217</sup>

As before, my first-order interest regarded only the coverage of Oppenheim's topics by his reviewers and not the connections that could be made between those reviewers. Since Oppenheim's text contains 4,120 sentences and the calculation of connections involves comparing every sentence in the review articles to every phrase in the book, calculation required several hours.

I implemented a cut-off threshold of 90% similarity so that greatly dissimilar phrases would be excluded. Even so, there are about 33,000 connections in the resulting graph. When I examined some of the lower-valued similarity relations, it seemed that a threshold of .90 was

---


<sup>215</sup>In this context, 'connectivity' means the number of SIMILAR connections a given node has to all other nodes.

<sup>216</sup>Available at <https://isac.uchicago.edu/research/publications/misc/ancient-mesopotamia-portrait-dead-civilization>

<sup>217</sup>In my first experiments, these preprocessing steps required a great deal of time, so much that some of my colleagues wondered if the effort was worth the result: would it not have been faster to simply read the book? But these are steps that can be dramatically sped up with algorithmic methods and computational optimizations. Therefore, I consider their present clumsiness to be an indication of the need to collaborate with an excellent computer scientist, not an insurmountable difficulty.

not precise enough for my purposes—the phrases connected at that threshold are too dissimilar, in actual linguistic fact, to be significant. A higher threshold, like .95, seemed preferable because the connections are much more significant and also much less numerous, which makes the interpretive process faster and more fruitful. Fortunately, it is a simple matter to filter the results for only those SIMILAR relations with a value above .95.

There are 1,225 SIMILAR relations with a value above .96. I extracted the index numbers for the phrases in Oppenheim’s text that correspond to those 1,225 connections, and wrote a script to count how many connections were on each phrase. The output from this program was then graphed using a spreadsheet, as shown below.



This image shows the relative frequency of connections as a function of phrase number in Oppenheim’s book. The eight reviewers are not distinguished; for the moment, we are only interested in the aggregate, how the group of reviewers covered the topic-space of the text. For reference, here is a list of the book’s chapter headings and sub-headings, presented as sentence numbers:

sentences 1-331	Prefatory note, Introduction: <b>Assyriology—Why and How?</b>
sentences 332-825	Chapter 1: <b>The Making of Mesopotamia</b> (The background, the setting, the actors, the world around)


sentences 826-1675	Chapter 2: <b>Go to, let us build us a city and a tower!</b> (The social texture, economic facts, “the great organizations”, the city, urbanism)
sentences 1676-2014	Chapter 3: <b>Regnum a gente in gentem transfertur</b> (Historical sources or literature?, an essay on Babylonian history, an essay on Assyrian history)
sentences 2015-2739	Chapter 4: <b>Nah ist—und schwer zu fassen der Gott</b> (Why a “Mesopotamian religion” should not be written, the care and feeding of the gods, Mesopotamian “psychology”, the arts of the diviner)
sentences 2740-3529	Chapter 5: <b>Laterculis coctilibus</b> (The meaning of writing, the scribes, the creative effort, patterns in non-literary texts)
sentences 3530-4120	Chapter 6: <b>There are many strange wonders, but nothing more wonderful than man</b> (Medicine and physicians, mathematics and astronomy, craftsmen and artists)

A first examination of this graph shows that between phrase 1000 and 2000, there is a dearth of connections. There is, on the other hand, a high density of connections between phrases 0 and 400, between 2000 and 2,600, and between 3,500 and 3,800. These numbers correspond with the preface, the fourth chapter, and the sixth chapter, which are indeed the topics most covered by the reviewers – the overview provided by the author in his preface, the surprising and controversial chapter on Mesopotamian religion, and the technical chapter about medicine, mathematics, and craftsmen. On the other hand, the chapters which are left less covered are those about writing (chapter 5) and the city (chapter 2).

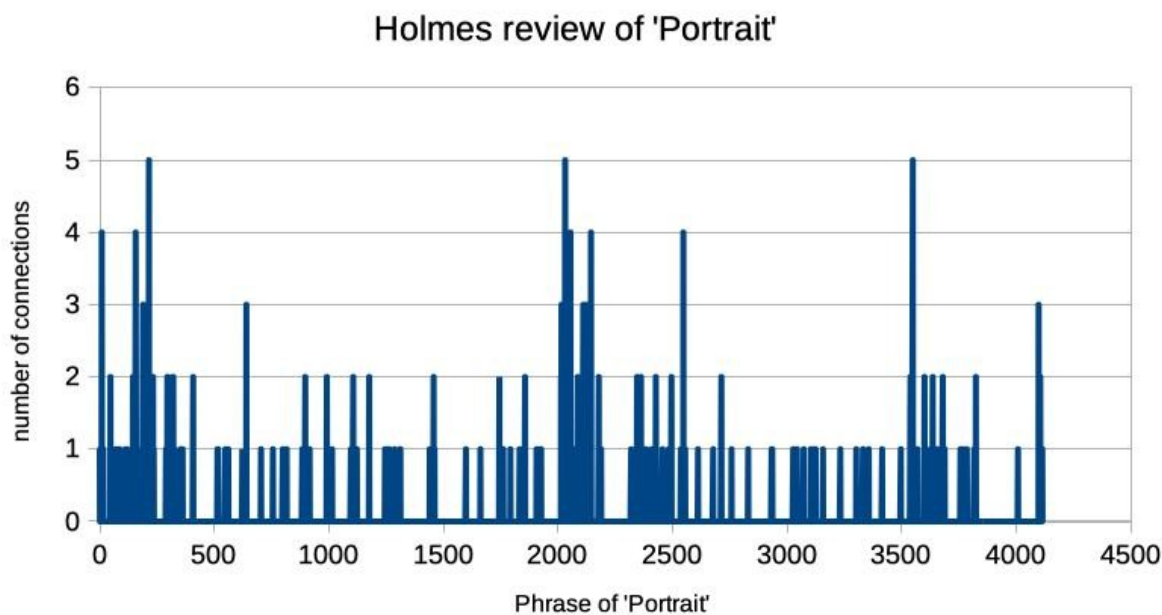
Next, I extracted a single reviewer and compared the signature of connections to Oppenheim’s text with the content of the review as I could understand it by reading it.

One reviewer, Forbes, dedicated about half of his review to the technical and artisanal details contained in Oppenheim’s book. If we except the introduction, which by its general

nature is well-connected to most of the reviews, we see that most of the connections Forbes' review makes with Oppenheim's book are concentrated in the fifth and sixth chapter. By looking at the signature of Forbes' connections, we would expect his interest to lie in writing and technology: and in fact, his review was published in the *Journal of the Society for the History of Technology*.



Elsie Holmes dedicated about half of her review to the topic of religion and, as can be seen in the graph of her connections to Oppenheim's text, the fourth chapter is most heavily represented.



These results are encouraging: even if the reviewers each have idiosyncratic styles, and although current natural language processing methods are far from perfect and do not correspond to an ‘understanding’ by the computer, the program is able, on the aggregate, to return to the user information that corresponds to the specificities of each of the writers. The automated extraction of topics can be used to guide the researcher toward relevant resources.

In this example, we only dealt with eight reviews and one book. This quantity of text could much more easily be read and understood by a human scholar than by writing programs and interpreting histograms. But what about a larger quantity of text?

## Expanding the experiment

The second experiment I conducted using this method regards the study of Mesopotamian religion from Leo Oppenheim’s time to the present. This subject, presented caustically by the author as “why a Mesopotamian religion should not be written”, was among the most controversial in Oppenheim’s classic. I wondered: how has Oppenheim’s *mot* fared, that a Mesopotamian religion should not be written? Was his advice valid for his time, but not for ours? Or has the proliferation of new data and interpretations only multiplied the details and

pushed synthesis even further away? To some degree, these questions ought to be amenable to study using my method.

Ten years after Oppenheim published his *Portrait*, his colleague at the University of Chicago, Thorkild Jacobsen, published a monograph entirely dedicated to Mesopotamian religion, entitled *Treasures of Darkness*.<sup>218</sup> Other studies of note in the following decades include a few studies by Bottéro<sup>219</sup>, Smith's analysis of monotheism and polytheism<sup>220</sup>, and Buccellati's recent monograph *When on High the Heavens*.<sup>221</sup>

My first test was to compare Oppenheim and Jacobsen. Does Jacobsen cover more ground than Oppenheim? In a sense, he must, because he wrote a monograph on the topic, whereas Oppenheim only wrote a chapter. However, it is possible that a long text not cover the entire 'thought-space' represented by another text. Which is the case?

In order to compare Oppenheim to Jacobsen, the first step was to input all of Jacobsen's book into the database and to create connections between it and Oppenheim's book.<sup>222</sup> Once I had cleaned the database, it contained approximately 7,000 nodes and 62,000 connections between them.<sup>223</sup> The sensitivity threshold was set at .955.

---

218Thorkild Jacobsen, *Treasures of Darkness. A history of Mesopotamian Religion*, Yale University Press (1976).

219Especially Jean Bottéro, *Writing, Reasoning, and the Gods*, University of Chicago Press (1992).

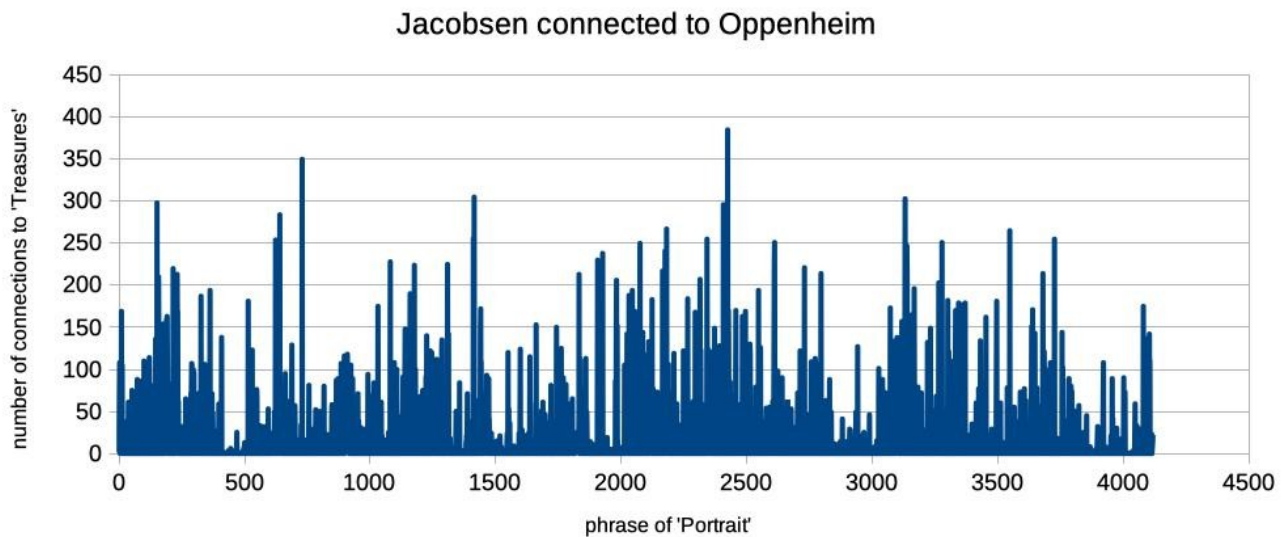
220Mark Smith, *The Origins of Biblical Monotheism: Israel's Polytheistic Background and the Ugaritic Texts*, Oxford University Press (2003).

221Giorgio Buccellati, *When on High the Heavens... Mesopotamian Religion and Spirituality with Reference to the Biblical World*, Routledge (2024).

222This process highlighted the problem of computational complexity with new urgency. In order to compare two books which each run about 3,500 sentences, the computer processing time was 32 hours. Clearly, with the program at its current level of efficiency, running on a laptop, it would not be possible to examine large corpora. But during the development of my program, I have learned many techniques to improve efficiency, and the program has steadily improved. In the last examples in this chapter, a far greater amount of text was successfully input and calculated, thanks to some of these improvements.

223Once the two books were input and connected, I noticed that some nodes (sentences) were connected to other nodes with 100% similarity. This made me suspicious, so I examined them more closely and realized that empty sentences were sometimes created at input time, which caused the program to calculate identical relations between sentences that were completely different. Luckily, this was easy to fix. I used a query to identify the source phrases of these relations, deleted them, and "stitched together" the surrounding phrases to recover the form of a book as one long series of sentences. This is the reason that the subsequent graphs contain some spurious components far from the main body of the text: if lower-numbered nodes are deleted at some point, importing text at a later time leads to non-sequential numbering of the nodes. Since the nodes involved are few, I simply removed them from the visualization instead of repeating the entire import and calculation process.

I then exported the list of connections between nodes based on the ID number of each node, which I input into a script I wrote called ‘count frequencies.py’. This script counts the number of connections on each of the phrases in one of the books. To begin, I calculated the connections of Jacobsen’s *Treasures* to Oppenheim’s *Portrait*. The resulting histogram is thus comparable to those we saw above. On the x-axis are the ID numbers for all the sentences in Oppenheim’s book, and on the y-axis the number of phrases in Jacobsen’s book that connect to each of them. The graph is a lot more complex than those we saw above, as we might expect, because Jacobsen’s book runs over 3,000 sentences, as compared with the few dozen in Oppenheim’s reviews. The pattern also shows what we might expect: Jacobsen’s book is densely connected to the chapter on religion in Oppenheim’s *Portrait* (sentences 2015-2739). There are also many other significant connections in other chapters.

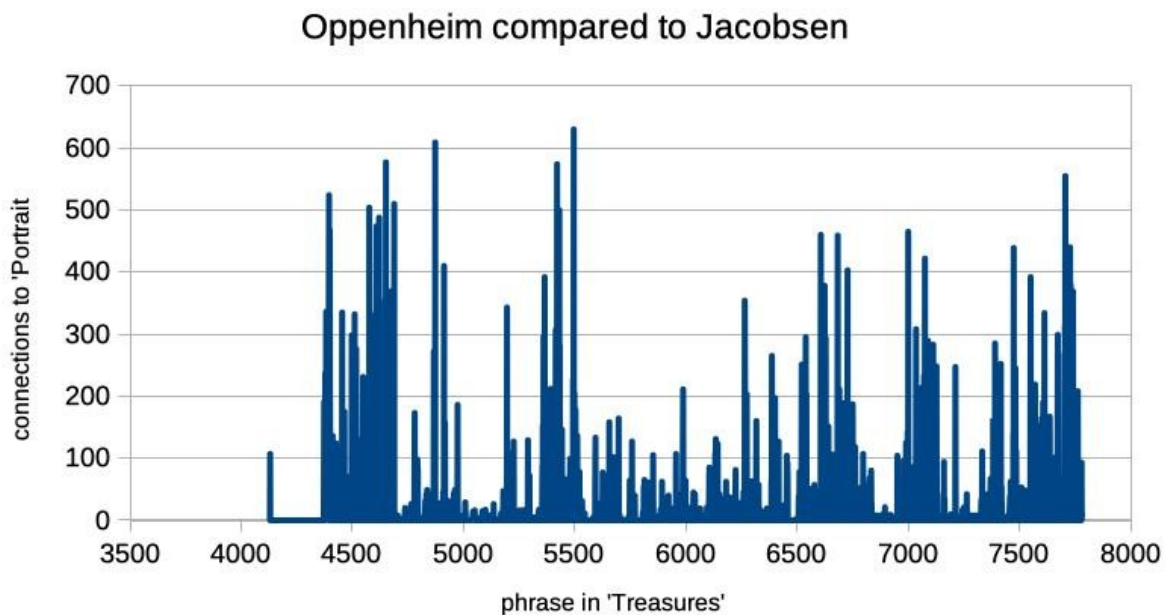


We can also view the data the other way around and put Jacobsen’s *Treasures* on the x-axis and put connections to Oppenheim’s *Portrait* on the y-axis, as seen in the next figure. This image is more useful because it shows that the topics Jacobsen deals with are similar to Oppenheim’s topics only in some of his chapters. In order to understand the reason for this difference, we recall that while the x-axis represents the sequence of sentences in one of the books, the y-axis merely represents the *quantity* of connections to that sentence, not their



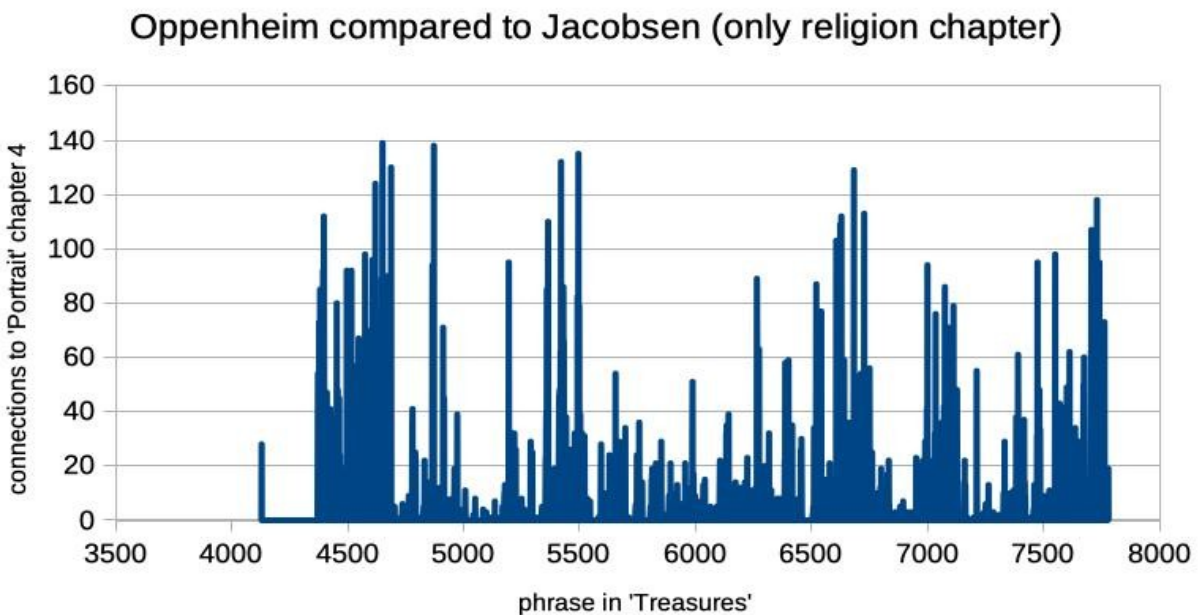
source. This means that while Oppenheim's chapter on technology, for instance, may receive many connections from Jacobsen, without further analysis we do not know where those connections come from. They could be from scattered phrases among the pages of Jacobsen's book, or they could represent the same few phrases repeatedly connecting to Oppenheim's text.

By inverting the axes and showing Jacobsen's book on the x-axis, we see that Jacobsen's second and fourth chapters are barely connected to Oppenheim's book. Upon examination, this is understandable. Jacobsen's second chapter, entitled 'The Gods As Providers: Dying Gods of Fertility' deals extensively with the Dumuzi cult, which was a theme of particular interest for Jacobsen, while Oppenheim only mentions it in passing. Similarly, Jacobsen's fourth chapter is an overview of ten Mesopotamian gods who he interprets as the embodiment of abstract principles such as authority, force, cunning, and so on. These details and the copious primary texts Jacobsen cites are also absent from Oppenheim's work. (Note that the sentence numbers in the x-axis do not start at zero because the index numbers in the database from zero to 4120 are occupied with the sentences in Oppenheim's work. In order to recover the proper sentence number in Jacobsen, we need to subtract 4120 from the number seen in the chart. This problem was resolved in a later version of the program, which we will discuss soon.)



This representation is still too coarsely grained. It shows the connections between Jacobsen’s entire book and Oppenheim’s entire book. This, however, is not a fair comparison, since Oppenheim’s book deals with all of Mesopotamian history, and Jacobsen’s book only deals with Mesopotamian religion. If we consider only chapter four of Oppenheim’s book, the chapter on religion, what difference would we see?

The next image shows Jacobsen’s book’s sentences on the x-axis and number of connections to Oppenheim’s *Portrait* chapter 4 on the y-axis. Again we see the dearth of connections in the two chapters mentioned above, so those “holes” are in fact robust indications of topics Jacobsen treats but Oppenheim does not.




I had initially hoped to load a larger series of books into the system in this second experiment, in order to map the development of the field over the last half-century. However, after processing just two of them I realized that this would not be possible with the current technical setup because it required too much processing time on a laptop. I settled for just one more full book-length comparison at this stage.

I chose a third scholar who studied at the University of Chicago and knew both Jacobsen and Oppenheim: Giorgio Buccellati. He recently published a monograph on Mesopotamian

religion, which I helped translate into English.<sup>224</sup> It was thus easy for me to input the complete text into my system, since the time-consuming pre-processing steps of extracting text from pdfs was not necessary. Following the same procedure as above, and noting again the x-axis shift (Buccellati's first sentence corresponds with the last of Jacobsen, or number 7750), we produce the following graphs:


Like we saw with Oppenheim, Buccellati does not extensively treat the Dumuzi cult in chapter 2, nor the ten gods in chapter 4. These topics remain the 'signature' of Jacobsen's study.




If instead we run the comparison the other way around, with Buccellati on the x-axis and Jacobsen's connections counted on the y-axis, we find a curious and significant structure:

---

<sup>224</sup>Giorgio Buccellati, *When on High the Heavens... Mesopotamian Religion and Spirituality with Reference to the Biblical World*, Routledge (2024).



Buccellati's book is extremely smoothly connected. It does not display any holes! If we carry out the same comparison to Oppenheim, we see the following signature:



Again, no holes. This result corresponds to what I would expect as a reader of the three books. Oppenheim's work is wide ranging, fairly complete, quite idiosyncratic in style, and shows a small number of 'holes' that correspond to small sections of chapters. Jacobsen's work

contains two long sections that deal with issues and ideas which are germane to his topic, but which are also very specific, like case studies. Buccellati instead presents a closely articulated book over 26 chapters, much more detailed and encyclopedic than either of the other two scholars. Furthermore, each chapter is of similar length and covers its topics in similar depth. This leads to a ‘signature’ that shows no holes.

At this stage, thanks to a conversation with my tutor Prof. Mauro Giorgieri, it became clear that mere histograms can show something of the overall structure of a book, but what a scholar would really like to do at this point is to explore the histogram in detail, and see exactly *which* phrases spike upward above the rest, and exactly which parts of a book remain unconnected to its peers. In order to address this desire, it was necessary to write a new part of my program.

The result, a “histogram explorer”, as I call it, is a way to interact directly with the histograms produced by comparing a baseline text with one or many other works chosen by the user. It offers a web interface with the ability to zoom in to close detail, and several tools that express which sentence in the baseline text corresponds to a given sentence number, and the contents of the sentences that connect to it. There is also the possibility to choose a threshold similarity value and exclude connections below that point. The details of its operation will become clear in the next chapter.

## Chapter 4 – Models of Historiography

Up to this point, I have set the stage: I have discussed the creation of models, declared my preference for “inner-referential” models and my intention to follow Oppenheim’s call to the “battle for synthesis”. I have also described several techniques used in digital humanism and the formal structure and interface of a computer program that aims to help us carry out our goals within the domain of historiography. To summarize the most important point: part of the information contained in text and across scholarly communities can be represented as a network. This network representation permits algorithmic access to the underlying information, and can improve interpretation by revealing layers of information that are hidden to the unaided eye.

In order to help my reader distinguish the network representation, which is a complex mathematical object with a specific ‘shape’ or topology, invariant under translation and rotation, from the tools that are used in building and accessing it, I will refer from here forward to it as the ‘Graph Object’: a shorthand, to mean the web of interconnected sentences produced through the methods outlined in Chapter 3.

The underlying structure of the Graph Object remains stable until new text is added to it. Future developments in language models may well improve the quality of the Graph Objects we can produce, by improving the details of sentence parsing and vectorization. Such developments are here considered to be external events that do not change my fundamental thesis. So too, improvements in interface design, the use of programming languages that are more computationally efficient than the Python code I have written, and new discoveries of ways to query and characterize the important nodes within a graph will all improve the results of the technique I am describing, without changing its essential characteristic: the representation of text information as a Graph Object.

## Writing the Mes-Rel Website

In order to clearly perceive the *cui bono* to which this dissertation tends, I dedicate the beginning of this chapter to further specify the context in which I became aware of the problem and began to develop the solutions I propose here. Although I have already done so from several theoretical points of view in the previous chapters, at this point it will be helpful to conduct the ‘battle for synthesis’ anew, enriched with the theoretical underpinnings we have seen so far. We loop back around to the place where we started, and discover that our progress is not circular, but rather like a spiral that returns to familiar themes in the plane, yet has progressed forward in a third axis.

This context is not only the chronologically first place I faced the issues treated in this thesis, it is also an example of the creation of a model of historiography within a specific domain. As such, it can help us to perceive the usefulness of the Graph Object that I propose. Once I have clarified the context, the rest of this chapter and the next will explore successively more sophisticated aspects of the information the Graph Object can help surface.

This thesis had its origin in the project of developing a multi-planar form of writing on the Mesopotamian Religion (Mes-Rel) website of which I am the associate editor. It may seem trivial to point out that this thesis is a linear discourse, and while reading it, the reader experiences time as a linear succession of moments and ideas. One starts at the beginning, moves through a development, and arrives at a conclusion. Even when an expert reader starts with the conclusion and jumps around within the text in order to decide whether the content is worth time and close attention, the experience of reading is still linear.

However, the information the dissertation draws on is rather web-like in structure. Every word and phrase is partially the echo of other, older words and phrases. Nearly every idea has been treated elsewhere and otherwise, many times already. One hopes there are a few newer ideas in the text which warrant the writing of the thesis—but the majority of its contents and its form are not new. And each idea in the text is related to other ideas, in a web of connections which, although not actually infinite, is as close to an actual infinity as anything else that humans can directly experience. In the case of the Mesopotamian Religion website, considerations about religion are related to issues about earlier hominid burial practices, to

issues of psychology and neurophysiology, and from there to questions of biology and chemistry, as well as to linguistics and the practice of metaphor and symbol, and also to issues about the computational limits of the human brain... the web has no end.

So, when writing any text, the inescapable question arises: where should boundaries be drawn around the topic, in order to define its outlines and limit it to its proper dimensions? An article should achieve the status of a “whole” in its own right, as well as that of a “part” of a larger work and of a field of inquiry. These issues are familiar to any writer. In the case of the Mes-Rel website, the question was not only a background aspect of the technique of writing, but was also one of the specific items of our inquiry. The website aims to be not merely a repository of text that could just as well have been printed on paper, but to be a specifically *digital* artifact that leverages the opportunities that digital writing affords. I begin with a closer look at the topic of boundaries.

## Wholes and Parts

Where are the boundaries of the “whole” which one aims to write? One answer to this question is the choice of a frame which defines the inside and the outside of a discourse. This is generally necessary in writing a work of any length, since the writer is conscious of more information than can be contained in the text. The writing process involves ordering and prioritizing, summarizing prior work, and drawing the reader’s attention to the elements that the writer considers most important in the present context. It also involves the complementary “negative” skills: avoiding sterile lines of argument and passing over some topics in silence. Not every related issue can be treated in the space of an article or a book, so the author must be discerning and concise. The choice of frame might be guided by tradition or by the specific formation of the author, which in some cases tends toward specialization, and in other cases tends toward encyclopedic breadth.

Another way to conceive of the relation between wholes and parts is the concept of cellularity. “The moment we consider fragments as cells, we emphasize their belonging to the whole”, writes G. Buccellati in his contribution to a volume on *Humanism and Digitization*.<sup>225</sup>

---

<sup>225</sup>*Digital Discourse: A view from the field. Festschrift James Walker, 2023. Forthcoming.*



Cells conjure the image of a living organism, which is made up of parts but which is also more than the sum of the parts. To apply this metaphor to writing, we can imagine that articles and books are “cells” within the living organism of a discipline, and have explicit and implicit connections to the rest of the components of that discipline. The larger discipline can be understood as a framing of a subset of knowledge, in which individual contributions are “cells” of that subset—necessary and insufficient components of a greater whole.

This metaphor faces some difficulties, however, if we probe more deeply. Is it true that academic disciplines are like organisms, perhaps like different species of animals? In an important sense, no: the disciplines are not self-standing. Important discoveries in mathematics modify physics and chemistry; important discoveries in neuroscience and psychology modify history, philosophy and sociology, and vice versa.<sup>226</sup> Perhaps the “disciplines as organisms” metaphor could be improved if we consider the various discipline-organisms as parts of a larger ecosystem, in which each organism has its own identity, and is connected to all the rest.

In my view, such an extended metaphor could be useful especially if we consider the ecosystem to be potentially infinite. I do not think that our current knowledge covers more than a small fragment of the total space of knowledge and reality. Perhaps the way the academic community currently understands wholes and parts, cells and organisms, is still akin to the situation at Aristotle’s time, when some remarkably robust distinctions had already been discovered, while many organisms and even the very principle of cellularity had not yet been directly observed. The history of human thought contains many shifts in perception, re-definitions of boundaries, and the attendant invention of new disciplines.

To give one recent example, recall that only a few decades ago, the relations between trees in a forest were understood mostly in terms of competition for scarce resources. Today, thanks to Suzanne Simard’s discovery of the importance of fungal networks at the end of the 1990s, our conception of forests has changed.<sup>227</sup> Where the divisions are drawn makes a difference:

---

<sup>226</sup>It is also true that many disciplines interact directly only rarely, if at all. See Katy Börner, *Atlas of Science: Visualizing What We Know* (Cambridge, Mass: MIT press, 2010), 13.

<sup>227</sup>Simard, S.W., Jones, M.D., Durall, D.M., Perry, D.A., Myrold, D.D., Molina R. (1997). “Reciprocal transfer of carbon isotopes between ectomycorrhizal *Betula papyrifera* and *Pseudotsuga menziesii*.” *New Phytologist* 137: 529-542.

considering the individual trees as “wholes” that compete or as “cells” that contribute to a larger organism can radically change the conclusions we reach downstream.<sup>228</sup>

I believe the same is true within intellectual work. There is certainly great utility to be found through specialization, but the narrowing of the researcher’s gaze can just as well lead to false conclusions if the disciplines are in fact connected by subterranean networks, like the trees in the forest. To return to our specific question and ground the metaphor: when writing our website, we aimed to produce a whole which has boundaries and which is also in relation to other wholes.

The frame of reference can also move in the opposite direction, toward the subdivision of “cells” into their smaller constituent components. The resulting atomization of material and metaphorical reality has led to important discoveries in many fields, and is clearly a viable and useful procedure.<sup>229</sup> Physics is the classic example of the success of this way of thinking.<sup>230</sup>

In his *Critique of Archaeological Reason*, G. Buccellati has extensively argued that archeology is characterized by the atomistic, “digital” nature of its primary source of information: the primitive data with which one has to deal is essentially fragmentary. In his view, it is the work of the scholar to discover links between the data and hypothesize the ways the fragments were once connected in their original whole environment.<sup>231</sup> Buccellati points out that the archeologist’s work would be incomplete if it were merely the collecting of fragments, and not also the offering of a proposal of a coherent view of how the fragments connect into wholes. The scholarly task is to achieve control of the fragments, and to interpret them within a

---

228The frame of reference also changes the actions of foresters who seek to improve the overall health of the forest. To wit: reasoning within a competition framework, foresters initially cut down the birch trees in a mixed Douglas fir and birch forest, believing that this action would improve the health and productivity of the fir trees they meant to harvest for lumber. However, after the birches were cut, the firs suffered as well. Then Simard discovered that the birches and the firs exchange important nutrients in a collaborative network connected by fungi: the “whole organism” was more complex than the individual trees, and the health of the trees depended upon more factors than were immediately apparent.

229A useful analog can be found in the study of phonology in linguistics. The breaking of words into their phonetic elements, famously by the Prague School, has helped linguists reach some important conclusions.

230It is worth mentioning that some eminent physicists, such as David Bohm, have argued that there is also an “implicate order” among the fragments which points toward the existence of unity as a prior reality, and to connections between the fragments that persist even under conditions of non-locality, such as in quantum entanglement.

231See in particular Buccellati, *Critique of Archaeological Reason*, Cambridge University Press (2017) chapter 2, and the accompanying website [critique-of-ar.net](http://critique-of-ar.net).

coherent narrative that remains transparent about what is known, and how well it is known. These characteristics, in Buccellati's view, make archeology a natively digital discipline.<sup>232</sup>

To move from the whole to the fragment and then back to a complete narrative is reminiscent of an analogous problem in linguistics. Words can be subdivided into phonemes, and discourse can be understood as a linear sequence of phonemes.

But linguists also emphasize that discourse is not only linear: in a discourse, words are put together in such a way that their references often leap over nearer words. In order to understand a phrase, linear word order is not enough, and the reader or the automatic system, as it may be, must also be sensitive to longer-distance dependencies and contexts.<sup>233</sup> Many poetic and rhetorical devices take the form of large repetitive structures, such as rhyme schemes, alliteration, or the chiasmic ABCBA form used in some biblical texts: these structures cannot be understood only in terms of a linear sequence.<sup>234</sup>

Even a cursory examination of this theme would take us far afield. For the purposes of this chapter, I limit myself to observing that while it is uncontroversial to note that a fragment of discourse can be understood fully only by taking its context into account, frequently what we mean by "context" cannot be strictly defined as locality within the linear time-sequence of signs or sounds making up the discourse. It may be difficult indeed to define the context for correct understanding of a phrase or concept. This fact implies that the boundaries of wholes and parts must be treated as if well-defined when one is writing—but that those same boundaries must also be considered as essentially open to modification in order to ever more closely approximate truth in what one writes.

We may extend this observation as a principle of *epoché* that accompanies larger-scale contexts as well. While it is necessary to postulate some limits to a field, and further subdivide that field into disciplines and then into the dimensions of books or articles, it is also necessary that we continue to consider that our decisions about boundaries are perfectible. Our current knowledge is amenable to correction on every level, including our presuppositions about where

---

232Ibid., p. 232.

233Much of the recent success of automatic text generation and translation systems like ChatGPT depends on their "transformer" architecture, which is specifically capable of detecting long distance dependencies.

234For instance, see Mark 2:27; Joel 3:17-21; Isaiah 1:21-26.

the boundaries of “wholes” and “parts” may lie. Like in Simard’s forest, the distinctions between ideas (or organisms) may seem clear at first, but occasionally they have to be radically rethought on the basis of new information.

This fact has a few important consequences for a project such as the Mesopotamian Religion website. On the one hand, we aim to do more than just produce a catalog of fragments. The fragments can be properly understood only within their context, and we attempt a synthesis that presents the known data together with an interpretation. At the same time, our interpretation should remain structurally open to the possibility that further thought and discovery will refine the boundaries between “wholes” and “parts”. Consequently, the Mes-Rel website, which aims to be truly “digital” in the sense Buccellati describes in the *Critique*, must simultaneously present a narrative or narratives, which implies that it posits certain robust boundaries around those narratives, and also contain mechanisms by which those very boundaries can be observed, critiqued, extended or restricted.

## **Digital Writing**

In this section, I describe several aspects of the website as they developed in my thought and practice, together with what I perceive to be the strengths and weaknesses they exhibit in their current form. Each of these aspects is an attempt to account for the relation between wholes and parts.

## The core narrative(s)

In its current form, the Mes-Rel website contains one “core narrative”, which is represented by two text-objects in particular. First, the core narrative is the printed book entitled *When on High the Heavens...* This work is a concise, detailed volume that articulates the claim that notwithstanding many similarities between the two, there is an irreducible structural difference between what the author terms Mesopotamian and Biblical “spirituality”. In each of the book’s 26 chapters, specific aspects of Mesopotamian spirituality are examined and placed in relief through a sustained comparison with Biblical analogies. At its basic level, the companion website acts as a technical apparatus of footnotes, bibliography, and indices to the printed volume.

In order to aid the reader of the website, there is a brief encyclopedic “history of the discipline” section to introduce the reader to some of the central issues that inform our interpretations.<sup>235</sup> There is also a “core narrative” section that presents approximately one-page summaries of each chapter in the book.<sup>236</sup> I wrote these summaries in such a way as to contain the key arguments deployed by Buccellati, but to also be very short and suitable for easy online reading. They aim to facilitate readers who may not have the physical book on hand, and their presence on the website also provides a detailed “table of contents” beyond what can be understood from the chapter titles and subheadings. They also provide a digital structure which supports hyperlinks to the more granular elements contained on the website, from thematic essays to footnotes and bibliography entries.<sup>237</sup>

---

235 See <https://4banks.net/Mes-Rel/history.htm>

236 See <https://4banks.net/Mes-rel/core.htm>

237 At one time, the “core narrative” pages also offered a platform on which readers could add annotations and gradually enrich the website beyond what was contained at the time of publication, through a third-party plugin called *hypothes.is*. Since the web pages were written in html, in order to update their contents it was necessary to have direct access to the underlying codebase and knowledge of the html programming language in order to modify or add hyperlinks. The *Hypothes.is* plugin was a solution that permitted the underlying page to remain unvaried, while permitting a potentially large number of competent scholars to contribute directly to building the web of interconnections between texts and data points. At the time of this writing, the plugin is no longer operative on the website, and the method of updating pages has been changed to a simplified programming language known as markdown. See <https://web.hypothes.is/>. This solution allows registered members to annotate the pages of the website directly, including the creation of hyperlinks to other pages. Readers of the website can toggle the visibility of annotations on and off. While all viewers can see annotations, it is possible to limit who is permitted to annotate the pages.

Although there is only one “core narrative” at present, it is our intention that future editions of the website will contain multiple narratives that refer to the same underlying bibliography and source data. This is a significant concept, and one that is only possible within the digital paradigm: the website will contain a plurality of narrative lines or planes that intersect and interact with a much wider web of information, explicitly and transparently.<sup>238</sup> To realize such a project would mean to recreate, within the scope defined by the website and its discipline, something of the “infinite” web of connections between data and interpretations alluded to above. The dual aim is to maintain editorial control of all annotations, text contributions, and the creation of hyperlinks, while facilitating broad collaboration with scholars who can contribute to enriching a multiplicity of narratives in dialog.

### **Bibliography and notes: in support of the core narrative(s)**

The bibliography of *When on High the Heavens...* was entrusted to the companion website instead of being printed in the book. This means that the bibliography can grow in time. It also

---

<sup>238</sup>Since this aspect of the site does not yet exist, it seems premature for me to write more than a footnote about it. But a few aspects can be intuited and discussed at this stage. The importance of the human community behind the research project was clear even at the outset of our work, which was characterized by a small group of scholars that met weekly to develop the site. Our meetings led to serendipitous discoveries, and the collaborative energy between us was sometimes like a fire igniting other fires. Sometimes, we disagreed about core issues and had to decide how to permit the coexistence of a multiplicity of interpretations. One important way to accomplish this goal is the simple, but crucial insistence by G. Buccellati that all contributions contain their author’s name. Everything from field observations, bibliography entries and notes, to longer thematic articles and monographs, bear the name of their author. Going forward, the issue of collaborative team dynamics and structures is worth deep consideration. In order to include multiple narratives in future editions of the site, the editorial team will need to manage permissions, and will probably need to render contribution as technically simple as possible, since hurdles like coding in html or in markdown would diminish the number of contributors who are well-versed in the topic and are also able to directly contribute in that format. Requiring all contributors to modify bare html or markdown files could also lead to an overall brittleness of the site, since errors in their formatting can easily cause entire pages to break. Some form of CMS (content management system) might be useful to solve this issue, although it would have the cost of instituting a layer of abstraction, with the attendant risk of technical obsolescence as programming languages and programs develop and decay. Alternative narratives could be difficult for the reader to parse if they reside at the same visual level as hyperlinks, so it might be necessary to distinguish them graphically in some way. Disputed interpretations of underlying facts or texts would need to be handled in a transparent way. The experience of other projects like Wikipedia might be useful as a comparison point for learning how to manage growing teams of volunteer writers. The present structure, in which permissions are granted directly from the editor to contributors, and the contributors must learn to use a specific markdown language in order to write their articles, is expressive of an underlying preference for a strong understanding of what “writing a website” entails: the narratives are the product of a small group of scholars who have been specifically invited to contribute, not the work of a larger and looser group of occasional contributors acting on their own volition.

permits a kind of bibliography that is rare in printed works: instead of listing only the most basic information about source material, each entry is enriched with a short summary of the work. The summaries aim to bring out the specific relevance of the work for the core narrative on the website, not as a generic summary of the entire book.<sup>239</sup>

Footnotes pose a particularly difficult choice for authors. Narrative linearity is important for the sake of readability, while footnotes allow for the creation of explicit lateral links to sources as well as the partial exploration of related issues, within the context of the larger narrative arc of an article or book. When the number and length of footnotes exceeds some minimum, and when the importance of the lateral links is sufficiently high, the reader will have difficulty experiencing the text as a sustained argument, and will instead be brought to consider diffusely some portion of the wider web of information referred to in the notes. This conundrum is familiar to any scholar, and is one of the issues that we hope to address in our attempt at writing “digital books”.<sup>240</sup>

In order for *When on High the Heavens...* to maintain a readable format, the author preferred to entrust the footnotes to the companion website instead of to the printed volume. This choice has several consequences beyond readability. The footnotes can grow in time as more scholars annotate the core narrative and more sources are added to the bibliography. Notes can also be linked from multiple places across the website, wherever they are relevant. Links to relevant notes can be found, for instance, in the bibliography, in the core narrative summary pages, and in thematic essays, as well as in a variety of indices.<sup>241</sup> As with all the other contributions to the site, notes carry the name of their author, which guarantees

---

239This fact means that future multiple narratives could refer to the same underlying source materials, and their bibliographic entries could be modified accordingly. This possibility is exciting, since it would help readers maintain detailed control of the data from high-level abstractions all the way to the fundamental data that supports interpretation. However, if this were to be undertaken, it would create issues of versioning and presenting information to the reader in a way that maintains transparency about authorship and interpretation. For instance, conversation between scholars could enter the bibliographic entries themselves, and point out disputed issues even at that granular level. Exactly how to grant permissions to modify these entries, and how to present a multiplicity of views, would have to be carefully decided in order to avoid a chaotic result.

240For a more detailed exposition of the theory behind our practice, see the Digital Discourse website at <https://d-discourse.net/>.

241In several places on the website, these links can be automatically created. For instance, notes that refer to works in the bibliography are automatically listed next to the relevant entry, allowing for bidirectional movement of the reader’s attention between sources and footnotes.

transparency and responsibility among the editorial staff. There is space for personal expression, and personal responsibility for the views expressed.

### **Excerpts, reviews, sources**

In some cases, source materials can and should be included directly in the printed book and on the website. A selection of texts presented in the appendix to *When on High the Heavens* aid the reader in understanding how some of the author's conclusions are supported by the primary sources. A similar, and broader, opportunity is offered on the website.

Not only does the site contain the sources present in the printed book (with all the usual advantages of digital text, like hyperlinking, automatic search and indexing tools), but it has also been enriched with alternative translations to improve the transparency of our interpretive choices. The site also contains other sources that were unavailable or had not been considered at the time the book was printed. Some of these are presented in the form of extended "reviews"<sup>242</sup>, like a longer and more articulate form of the summaries present in the bibliography. Others are presented as collections of direct "excerpts"<sup>243</sup> which offer the reader convenient access to relevant passages from several classics of the field. This too is a way to control the details of the available data, present them transparently to the reader, and weave them into a digital discourse that is hyperlinked together.

### **Themes and monographs**

As we worked on the bibliography and notes, we noticed that some topics required deeper treatment than had been given in the printed volume. Here too we find an advantage in writing a "digital book": such topics can be explored both by the author of the core narrative and by other scholars at a later time. See for instance the thematic exploration of the concept of the "affecting presence",<sup>244</sup> which was written years after the publication of *When on High the Heavens...* because the topic seemed insufficiently expressed in the written work. Such addenda have traditionally been possible only in subsequent printed editions; on a website, there can be a continuous enrichment of the underlying argument, as well as the establishment

---

242<https://4banks.net/Mes-Rel/synopses.htm>

243<https://4banks.net/Mes-rel/excerpts.htm>

244<https://4banks.net/Mes-rel/temi/affecting.htm>



of versions which constitute new editions of the site.<sup>245</sup> And, in the future, some themes could be further expanded into the dimensions of monographs in their own right, or even new “core narratives”.

## **Indices and ‘maps’**

A final feature of the website, which brings us to one of the specific focal points of this chapter, is a series of indices created to offer the reader an additional way to traverse the web of information collected on the site. Keywords chosen by the authors of bibliography entries and footnotes are collected in alphabetical, author name, and chronological lists. An additional set of keywords automatically generated by Natural Language Processing software are presented as well in a “multinodal” format described in detail below.

While building the linear indices, I noticed that the keywords chosen by the author of a summary barely cover their semantic space. It seemed important to find a better way to parse the data contained in these files, and create more useful indices. At the same time, it seemed useful to find a way to represent the connections between entries, as well as their internal data. A linear index represents information atomized into constituent parts; another useful tool would be a representation of the similarities and connections between entries.

When the eye scans a linear index, it can quickly sort through the words of interest and those that are not related to the questions that currently interest the reader. But the list of keywords is “flat”, containing little information about the relations between words. It is true that the number of references carried by a keyword give some indication of its relative importance in the present context, but a second layer of information is entirely hidden. Keywords form clusters: some are closely related, and define a domain of information. Others are simultaneously present in the overall website, but never co-occur in an article or bibliographical summary. This fact led to the development of the “multinodal index”. By seeing the connections, and the lack of connections, between keywords, it is possible to navigate the site with greater depth and insight. Such a visualization provides a “map” to the subterranean connections between ideas, authors, and books.

---

<sup>245</sup>On this topic, and on the related “ephemeris” idea, see Buccellati’s explanation at <https://4banks.net/Mes-rel/archives.htm>

The multinodal indices are tools to navigate the underlying bibliographical data of the website. They are built by automated extraction of keywords from the summaries present in the bibliography. Keywords are ranked by co-occurrence in other summaries that contain the same keyword, and ranked by strength as measured by the number of occurrences in a given reference. Then they are plotted as a force-directed graph. This visualization permits the user to explore connections that are implicit in the data, but not explicitly coded in the form of notes or hyperlinks.<sup>246</sup>

## Automation

Each of the *fragments* we have included in the website, from the tiny notes and bibliographical summaries to the longer thematic articles, is a contribution to a *whole* discourse that is the product of the editorial team. Our core narrative draws on the fragments and weaves them together into a multiplanar discourse, which attempts to take all the data into account in its overarching narrative. Some of the fragments can be handled in a way that mixes direct human processing with automation. This is the case of the indices.

We can think of the creation of our indices as an instantiation of an “artificial language”. Consider the case of the keyword indices. The “natural language” used by the authors of the site has been enriched in some cases with manually coded keywords. Keywords can then be ordered alphabetically, by using algorithms described in programming languages. These languages are an intermediate abstraction to help humans command computers to carry out specific operations: the input information (keywords) are automatically ordered according to criteria established by the human writers of the website. In this case, automation is a way to create a dynamically updated index that remains ordered even as new information is added to it

---

<sup>246</sup>The site also contains other forms of indices. For instance, the author search function initially returns the “keyword space” of a given author, and links to other authors who use the same keyword in their works. Most elements are clickable, returning more information or direct links to the relevant bibliographic entry. The author comparison visualization overlays two authors over the entire “keyword space” that represents all authors contained in the bibliography. It can give an indication of how completely an author covers the entire space, and how closely two authors’ interests intersect. It is important to note that this representation is based only on the bibliographical material present on this site: therefore, it does not accurately represent the actual content of the author’s body of work, but only the subset present on this site, filtered through the lens of the scholars who summarized the work(s) in the bibliography. When the comparative tool indicates an interesting intersection or non-intersection of two authors, the reader can follow the bibliographical links for the context in which to interpret the results.

over time, unlike the situation of a printed book. A human could certainly produce the ordered list of keywords, but it would be a mind-numbing task to collate such a list, given that the number of distinct keywords runs in the thousands. Better to have the computer keep the list in order!

In the multinodal index keywords have been *automatically* extracted from the bibliographical summaries.<sup>247</sup> These keywords can also be manipulated automatically. We have chosen to represent them as a web of concepts that interconnect the contents of the bibliography. It would be difficult for a human reader to perceive this web of semantic connections without an automated visualization: here, automation serves not only as a way to speed up a process that humans could do unaided, but also to offer a tool to see connections that would otherwise remain hidden.<sup>248</sup> Automation can help surface fragments and their connections to the whole that may not be easily found by the reader. It can provide an additional level to the perception of the whole, without replacing the central importance of the humanist responsible for creating the narrative, regardless of how digital their tools may be.

The distinction between natural and artificial languages can help keep this point clear. Artificial languages like formal logic and programming languages produce deterministic results from their inputs, and can be used to carry out operations to order and cluster large quantities of information. Natural languages, instead, are the domain of humans who are engaged in weaving a whole out of the fragments.<sup>249</sup>

As the fragments proliferate, it becomes more and more difficult for scholars to produce a “whole” that takes into account all the relevant information. Some form of automated process

---

247For details about how this automatic extraction is accomplished, see

<https://4banks.net/Mes-rel/multinodal.htm#keywords>

248A further use of automation could be the grouping of keywords by semantic similarity, a task now well within reach of many modern “large language models” (LLMs) and word embeddings.

249The recent development of ChatGPT and similar LLM technologies has caused much debate and some consternation about just how “natural” human language is, if it can be quite successfully imitated by statistical processes running on computer programs. In this connection it might be worth remembering the Canard Digérateur created by Jacques de Vaucanson and unveiled on 30 May 1764 in France. De Vaucanson’s duck “ate” food and produced faeces—and so seemed to be a successful simulation of the digestive tract of a duck. However, the internal process was completely different from that of a duck. The fact that a language model can take input and produce output that is similar to what a human would produce does not necessarily imply that the real process of human language production has been discovered and instantiated in a computer. Joseph Weizenbaum’s program ELIZA, an automated psychotherapist modeled on the Rogerian school, was a similar example of the simulation of understanding and human communication.

could help serve as a “telescope” with which to observe the overall information present in the field of study under examination. Naturally, a telescope is only useful if it is directed toward relevant objects and interpreted by a competent professional. The tool cannot direct itself. But without the tool, a great deal of relevant information remains invisible, and human proposals of “whole” theories can easily be incomplete or false.<sup>250</sup>

This chapter began with the observation that while reading takes a linear form, writing draws on a much more complex web of information to inform its linear distillation. That underlying web can be made explicit in a way such as we have done on the Mes-Rel website, which contains direct hyperlinks between supporting data, bibliographic sources, and footnotes by a research team which can be algorithmically manipulated by programs. If we can imagine an article “bristling” with footnotes, we can also imagine the further connections between more granular elements of the footnotes as a complex web made up of the semantic similarity between words and concepts, or the representation of related works that interpret the same issue in alternative ways, or authors who take some information into account and ignore other details, or who have developed and modified their interpretations over time...<sup>251</sup> We can further intuit that the webs that pertain to two competing “core narratives” could be compared directly through a variety of metrics, including their invariant shape.<sup>252</sup>

This intuition brings us back to our Graph Object, which we can now perceive as a complex reality that lies behind the many different forms of digital writing I have described above. In the Mes-Rel website, the Graph Object is a detailed map of the domain we have

---

250The classic case in point, of course, is the overturning of the geocentric view of planetary motion, in favor of the heliocentric model. This discovery was made possible by close (telescopic) observation of planetary motion over time by Galileo, Copernicus, and Kepler.

251Graph databases which model the connections between “atoms” of information have become popular as a way of representing networks of nodes, whether they be social networks, networks of citations and bibliographical references, networks of travel and trade, or still other forms of interconnected data. For instance, some businesses use graph databases in order to detect fraud by identifying the structures that are typical of legal business transactions and the relations between entities and persons involved in such transactions, and transactions which are most likely fraudulent. Fraudulent transactions may have the appearance of being normal transactions when seen from the level of a bank teller, but when viewed with the “god’s eye view” of a graph database they can be seen to contain illegal transactions. Similarly, a citation network can distinguish authors who are prolific but isolated, from authors who are central to a field.

252Some of the most interesting aspects that concern our topic involve the definition and measurement of “holes” in the network or graph, the calculation of the relative centrality of nodes which act as a sort of “gatekeeper” of information flow in the graph, and the relative complexity of interconnections overall.

encoded in bibliography, footnotes, primary sources, and core narrative. It is so detailed that our physical senses cannot perceive it all at once. But if we slice it one way, we see the outline of a single book: all the nodes connected by “PRECEDES” relations. If we slice it another way, we see clusters of keywords in an abstract space that represents the closeness and distance of the topics in our domain. Slice the Graph Object a third way, and we can view the branching logical structures that lie behind narratives, and determine downstream consequences of which their authors may not even be aware.

Perhaps another allusion to *Flatland* is in order. Higher dimensional objects reveal a variety of shapes when “sliced”: depending on the angle of the sectioning, cones can be sliced into lines, circles, ellipses, parabolae and hyperbolae. All those ‘conic sections’ can be seen as the two-dimensional shadows of a three-dimensional object. How many more structures might we find within the 1536-dimensional web of connected word vectors in our Graph Object!

Some work of this nature has already been done in the digital humanities, such as the creation and representation of citation networks, but many more developments could be imagined to give the humanist a view of the information available that is both broader and more detailed.<sup>253</sup> Distinguishing wholes and parts, and writing in such a way that one’s product is a “whole”, however small, is a crucial characteristic of scholarly work. It is becoming increasingly difficult as the overall amount of relevant information grows. One solution is to address tightly-defined questions in order to preserve the wholeness of one’s production, but this road is not open to those who wish to propose a large-scale interpretation of a culture, such as the authors of the Mes-Rel website attempt. We present an attempt at developing a coherent, “whole” argument by giving the reader transparent access to our sources and the granular considerations that inform our interpretations. Some of its aspects are electronic versions of mature analog technologies, and other features are attempts at creating a form of writing that directly leverages the affordances of digital technology to further the aims of scholarship<sup>254</sup>. This project provides an important test case for the Memex described in Chapter 3.

---

253 CitNetExplorer (<https://www.citnetexplorer.nl/>) is an online tool that allows viewers to explore citation networks directly. Other similar tools are available on online repositories like PubMed and arXiv. Much research using these tools exists. See for instance Liu, H., Kou, H., Yan, C. et al. “Link prediction in paper citation network to construct paper correlation graph.” *J Wireless Com Network* 2019, 233 (2019).

254 Make sure this is covered earlier

## Validating the Mes-Rel Website

For the space of more than a year while writing the Mes-Rel website, I met weekly with a team of scholars who engaged in writing footnotes to the text and adding sources to the bibliography. Our team hailed from several disciplines: history, archaeology, philology, and theology. Our conversations aimed at locating controversial points in Buccellati's argument and rendering them transparent and well-supported in our critical apparatus.

Insofar as the website is a broad comparison between Mesopotamian and biblical spiritualities, there are scholars who hold different interpretations on every significant point. To return to the terminology I used in the first chapter, first our research team created mental maps of the patterns of the information we had available, then we created narratives and models of the information, and finally we compared these models with those created by other scholars. We had to argue our position from the basis of the known facts, and extrapolate from them to a plausible interpretation. We also had to arrive at a fairly unitary, compact reading, the conclusion of a long process of study and interpretation.

Once I began to see the complexities involved in this project, and the difficulty in remembering all the different interpretations and patterns of reasoning, I wished for a system that would help us to remember the granular data together with our broader interpretations, and grow along with our work of pattern-comparison. It seemed to me that an important aspect of scholarship, beyond the specific site we were building, was to retain information from every source and patiently construct an interlocking web of information that is itself a model, or map, of the domain, and contains within itself the interpretations we developed and continually improved.<sup>255</sup>

Sometimes, the conversation between team members led us to identify ideas or authors we needed to include in our bibliography. Other times, we realized that we needed to seek alternative points of view to fill out a complete and intellectually honest vision of the topics

---

<sup>255</sup>Our work seemed to me to be an attempt at creating an all-encompassing model, an interpretation that had no "outside". Insofar as we aimed to map the domain of Mesopotamian religious studies sufficiently to support our narrative, we aimed at completeness; insofar as we positioned our interpretation as a contribution to scientific discourse, we did not argue from theological or faith-based positions. However, we did accept that an "outside", a world of the spirit, might very well exist and be interoperative with the "inside" of our experience, as religious people of the past have believed.

under discussion. And sometimes, our collective intelligence wandered among the connections we could intuit between authors, archaeological evidence, ancient texts, etymologies, and mythologies. It was a true experience of a community of scholars.

We were building the [4banks.net/Mes-Rel](http://4banks.net/Mes-Rel) website, an online companion to the book *When on High the Heavens*. The site contains an extensive bibliography related to the theme of Mesopotamian religion, and as its introductory page clarifies, it addresses studies about

(1) Mesopotamian religion, understood as a cultural system of beliefs about the absolute, that can be defined in terms of the external aspects of expressed principles and identifiable elements, as well as in terms of the actions that are aimed at establishing a relationship with these elements, for both the individual and the community; and (2) Mesopotamian spirituality, understood as the underlying and unifying perception that can be interpreted as both the driving force and the goal of the religious system. The Mesopotamian system of religion and spirituality is contrasted with its counterpart in the biblical world, seen as a cultural phenomenon that is at odds with the one elaborated in Mesopotamia even while depending on it for its cultural aspects. It is narrative based, in the sense that it uses at its core a frame of reference to which the individual studies are linked as parallel arguments. It also expands selected topics into larger treatments. Serving as a point of reference for the field, it remains active through the work of a Research Group that provides input on an ongoing basis.<sup>256</sup>

Following a procedure developed by Buccellati and his team in previous websites, in particular [critique-of-ar.net](http://critique-of-ar.net), we encoded bibliographical information about the hundreds of books and articles pertaining to the overall theme of the website that we were studying, in a format suited for the program DABI (“Digital Analysis of Bibliographical Information”) written by Bernardo Forni for use in this website and others like it. Each DABI file contains bibliographic codes for author name, title, publication information, as well as a summary of the work in question. There are also codes for keywords and for notes, and the use of “markdown” to encode more complex formatting that cannot be directly supported in an ASCII text file.<sup>257</sup>

A typical example taken from my contributions to the site follows:

---

<sup>256</sup>Written by G. Buccellati. See [4banks.net/Mes-Rel/home.htm](http://4banks.net/Mes-Rel/home.htm)

<sup>257</sup>Briefly, markdown entails the use of simple character codes to represent formatting. In the above example, asterisks enclose **bolded text**. For details on markdown, see [www.markdownguide.org/](http://www.markdownguide.org/). The reasons for using a simple text file to encode our information has mostly to do with portability and durability issues. Computer systems rapidly change and are updated, and many programs and formats that were in wide use only a few years ago are now obsolete. Since a study of history and archeology takes a long view, we prefer to use formats that are unlikely to ever go out of date, and in any case can be easily ported to new systems when that becomes necessary. Although this is perhaps not possible with ever-updated programming languages, it is possible for data files such as our DABI files. In fact, the bibliographic material encoded twenty years ago in this format is still interoperable with the latest version of the software.

AU Mendenhall, George

Y 1975

T The Conflict Between Value Systems and Social Control

P in Goedicke and Roberts, \*Unity and Diversity\*

@@@R

SA jJL

SD March 2020

TO Monotheism; Polytheism; Covenant; Law

This insightful article distinguishes sharply between **Biblical monotheism** (particularly in its moral structure) and **Mesopotamian polytheism** (which Mendenhall likens to contemporary North American society, and liberalism in general).

The <<structures of control systems>> (political or social power) and values, while opposites, are not necessarily at war with each other. <<The clear message of the Bible is that social systems can continue to exist only where a minimally tolerable **value system** has already become operative>> (p. 171). Further, <<the mainstream of biblical faith is ... the affirmation of a real factor in human life and experience that is independent of, not produced by, but ultimately essential to the existence and satisfactory operation of any social control system.>> (p. 171)

The second part of the article is an outline of what could be a <<monumental ten volume work>> if the author had the energy to complete it. He compares <<Covenant>> with <<Law>> along 10 different axes, and demonstrates the radical difference between the two. **Covenant**, in brief, creates *ex nihilo* a society oriented to the good; **law** presupposes a society and attempts to exclude the negative by enacting a war between society and the transgressor.

The competition between power structures is insoluble in their own terms alone. <<A sense of **justice and right** ... must take precedence over selfish interest. The permanent symbol of the necessity as well as the reality of that Rule of God is the crucifixion of Jesus--the equally permanent affirmation that <<winning>> in the jungle of social manipulation and social competitiveness cannot be the controlling motivation of those who hunger and thirst for righteousness and peace.>> (p. 178).



This file pertains to a 1975 article by George Mendenhall. The “@@@R” code means it is part of the religion website (the distinction is necessary because there are several other topical websites connected by the same underlying system). Our website contains more than 550 such files, which describe as many sources. One reason to create the DABI program was to dynamically update the website as new data files are added. Thus, the bibliography is modified and reformatted in proper order and page layout every time a new file is loaded into the system. Another reason to use the program is to create dynamic indices which list information in alphabetical, chronological, or keyword-based order, updated when a new file is uploaded.

As our collection of files grew, the indices grew apace. At the time of writing, there are 2721 keywords listed. Keywords are chosen by the author of each summary (as can be seen above in the line “TO Monotheism; Polytheism; Covenant; Law”, where TO stands for “topics”).<sup>258</sup> It is common for different authors to use closely related, but different, words or phrases for similar entries in the index. Even the same author can forget which specific keyword they used weeks or months before. Therefore, “intransitive” and “intransitivity” appear as distinct entries, as do “Asherah” and “asherah”. More sophistication in the programming could eliminate some of these duplicates, but at the level of detail that initially interested us, they are not too important. They would become important, however, if we wished to automate the creation of links between keywords, which a computer would recognize as identical only if they were indeed identical, down to the upper- or lower-case of the characters.

## **The Memex as “X-Ray”: The Multinodal Index**

Given the large amount of text that makes up the critical apparatus of the Mes-Rel website, indices are a useful way to navigate through the complexity and quickly find relevant information. Yet the few keywords chosen by the author of a summary barely cover the semantic space of the summary itself. Above, the entry for Mendenhall’s 1975 article is poorly characterized by the four keywords I chose. The keywords are so generic as to be applicable to nearly anything in the bibliography!

---

<sup>258</sup>See [4banks.net/Mes-Rel/sub-indice.htm](http://4banks.net/Mes-Rel/sub-indice.htm).

I wondered if we could find a better way to parse the data contained in these files to create more useful indices. At the same time, I wondered if we could find a way to represent the *connections* between entries: their *structure*, as well as the granular data. I wanted to create a sort of “X-Ray machine” to view the “bones of thought”, the hidden structures linking items in the bibliography.

A linear index represents information atomized into topical keywords. When the eye scans a linear index, it can quickly sort through the words of interest and those that are not related to the questions that currently interest the reader. But the list of keywords is “flat”, containing little information about the relations between the words and their sources. It is true that the amount of references carried by a keyword give some indication of its relative importance in the present context, but a second layer of information is entirely hidden. Keywords form clusters: some are closely related, and define a domain of information. Others are simultaneously present in the overall website, but never co-occur in an article or bibliographical summary. These facts led me to develop what we call the “multinodal index”. By seeing the connections between keywords, it would become possible to navigate the site with greater depth and insight.

In order to create this tool, I developed an automated method for extracting keywords from the text of the DABI files themselves, as a supplement to the keywords chosen by the summary authors. Two distinct indices could then be generated: one reflecting each scholar’s judgment about what is most important, and another that contained an algorithmically generated pathway to every significant word contained in the entire corpus on the website.

Extracting keywords from plain text is not an unusual task, and many methods exist to do it. One method consists in breaking a text into words at every “whitespace” character, and throwing away “stop words” like articles and prepositions. A further step might be to make all words lowercase, so as to eliminate one of the errors mentioned above, where “Asherah” and “asherah” appear as distinct keywords. However, this procedure is unable to account for plural and singular, for proper names, and for words like “through” that are complex enough to often escape filtering for stop words, but which are not real keywords. One might also prefer to filter

adjectives and adverbs from a keyword list, with the help of the language models I discussed in Chapter 2, and the SpaCy package that simplifies their implementation.

A few lines of code provide the basic functionality:

```
nlp = spacy.load("en_core_web_md")
doc = nlp(sentence)
for t in doc:
    if len(t)<3: #get rid of short words
        continue
    if t.pos_ == 'NUM' or t.pos_ == 'AUX' or t.pos_ == 'ADP' or t.pos_ == 'SYM': #get rid of numbers,
        auxiliary verbs, prepositions, symbols
        continue
    if t.pos_ == "PROPN": #put proper names in the list with capital letter
        topic_list.append(t.lemma_)
        continue
    if t.pos_ != "NOUN": #keep only nouns
        continue
    topicToAdd = t.lemma_
    topicToAdd = topicToAdd.lower() #make everything except proper names lowercase
    topic_list.append(topicToAdd) #append the lemma of each word to the list
```

In this code, first the English language medium sized model is loaded as the `nlp` object. I chose the medium sized model as a good compromise between speed and accuracy, and because the smallest model does not contain word vectors, which are essential for my purposes. Next, each sentence of the document (previously parsed) is passed through the `nlp` model and assigned to the `doc` variable as a list of words with part-of-speech attributes, word vectors, and other syntactical information. All of this is hidden below the surface, which makes SpaCy exceptionally easy and clean to use.

It is then possible to iterate through each word in the sentence: for each word token (referred to as `t` in the above code), the program carries out a series of tests. If the word is only one or two characters long, the program skips to the next word without further action (`continue` means to skip over all subsequent instructions, and move to the next word in the sentence). If

the word's `pos` (part of speech) is Number, Auxiliary verb, Preposition, or Symbol, it is skipped. Proper names are included in the list of keywords as lemmas, maintaining the capitalization of the first letter. Then, all words that are not nouns are skipped, which leaves only nouns to be included in the list of keywords, since the last few lines of the loop are only reached in this case. They are first lemmatized, so that prefixes and suffixes do not distract the program into considering identical lemmas as different keywords. Finally, they are converted into lowercase and added to the variable named `topic_list`, which is the list of keywords.

It may be noted that the logic of this simple program is partially redundant: could it not distinguish nouns from all the other parts of speech without the first line? In testing, I found that it was more error prone to place the major discrimination element entirely on the shoulders of one command, and that by first eliminating some parts of speech I could more accurately identify the nouns. On its own, SpaCy is only 87% accurate in labelling the part of speech of each word.<sup>259</sup> I reasoned that it was more important to avoid spurious keywords than to ensure that all keywords were represented, since the goal of this index is to help a human user find useful information. Other use cases might choose a different priority.

The rest of the program manages the connections between information in order to generate a table such as this (the actual table contains hundreds of rows and dozens of keywords per row; this table is truncated for brevity and clarity):

	<b>filename</b>	<b>author</b>	<b>year</b>	<b>keywords</b>
0	Nissinen2019Divination	Nissinen, Martti	2019	[divination, essay, Eastern, Prophecy, prophec...
1	Cohen2018MoralityInAntiquity	Cohen, Yoram	2018	[morality, Antiquity, tradition, proverb, comp...
2	Scheil1913Memoires	Scheil, Jean Vincent	1913	[Mémoires, mission, Susiane, Vol, dream, Appen...
3	Rubio1999Substratum	Rubio, Gonzalo	1999	[Alleged, Substratum, substratum, problem, Sum...
4	CDLI	None	2021	[Cuneiform, Digital, Library, Initiative, cune...

<sup>259</sup>See [spacy.io/usage/facts-figures#benchmarks](https://spacy.io/usage/facts-figures#benchmarks) for benchmark figures regarding performance, which depends on the size of the model used and the specific task performed.

Some errors are clear in these lines: “Alleged” should not have been included, and “Substratum” and “substratum” are duplicates that were not properly united, probably because the model incorrectly labeled “Substratum” as a proper name. Also, the French language words in Scheil’s text confuse the system because it only is able to use one language model at a time, and cannot distinguish French words from the English in the surrounding text. Thus, the French text is incorrectly labeled. Still, it is an acceptable rendering of the keywords of each bibliographical entry, not messier or more redundant than the keyword list generated by our research team.


The next step in creating our multinodal index was to calculate the strength of connection between each node representing a single work in the bibliography. How to do this is a question of how to model the data, and is a crucial interpretative step. I decided to consider two nodes (bibliography entries) as “connected” if they share at least five keywords, and not connected if they share less than five. The strength of the connection can also be represented on a sliding scale based on the number of shared keywords. Similarly, the size of each node is determined by the number of keywords in it. Larger nodes represent more detailed bibliographical entries, with more keywords.

Once this calculation is performed, it is straightforward to plot the nodes and their connections as a force-directed graph. I chose this representation because it offers a simple and intuitive map of the density, or lack thereof, of groups of nodes. I was interested in seeing which bibliographical entries covered a similar “semantic space”, and which were isolated from the rest.

Force-directed representations are based on the physics of charged particles connected by springs. Each “particle” has the same charge, and repels all the others, and each “spring” attracts the two particles it connects. Thus, in my representation, each node (or bibliographical entry) has a “charge” that is proportional to the number of keywords it contains. Left to themselves, each particle will repel all the others and fill the space with a maximum distance between all particles. But many of the nodes are connected with “springs” that pull them together, with a force in proportion to the number of keywords they share. This means that the

space will be filled with clusters of related nodes, and less-related clusters, or single nodes, will be pushed away, where they can easily be identified.

A simple force-directed graph is below<sup>260</sup>:



From this representation, one can immediately see that two light orange nodes on the left are unconnected to the rest of the graph. It might be relevant to understand what they are, and why they are so eccentric with respect to the rest. Similarly, one can see that a small number of large nodes act as “gatekeepers” holding the network together, in a sense controlling “traffic” between less well-connected nodes. In my representation, this might mean that they are among the more important works in the bibliography, since they have been described at greater length<sup>261</sup>.

I applied this method to the table of bibliographical entries and keyword lists on our Mesopotamian religion website, and filtered the network to represent only those works that contain a given keyword. Then, with the help of Ann Blevin of Blevmore Labs, I created the

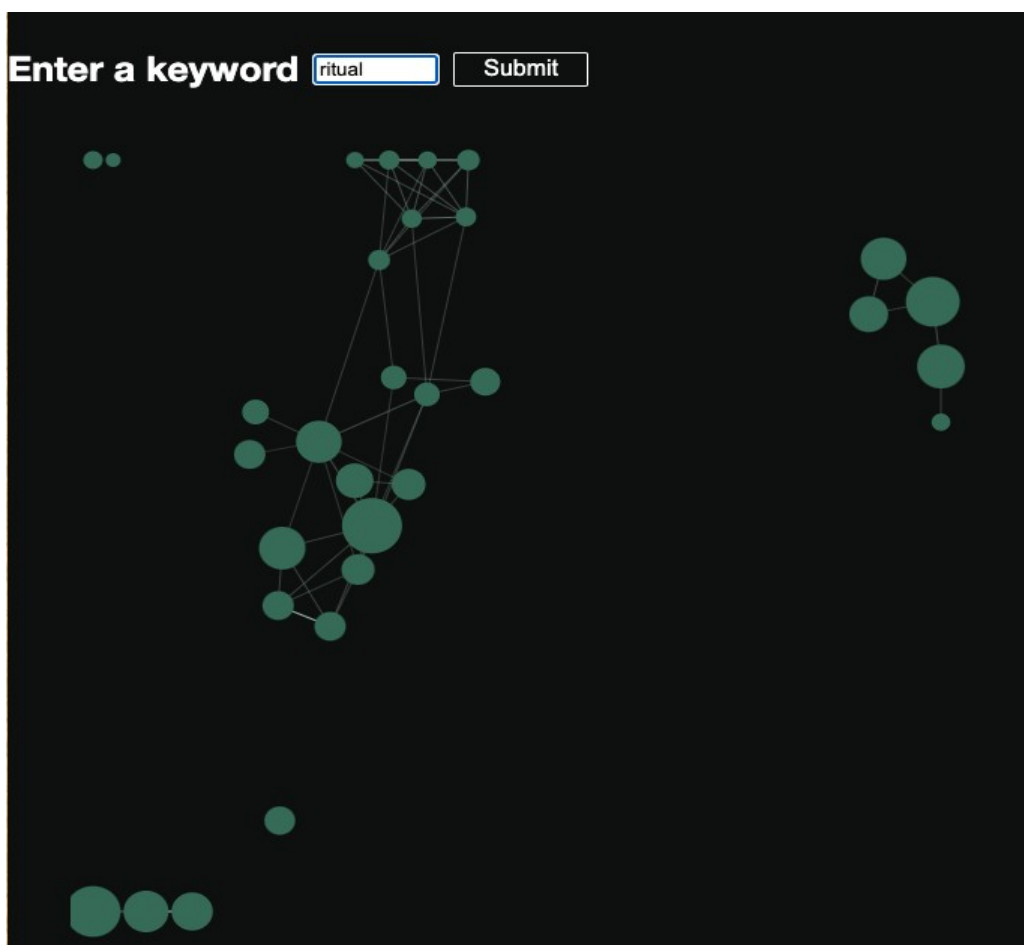
<sup>260</sup>Source: <https://repository-images.githubusercontent.com/140739735/d9f6bc80-ae30-11e9-86e4-f8bd1b99a4aa>

<sup>261</sup>Of course, it could also be the case that they are longer because they were summarized by a more verbose scholar! At each step in this process, many factors contribute which are not necessarily significant to the overall interpretation. This is inescapable, and must be kept in mind.

visualization software.<sup>262</sup> The keyword search function navigates the lateral connections that exist when the same keywords are used by a group of authors. The force-directed graph of keywords can be interpreted as a visual representation of the strength of connection between the works contained in the bibliography. Each node represents a work in the bibliography that contains the search term. The size of the node represents the number of keywords contained in that work. This implies that larger nodes have more detailed entries in the bibliography.

Connections between nodes are formed when two works share at least five keywords. The connection strength (=number of shared connections) is graphically represented by the color of the line connecting two nodes. The brighter the line, the stronger the connection.

Below is an example of the keyword representation of all works containing the word “ritual”:



---

262 See [www.aesizemore.com/design.html](http://www.aesizemore.com/design.html)


The network structure gives us an intuitive impression that corresponds well to the contents of the bibliography. The largest node in the largest cluster refers to the chapter on religion and magic in von Soden's classic *Introduction to the Study of the Ancient Near East*. The more tightly connected components of that cluster deal with the topic of ritual and religion in the ANE, whereas less tightly connected components are more specialized: most of the seven-node cluster at the top of the image deals with the *namburbi*. There are also two unconnected clusters that contain works about other religious traditions, the traditions of Çatalhöyük and Christianity.

In this case, the network visualization offers a guide to the content of the bibliography which can enrich the work of the researcher: it indicates that the current bibliography focuses on the ANE, while considering other traditions as well, and directs the researcher to the clusters of relevant works, whereas a linear index does not reveal the subclusters of works that contain a given keyword.

The force-directed graph representation, then, allows one to see the hidden structure of the thematic connections between distinct works in the bibliography. It offers an intuitive navigation of the bibliography that contains more information than a simple keyword search could provide. Remember that in a keyword search, or in a linear index, works are not distinguishable based on the "semantic space" they cover. All works that contain a given keyword are on the same level, so to speak. By using a force-directed graph representation, it is possible to give the user a bit more information about which works are most relevant, important, and well-connected to the rest of the present corpus.

In some cases, the network visualization does not add much to the experience that could be had by using the traditional linear index. For instance, searching for the word "tradition" returns a densely connected cluster of related works that is essentially identical to the linear index:






The only additional information in this case is found in the two unconnected nodes, which reference Diels-Kranz’s *Presocratic fragments* and the *Database of Religious History*, which indeed are tangential to the main contents of the site. In many queries, unconnected nodes appear at the periphery of the visualization without links to the core cluster. These nodes refer to elements in the bibliography that contain the search term, but do not share more than five keywords with any of the other works. They can be interpreted as “hapax” elements which are not corroborated by other sources in our bibliography. (It should be noted that in all cases, the content of the linear and multinodal indices is determined not by the works themselves, but by the summaries and keyword tags our research team has produced.)

The Author “keyword space” visualization shows a network of keywords found in the bibliography, filtered by author. In some cases, this visualization is either too simple or too complex to be of use as an intuitive guide. If there is only one work by an author in the bibliography (as is the case with the majority of authors), the keyword space visualization is no more useful than a simple word cloud. At the other extreme, if there are too many works by an author, the visualization is too complex to admit intuitive use—see for instance Buccellati’s


graph, which is highly complex because the bibliography contains 28 entries on his works. This level of complexity requires mathematical network analysis to provide useful information.

However, the Author keyword space visualization is useful in intermediate cases. For instance, Verderame's network appears thus:



From this image we can deduce that Verderame's interests cluster into discrete groups that do not share keywords. This reflects the contents of the bibliography, which contains three of his articles on very different issues: magic and divination, calendars and measuring time, and writing. We also notice eight unconnected nodes in the top left corner, which represent keywords that are not present in more than one of his articles, and therefore are not connected to the rest of the network.

In the case of Erica Reiner, we see a more complex and well-connected graph:




Reiner also covers a variety of topics: Šurpu incantations, the “seven sages”, wisdom literature, and poetry. But in her case, these topics share keywords and present as a connected graph. We can further notice that the topic that holds most of this network together is the word “king” (highlighted in the image).

The 4Banks cluster of websites aims to create a form of writing that is multilinear and multiplanar. In the context of this aim, the Multinodal Index visualizations were developed in order to permit the reader of the site to explore connections that are implicit in the data, but are not explicitly coded in the form of notes or hyperlinks. The Multinodal Index is thus a layer of algorithmic multiplanarity that supplements the interconnected layers that were explicitly written by the authors of the site.

The author search function initially returns the “keyword space” of a given author, representing it as a network of nodes (keywords) connected to other keywords that co-occur in

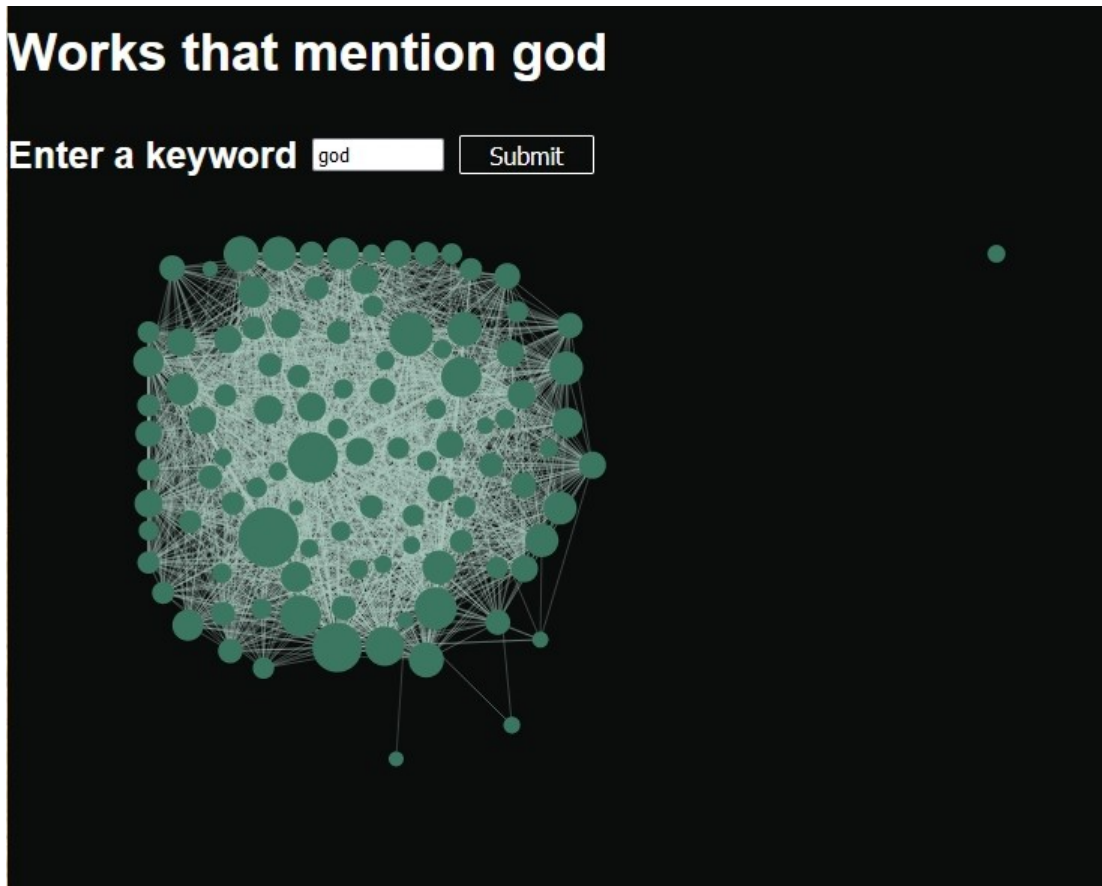
the same text. On the side of the image, there is a list of links to other authors who use the same keywords in their works. This visualization can be interpreted as a network of keywords found in a single author, across all their works contained in the bibliography. For example, Thorkild Jacobsen's contributions are mapped in this graph:



The largest, most central node is “Mesopotamia”, and we can see a complex web of connections to other words starting from that point. This representation can be interpreted as a semantic map that not only lists the main keywords in the bibliography of that author, but also provides a way to access the relative importance of those terms, on the basis of the size and centrality of the words represented. It can also be used as a way for the reader to notice peripheral terms that are not well connected to the central cluster, and thus represent issues that the author does not deal with closely. As usual, a caveat applies: since these representations are automatically generated from the material present in the site bibliography, they do not

necessarily represent the authors, so much as the website research team who have summarized their works.

Sometimes, these visualizations are too complex to be intuitively understandable. This is particularly true when there are so many keywords, or so many connections between the works of a prolific author, that the cluster appears messy and impenetrable to the human eye. For instance, a search for the term “god” produces the following image:



One feature that can be intuited in this case is the existence of an outlier displayed here at the top right, which contains the search term but is otherwise unconnected to the rest of the cluster. We can also identify several large nodes containing the search term, and infer that they are probably important works in the bibliography, because they contain many keywords.

More advanced analysis is also possible, although our website does not yet contain the tools to perform it directly. The “translation” of text information into visualizations and mathematical structures such as networks and UMAP dimensional reduction offer additional

tools to the historian which I will describe in the next chapter.<sup>263</sup> Using network analysis, topic clusters can be identified; central nodes (for instance, seminal works in the field) can be located and interpreted as “gatekeepers” that determine the flow of information in the network; and more complex topological features such as cavities can also be identified and interpreted.

As digital writing develops the capacity to explicitly render multiplanar discourses as well as the linear discourse we have used for the last six millennia, humans will need algorithmic tools to navigate the complexity of such discourse. The multinodal indices instantiate another aspect of my Memex, a tool to help scholars extend the capacities of their minds and see more deeply into their domain of study. They are like an x-ray machine, allowing a view of the database that is inaccessible through traditional indices because they do not transmit information about the connections between keywords.

## The Memex as “Bird’s Eye View”: Validating the Bibliography

After completing the first edition of the website, which comprises more than 550 sources in the bibliography, and several thousand footnotes, the question arose: how well does this website and its bibliography cover the topics treated in the volume *When on High the*

---

<sup>263</sup>Briefly, UMAP is a way to project the “keyword space” onto a two-dimensional surface. It can be likened to a two-dimensional shadow of a higher-dimensional structure, and like the shadow of a sculpture, it contains something of the original structure, while also being an impoverishment of it. Our visualization produces a map of the keywords present in the bibliography, and assigns them a location on a two-dimensional surface in such a way as to retain some of the high-dimensional structure of the keywords. In our case, the spatial closeness between keywords can be interpreted as a measure of their semantic similarity. Representing high-dimensional data on a two-dimensional screen is the object of intense research in data science. One of the first methods was called PCA. Another popular method, known as t-SNE, was developed by van der Maaten and Hinton in 2008. See [distill.pub](https://distill.pub) for an introduction to the method, with interactive visualizations that help see how powerful the procedure is, and how easy it is to misinterpret its results. Our multinodal index uses a UMAP projection of the keywords, which is a recently developed procedure (2018) based on t-SNE. An introduction, with interactive visualizations, written by Andy Coenen and Adam Pearce can be found at [pair-code.github.io](https://pair-code.github.io). The examples on that page offer an easy way to intuit the nuances of the problem of projecting high-dimensional semantic structure of the list of keywords to two or three dimensions. While our UMAP projection currently adds only a little to the reader’s ability to intuit the distribution of topics across the website, it is likely that further research will augment our ability to interpret the meaning of the projection space and increase the usefulness of this “map” of our website. Several recent studies have applied this technique to map domains of textual information, and it seems to be a promising approach. See Rita González-Márquez et al., “The Landscape of Biomedical Research,” preprint (Scientific Communication and Education, April 11, 2023), <https://doi.org/10.1101/2023.04.10.536208>. for an analysis of this approach applied to biomedical research. To my knowledge, I am the first to apply it to the study of history.

*Heavens...?* As the assistant editor, I was interested to know if there were topics in the book that were not adequately represented in the digital bibliography. I hoped an algorithmic approach could give me a view from above, a “bird’s eye view”, with which to view and assess our progress.<sup>264</sup>

In the last chapter, we saw a few simple examples of ways to use the Graph Object to characterize and compare groups of text. However, those examples were not particularly convincing, because their results could more easily have been reached by simply reading the (few) texts in question. In the present context, because of its dimensions, the Mesopotamian Religion website offers a serious test-case for employing the Graph Object as a useful method for achieving greater intelligence of history.

The “core narrative” volume examines eight aspects of individual religious practice (morality, divination, prophetism, apparitions, meditation, magic, prayer, and materializations) as well as eight communal aspects of religion (politics, narrative, representations, history, the temple, proclamation, worship, and the sacred). In the space of a few hundred pages, the sheer quantity of topics addressed is remarkable. Because of its size, simply reading this mass of material is an uncertain guide to the scholar who wishes to assess how well the critical apparatus serves the narrative.

In order to compare the book to the website’s coverage, I therefore input Buccellati’s volume and all of the supporting evidence gathered in the bibliography into a new Graph Object<sup>265</sup> by parsing the sentences and connecting them with relations as described in Chapter 3.<sup>266</sup> Because I was dealing with simple text files, it was a straightforward procedure to write a

---

264The next chapter will examine some of the further topological features that can be used to characterize the network of information gathered in our system. We will work up to that level of complexity by examining the three examples in this chapter, which explore the results of our method during research in order to find a fresh point of view on the work done so far; and as a validation check, after having written a text, to see how it compares with classic works on the same topic.

265At this stage, I considered only the summaries present on the bibliography website, not the full text in each work. This implies that the measurement of connections regards only what our team collected and explicitly wrote on the website, and cannot be taken to mean that the authors of those texts would recognize themselves in the considerations that followed. Still, as a proof of concept, this exercise offers a significant step forward in regard to the previous.

266This process may sound daunting, but it can be accomplished automatically with a simple program (in the appendix) that loops through the sentences in the text and inputs them into the database. The main difficulty in such a program is correctly distinguishing sentences. Common abbreviations like Ph.D or i.e. contain periods, so that character cannot serve as a univocal sign of the end of a sentence. The problem has been



script to load the texts into the graph automatically.<sup>267</sup> Once the graph was created, I began to explore its features. I was considering a large network and many writing styles in order to arrive at actionable information: if I were to find that some section of Buccellati's text is poorly represented in the bibliography, this information could guide our decisions in the near future while we prepare the second edition of the website.

The texts included in the present analysis are the bibliographical summaries contained on the website, of which two entries, including the Mendenhall entry discussed above, can be seen here:

The screenshot shows a web browser window with the URL <https://4banks.net/WEBSITE/DEV/Mes-rel/bibl.htm#M>. The page is titled "7. Bibliography" and features a navigation menu on the left with sections like "MESOPOTAMIAN RELIGION", "Introduction", "I. THE ARGUMENT", "II. THE RECORD", and "III. UTILITIES". The main content area displays two bibliographic entries:

- Mendenhall, George**  
 1975 **The Conflict Between Value Systems and Social Control** in Goedicke and Roberts, *Unity and Diversity*  
 This insightful article distinguishes sharply between **Biblical monotheism** (particularly in its moral structure) and **Mesopotamian polytheism** (which Mendenhall likens to contemporary North American society, and liberalism in general).  
 The «structures of control systems» (political or social power) and values, while opposites, are not necessarily at war with each other. «The clear message of the Bible is that social systems can continue to exist only where a minimally tolerable **value system** has already become operative» (p. 171). Further, «the mainstream of biblical faith is ... the affirmation of a real factor in human life and experience that is independent of, not produced by, but ultimately essential to the existence and satisfactory operation of any social control system.» (p. 171)  
 The second part of the article is an outline of what could be a «monumental ten volume work» if the author had the energy to complete it. He compares «Covenant» with «Law» along 10 different axes, and demonstrates the radical difference between the two. **Covenant**, in brief, creates *ex nihilo* a society oriented to the good; **law** presupposes a society and attempts to exclude the negative by enacting a war between society and the transgressor.  
 The competition between power structures is insoluble in their own terms alone. «A sense of **justice and right ... must take precedence over selfish interest**. The permanent symbol of the necessity as well as the reality of that Rule of God is the crucifixion of Jesus—the equally permanent affirmation that «winning» in the jungle of social manipulation and social competitiveness cannot be the controlling motivation of those who hunger and thirst for righteousness and peace.» (p. 178).  
 — Jonah Lynch, 2020
- Merlo, Paolo**  
 1994 *«L'ašerah di YHWH a Kuntillet 'Ajrud. Rassegna critica degli studi e delle interpretazioni»*  
*Studi Epigrafici e Linguistici sul Vicino Oriente antico* 11, pp. 21-55  
 This paper aims at better defining the figure of the **Ašerah**, sometimes interpreted as a possible pareadra of Yahweh, through the analysis of a depiction of this entity on a vessel from the site of **Kuntillet 'Ajrud**, in Syria (for some pictures of this site and of the vessel itself, see a paper of Ze'ev Meshel.  
 PDF available [here](#)  
 — Marco De Pietri, 2020

The summaries were written by the team of scholars, and show significant variety in their length and level of detail. In comparing these summaries and the core narrative we must keep in mind that the summaries reflect the judgment of the summary author as well as the contents

solved in a variety of ways; I chose to use the SpaCy package previously described, which does quite a good job of parsing sentences.

267The advantage of computer programming and automation here is that once the program has been written, it is an easy task to input even very large amounts of data into the system, and the speed of input depends mostly on the speed of the available computer—which is above all an economic issue, since powerful computers are available as cloud instances for a modest price considering their power. It is also an issue of technical complexity, since adding a further layer of technology (the cloud computer) increases the difficulty of the operation. But these are ordinary tradeoffs in research: some projects require sophisticated equipment.



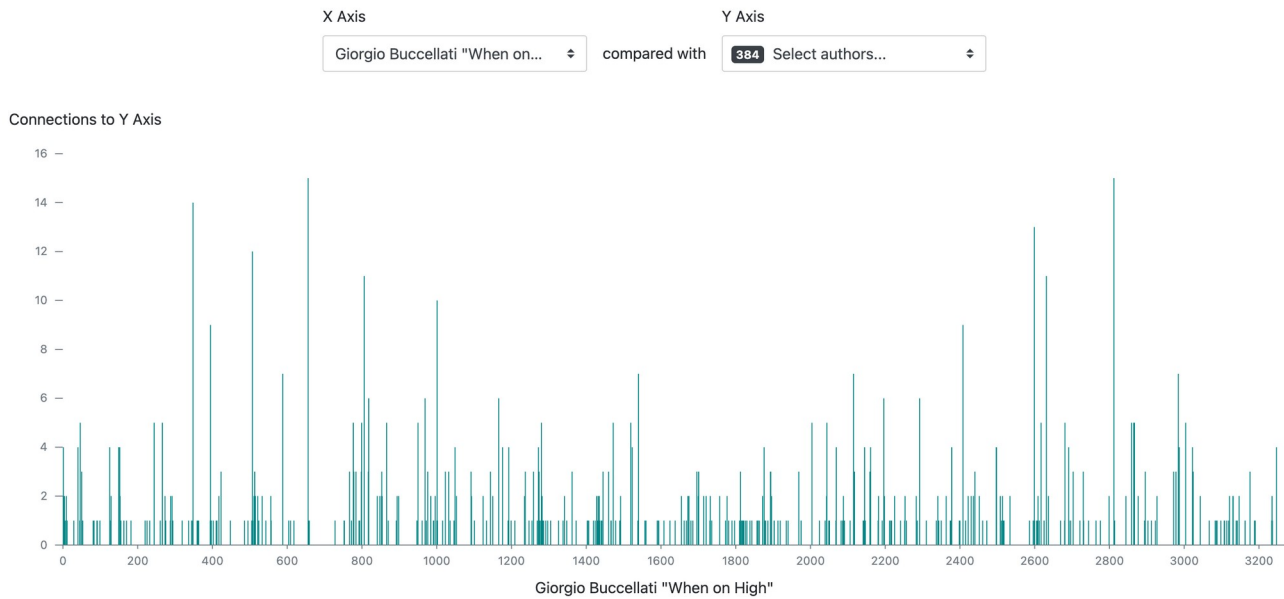
of the works they describe. Still, they do reflect to some degree the depth and specificity of the coverage of the core narrative carried out by the team of scholars.

How can these hundreds of bibliographical summaries be compared with the core narrative? One way would be to extend the histograms used in Chapter 3 to the dimensions of the entire website. I developed a ‘histogram explorer’ visualization software package in order to view and explore a histogram of the comparison between any one text (on the x-axis) with some or all of the other texts in the database on the y-axis. This visualization allows a rapid identification of which areas are more or less connected to the core narrative. By clicking on a sentence, the user can see exactly which other texts are connected to which affirmations. The user can also adjust the strength of connection in order to visualize only those relations above a threshold value.<sup>268</sup>

Below, I present a first view of Buccellati’s entire text, connected to the 384 authors and more than 550 works in the bibliography, using a cutoff threshold of similarity 0.96. The horizontal axis refers to the sentences in *When on High the Heavens...*, numbered sequentially (the book contains about 3200 sentences). The vertical axis refers to the number of connections found between each sentence in Buccellati’s text and all of the bibliography entries. The value of this number depends on the sensitivity threshold: if instead of .96 we had chosen .94, for example, there would be many more connections. This is a crucial parameter which must be chosen carefully to preserve the significant structures of clusters and empty spaces, while not multiplying connections too much to be analyzed. Here is the histogram representing all the connections between bibliography entries and the core narrative:

---

<sup>268</sup>This threshold value is the reason that some visualizations show dozens or hundreds of connections, and others show only a few.

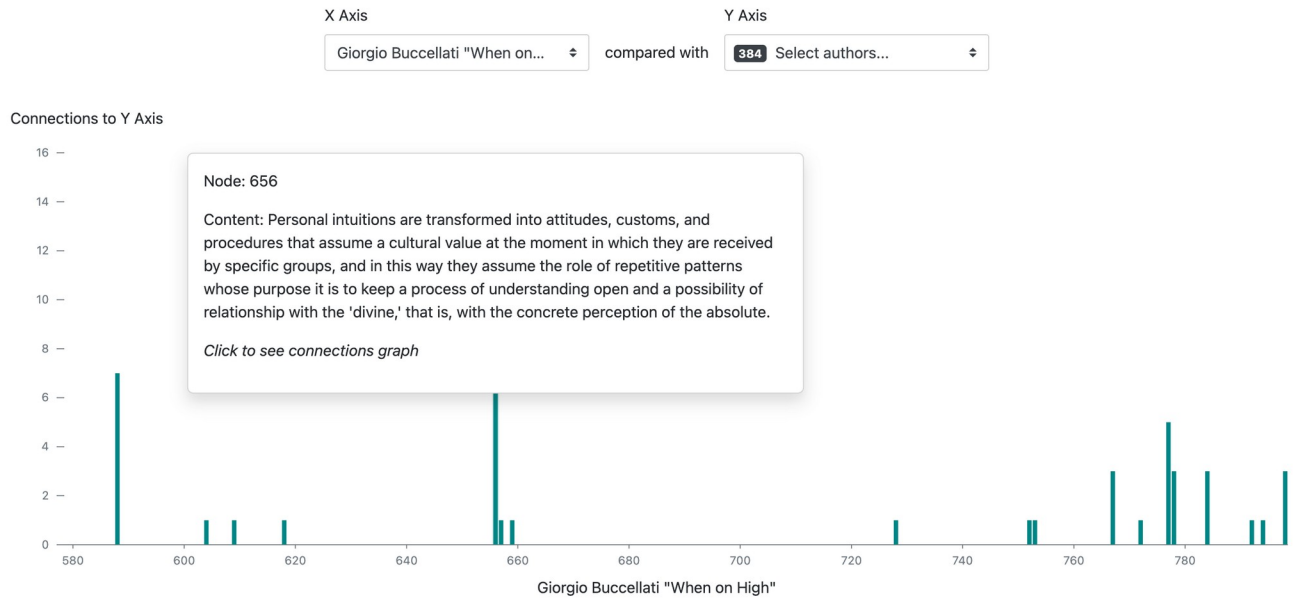


According to the histogram, the general chapters at the beginning of the book, in which the author defines his terms and develops his theory of the structural differences between Mesopotamian and Biblical religion, are relatively more represented in the bibliography than the rest of the book, which deal with specific single issues in religious practice. To interpret the results, we must divide the horizontal axis into chapters. In a future version of the software, I hope to automate the creation of chapter headings on the x-axis, so that it can be more easily understood as a table of contents of the book: each number corresponds to a sentence, and each chapter corresponds to a range of sentences along the axis. For now, we must make do with a table:

Sentence range	Chapter number and heading
1-70	Foreword
71-259	1. Religion and spirituality
260-535	2. Mesopotamia and the Bible
536-544	Part I: The divine element
545-651	3. The concept of the divine
652-729	4. The encounter with the divine
730-817	5. Structure of the divine
818-1019	6. Diachronic developments
1020-1030	Part II: The human element in his relationship with the divine

1031-1247	7. The “affecting presence”
1248-1263	Section One: The Individual A. The divine manifestation at the individual level
1264-1374	8. Morality
1375-1503	9. Divination
1504-1601	10. Prophetism
1602-1755	11. Apparitions
1756-1763	B. The search for the divine at an individual level
1764-1860	12. Meditation
1861-2099	13. Magic and rituals for the individual
2100-2225	14. Individual prayer
2226-2263	15. Materializations
2264-2284	Section Two: The community A. The divine manifestation at the community level
2285-2403	16. Politics
2404-2553	17. Narrative
2554-2615	18. Representations
2616-2684	19. History
2685-2694	B. The search for the divine at the community level
2695-2784	20. The temple
2785- 2860	21. Proclamation
2861-3024	22. Worship
3025-3060	23. The ruin of the sacred
3061-3067	Conclusion: The apex of desire
3068-3130	24. Them and us
3131-3224	25. Fatigue and catharsis
3225-3267	26. Afterword

We note a certain abundance in the bibliography histogram regarding issues like the question of divine kingship, magic, and ritual. Other topics, like morality, divination, and prayer, receive scarcer attention in the texts of the bibliography. As before, we note some holes in this graph. But now we can also zoom in to explore. For instance, between sentences 600-800 there are few connections. We can take a closer look, and see the content of Buccellati’s book at sentence 656 for instance:



The sentence shown above is at the beginning of chapter 4. This chapter is entitled “The Encounter with the Divine” and deals with some of the crucial underpinnings of Buccellati’s argument. Yet, as we learn by noticing the holes in the graph above, the bibliography contains very few works dealing with its issues directly. Although the research team studied this chapter for several weeks and spoke about it for hours, produced more than twenty footnotes to it and wrote thematic essays on some of its central issues, querying the Graph Object reveals that there is a lacuna in our bibliographic treatment. The tool is unable to tell us why this hole exists: perhaps there is a dearth of scholarly studies on the issue? Or perhaps our team was not sensitive enough to this point, notwithstanding its importance in Buccellati’s thought?<sup>269</sup>

Interpretation requires the intuition of the humanist, but the Memex can help direct and sharpen our intelligence by offering a “bird’s eye view” of the work done thus far, and of the work that remains to be done.

<sup>269</sup>Operationalizing the concept of ‘void’ is a great deal more complex than it might seem at first. Achille Varzi dedicated a book to this question. He uses examples to point out the complexity of definitions. For instance, given a stemmed wine glass, the concept of concavity is insufficient to describe the difference between a point within the cup of the wine glass and a point outside the cup but near the stem, which are both within the space defined by the object, but only the former would be considered within the glass. See Roberto Casati and Achille C. Varzi, *Holes and Other Superficialities* (Cambridge, Mass: MIT Press, 1994).

## Conclusion

In chapters 3 and 4 we have seen several examples of the information that can be extracted from our method of processing, storing, and visualizing text data in a Graph Object for a historical study. In the cases examined in this chapter, our Memex and its internal method has shown some promise both in representing the internal connections between the works included in a bibliography, and in comparing the overall bibliographical coverage to a baseline “core narrative”. The method effectively can guide future research toward topics that will even out the overall coverage and improve the quality of our bibliography.

In the next chapter we will face a few more examples that lead directly to the issue of ‘information glut’ to which I have alluded throughout this thesis. What aid can my system offer at the first stage of research, when a stack of too-many articles has to be culled and specified? And what other ways can the structure of information be visualized and interpreted, in order to help the scholar reach greater intelligence of the data?

## Chapter 5 – Further Models of Historiography

The examples we saw in Chapter 3 have introduced a few basic possibilities afforded by querying the Graph Object. We have seen that the vector representation of words permits several forms of statistics to be applied to small- and medium-sized corpora of historical or historiographical texts. The visualization of this information as histograms offers us a way to characterize each text in a corpus by its “signature” as compared with other texts, and to find sections of a text that were closely or sparsely connected to other texts, as a measure of the relative importance others had given to the affirmations contained in the baseline text.

In Chapter 4, we applied this methodology to a larger corpus—so large and varied that even the scholars who worked most intensely on the corpus might be unable to say reliably whether their efforts had covered the topics they set out to document. In this chapter, we will explore several more ways to use the Memex and leverage the vector representation of text in order to handle large amounts of information, and extract insights that are not readily accessible to the reader. The first of these is the final part of the Multinodal Index described in the last chapter.

### **The Memex as “Fisheye Lens”: the Author Comparison Tool**

The final tool I developed as part of the multinodal index on the Mesopotamian religion site is a mapping of keywords across the entire bibliography, in such a way as to permit direct comparison between two authors across the “keyword space” of the entire site. One could think of it as a fisheye lens, which due to its shape can receive light from an entire hemisphere of three-dimensional space and project all the stars in the night sky down to a two-dimensional plane. In a similar way, the Author Comparison tool receives information about all the keywords in the bibliography and maps them to a two-dimensional space, which allows an intuitive comparison between authors based on how the keywords they use are distributed.

The author comparison visualization<sup>270</sup> overlays one or two authors on the entire "keyword space" that represents all authors contained in the bibliography. The question this visualization helps to address is: how well do authors cover the "keyword space" of the entire corpus present on our site? It can give an indication of how completely an author covers the entire space, and also how closely two authors' interests intersect. It is important to note that this representation is based only on the bibliographical material present on the site: therefore, it does not represent the actual content of the author's body of work, but only the subset present on the site, filtered through the lens of the scholar who summarized the work(s) in the bibliography. When the comparative tool indicates an interesting intersection or non-intersection of two authors, one can follow the bibliographical links for the context in which to interpret the results.

One way to compare across a large space with thousands of interconnected keywords, such as we saw in the Multinodal Index, is to show only the nodes, but place them on the two-dimensional screen in a way that retains some of their underlying structure.<sup>271</sup> One way to accomplish this is to place nodes close to the nodes they co-occur with in documents. By highlighting the nodes belonging to two authors in two different colors, the two can be compared.

I used the UMAP algorithm to place the keywords on the plane according to their embedding vectors in the Memex.<sup>272</sup> Recall how our text has been encoded: it has been passed through a neural network language model and converted into a vector that represents its content as a series of numbers. These vectors have various dimensions, according to the complexity of the model they are based on. In the case of the OpenAI word vectors I am using at the time of

---

270 See [4banks.net/Mes-Rel/multinodal/templates/pairing\\_authors.html](http://4banks.net/Mes-Rel/multinodal/templates/pairing_authors.html)

271 How precisely to do this is the object of intense research in data science. One popular method, known as t-SNE, was developed by Laurens van der Maaten and Geoffrey Hinton in 2008. See [distill.pub/2016/misread-tsne/](http://distill.pub/2016/misread-tsne/) for an easy introduction to the method, with interactive visualizations that help see how powerful the procedure is—and how easy it is to misinterpret its results. My project uses a UMAP projection of the keywords.

272 For more information about this recently developed procedure (2018), see the foundational article Leland McInnes, John Healy, and James Melville, "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction" (arXiv, September 17, 2020), <http://arxiv.org/abs/1802.03426>. An easier introduction, with interactive visualizations, written by Google's Andy Coenen and Adam Pearce can be found at [pair-code.github.io/understanding-umap/](https://pair-code.github.io/understanding-umap/). See in particular the three-dimensional image of a mammoth mapped to a two-dimensional representation: it offers an easy way to intuit the nuances of the problem of projecting high-dimensional structures to two or three dimensions, and also a direct comparison of t-SNE vs. UMAP on the same dataset.

writing, the vectors have 1536 dimensions. In this context, ‘dimension’ is a measure of the complexity of the mathematical object, and means that there are 1536 independent parameters that describe the syntactic and semantic content of the word or phrase represented by the vector.

We are familiar with the concept of dimension as regards the three-dimensional space. We could imagine a two-dimensional world, as Edwin Abbott did in his delightful *Flatland*.<sup>273</sup> We may also go in the other direction, and construct arbitrarily high-dimensional spaces. Like a two-dimensional space is constructed by extending a line in a direction orthogonal to it, and a three-dimensional space is constructed by extending a plane in its orthogonal direction, we can arbitrarily continue extending the idea of each successive dimension being orthogonal to all the preceding dimensions. It is usual for mathematicians to treat word vectors as inhabiting a high-dimensional space which can be explored and manipulated using geometrical rules. But humans are incapable of intuitively imagining such a monstrosity as a 1536-dimensional space and the interactions between vectors inside it, so techniques of dimension reduction have been invented to allow a more immediate understanding of the dynamics of the otherwise inaccessible space.

The recently developed procedure called UMAP is a sophisticated approach to this problem. It locates a low-dimensional manifold that preserves some of the high-dimensional structure in a data set, and maps it to a two or a three-dimensional representation (McInnes, Healy, Melville 2018)<sup>274</sup>. In the literature about the UMAP algorithm, there is an example that can help us to clarify its purpose and methodology. A mammoth skeleton at the Smithsonian Institute was 3-D scanned and visualized as a rotating point cloud on a webpage.<sup>275</sup> The researchers also projected the skeleton into a two dimensional static plot, following the UMAP procedure.<sup>276</sup> The two representations clearly demonstrate the variety of ways three-

---

273The first edition text is available at [https://en.wikisource.org/wiki/Flatland\\_\(first\\_edition\)](https://en.wikisource.org/wiki/Flatland_(first_edition)).

274UMAP is an improvement over the basic idea of t-SNE, and serves a similar purpose to the older PCA method.

275For the visualizations mentioned in the text, which give an intuitive understanding of this point, see the excellent tutorial at <https://pair-code.github.io/understanding-umap/>.

276See Coenen and Pierce, “A deeper dive into UMAP theory”, <https://pair-code.github.io/understanding-umap/supplement.html>



dimensional information can be represented on a flat screen through the use of UMAP, as well as the impact of the tuning parameters in the algorithm.

In a sense, the UMAP projection is like the use of perspective in volumetric drawing, which allows a draftsman to represent a volume on a two-dimensional page. But instead of settling for a single point of view, which shows some aspects and hides others, as an artist does when drawing in perspective, the UMAP algorithm attempts to show the information about where points lie and how they are connected to each other in a more complete manner, and for an arbitrarily high number of dimensions, not just the three that can be represented using classical perspective drawing.

Several parameters affect how the points are plotted, and there is always a tradeoff between representing local accuracy and global structure. As a viewer can quickly realize by playing with the parameters on the website, there is no “right” answer regarding the best values, because there is no mapping that does not lose information. There are only mappings that preserve certain kinds of information and lose other kinds.<sup>277</sup> Even so, the viewer can easily intuit that *something* crucial about the original 3-dimensional structure is preserved in the 2-dimensional mapping.<sup>278</sup>

Because of its characteristics, UMAP offers a way to represent the word vectors in our Memex on the plane, and thus form another kind of map to gain insight about the contents of

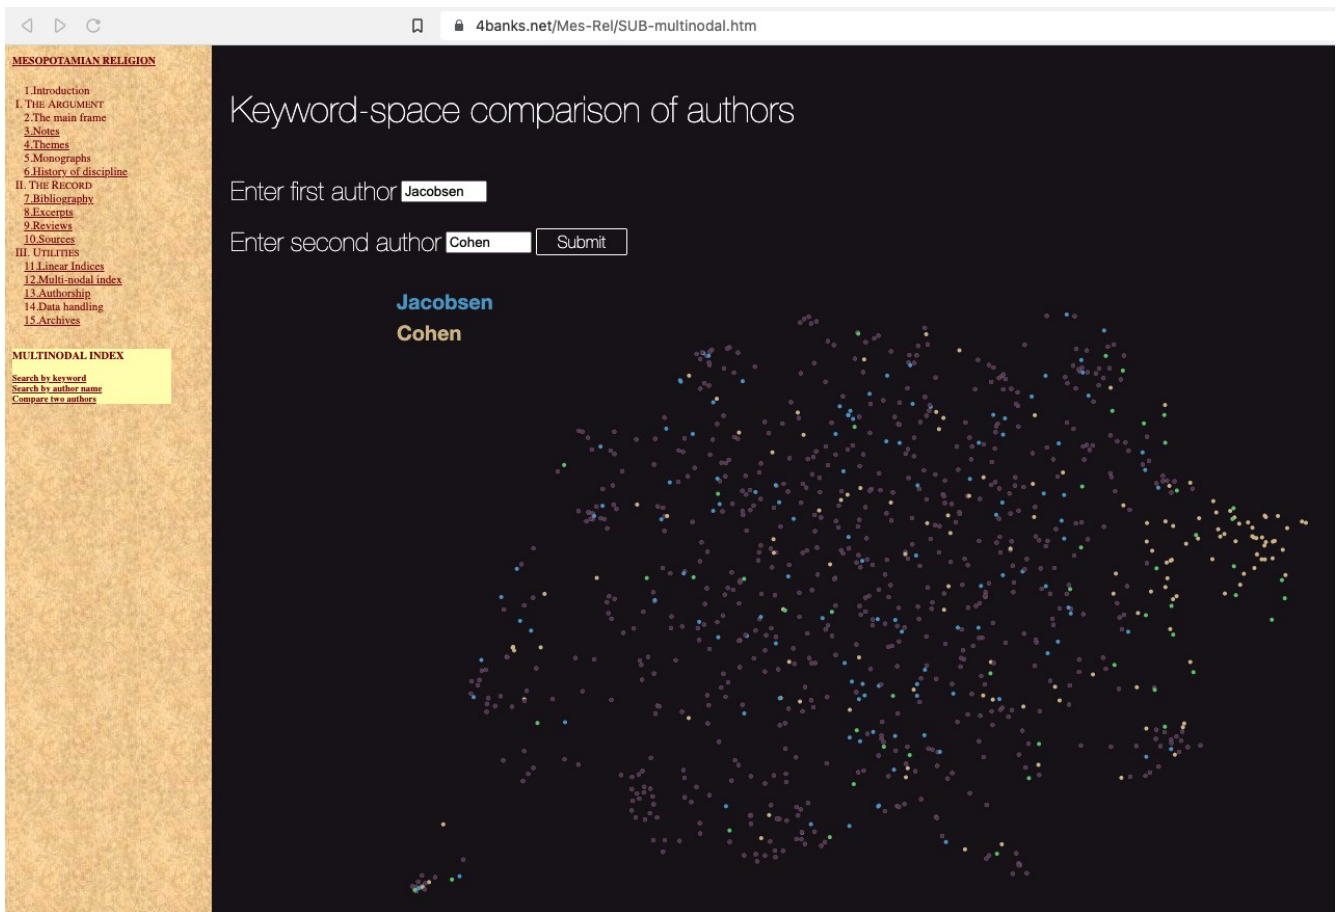
---

277The crucial parameters for a UMAP projection are the minimum distance and the n-neighbours, which determine how well local and global structure are balanced in the low-dimensional projection. In other words, they are a measure of just how fuzzy each datapoint should be considered. The diameter of “fuzziness” of each point is also determined by its context: if it is far from other points, it will have a wider circle of relevance.

278We are well used to mapping three dimensional objects to plane surfaces. The history of painting and cinema is full of solutions to this problem: one of the classical solutions is the invention of perspective. We could think of text publication as a form of dimension reduction from the multidimensional graph of the many references and ideas in the mind of a scholar, to a linear, “one-dimensional” article with a second dimension (footnotes) to suggest the richness that lies beyond the explicit narrative. A written document is like a projection of a richer space, from the hyperuranic space of eternal thoughts connected in a hypergraph of “reality”, to the space given by language and the human linear experience of time. Other ways to project a higher dimensional text space exist too. Pliny is one interesting example from within studies in the Digital Humanities, based on making connections in a 2d space, by manually placing objects close or far, above or below. See Bradley 2008. The Macintosh computer “desktop” was a first instantiation of this idea, based on the older physical manipulation of paper notecards on a desktop. See Bradley and Pasin (2017) at <http://www.digitalhumanities.org/dhq/vol/11/1/000279/000279.html>. See also the research page on the Pliny project at <https://www.kcl.ac.uk/research/pliny-project>

our data. As mentioned at the outset of this thesis, the translation of text information into mathematical vectors unlocks access to several affordances, including this ability to visualize distributions of text information within a space, and notice correspondences and differences that are not visible from the text itself.

As an example, consider a comparison of two prominent authors on the Mes-Rel website, Thorkild Jacobsen and Yoram Cohen. Jacobsen is one of the most illustrious scholars of Mesopotamian history who worked at the University of Chicago a half century ago. He famously took up the challenge posed by Oppenheim—who asserted that a Mesopotamian religion “should not be written”—and penned one of the classic works dealing with Mesopotamian religion. Cohen is a contemporary scholar who focuses on wisdom literature and has the advantage of access to many texts and pieces of information that had not been discovered in Jacobsen’s time. Thus, a comparison between the relative coverage of our website’s semantic space by these two authors might show some interesting features:



In this image, the yellow color indicates the keywords present in Cohen’s work, the blue color indicates the keywords in Jacobsen’s work, and green keywords are those that are shared between the two authors. Purple dots indicate other keywords that are present in the bibliography, but not found in either author in the present comparison. There is significant overlap between the two authors (the green nodes are well-dispersed); and both authors cover the field fairly completely, although Cohen clusters to the right side of the field. We thus have four sets visible simultaneously through the “fish-eye lens” of the Memex: the entire “galaxy” of keywords present in the site (purple); those that are shared by the two authors (green); and the two sets that are exclusive to one or the other author (blue and yellow).

The elements are clickable, returning more information or direct links to the relevant bibliographic entry, the reader can use this map as a way to navigate the bibliography by topic instead of by author name. When the comparative tool indicates an interesting intersection or non-intersection of two authors, one can follow the bibliographical links for the context in

which to interpret the results. It is difficult to adequately represent the usefulness of this tool on a static page, and I encourage my readers to try it for themselves at [4banks.net/Mes-Rel/SUB-multinodal.htm](https://4banks.net/Mes-Rel/SUB-multinodal.htm).

Because of the spatial distribution of words, we might be tempted to try to interpret not only the overall distributions, but also the specific locations of words in the space. However, the specific locations in the space of a UMAP representation do not necessarily have a clear interpretation, because some distances are not preserved during the dimension-reduction procedure.<sup>279</sup> Therefore, it seems to not be possible to directly interpret what it means that Cohen's contributions cluster in the right-hand part of the space. UMAP does offer a way to *compare* two distributions, and the difference between two authors can be significant.

Another important clarification is that each time the UMAP projection is calculated, it returns different patterns. In reducing thousands of dimensions to just two, the algorithm depends on “fuzzy simplicial complexes” and random numbers in its calculation, and returns similar, but not identical results upon repeated iteration. So how could it serve as a map? The crucial point here is that the meaning of clusters and voids in the UMAP plot are not determined by their specific location in the two-dimensional space of the plot, but by the *interrelationships* that obtain between clusters on the plot itself. The details regarding the degree to which distance and density are preserved in the UMAP transformation are complex.<sup>280</sup> For our purposes, they are preserved enough that we can assert that, even though we cannot assign a static meaning to the x and y axes of a UMAP plot, we *can* distinguish clusters of information that have some semantic similarity, and we can understand their meaning through comparison with other clusters of information.

---

279The algorithm seeks to preserve *local* structure, in order to represent clusters of related data points, but the inter-cluster distances do not clearly reflect the *global* distances present in the original data. See N. Oskolkov, “tSNE vs. UMAP: Global Structure. Why Preservation of Global Structure is Important”, <https://towardsdatascience.com/tsne-vs-umap-global-structure-4d8045acba17>

280A good starting place is the discussion on Stack Exchange regarding t-SNE, the predecessor to UMAP: <https://stats.stackexchange.com/questions/263539/clustering-on-the-output-of-t-sne> Note in particular the last answer by username ‘amoeba’, who convincingly argues that the dimensionality reduction procedure does preserve density and clustering. See also Martin Wattenberg, Fernanda Viégas, and Ian Johnson, “How to Use T-SNE Effectively,” *Distill* 1, no. 10 (October 13, 2016): 10.23915/distill.00002, <https://doi.org/10.23915/distill.00002>., available at <https://distill.pub/2016/misread-tsne/> Of course, this discussion relates to t-SNE, not UMAP, which has introduced several improvements. See the interactive explanation at <https://pair-code.github.io/understanding-umap/>

This fact might be described through a “theory of relativity” of the space of the plot. There is no univocal meaning that can be attributed to a location on the plot; all meaning is generated through the relations between elements of the plot, not their location in immutable space. The meaning of the UMAP plots is “inner-referential”.

The point can be more clearly seen through an analogy. In an article entitled “The semiotics of ethnicity”, G. Buccellati points out that “Ethnic identification is based on shared recognition of specific signs. There is no identifier that is in and of itself “ethnic,” the way a sound may be measured acoustically or a color chromatically.”<sup>281</sup> Ethnic signs receive their significance in relation to each other:

The attributes of an ethnic marker are exquisitely “-emic” in nature, meaning that each acquires its distinctive valence from its opposition to another, the way a phoneme is a sound (-etically) charged with a specific valence because of its opposition to another. Just like any sound may emerge as a phoneme in a given contrastive system (a language), so an “ethneme,” as it were, is any cultural trait that assumes contrastive valence as a group identifier. Accordingly, no single sign is ethnic, nor is one more ethnic than the other. Any sign may become ethnic to the extent that, precisely, it signifies symbolically something not immediately coterminous with its typological contours. [...] The linguistic analogy is once again enlightening – for no single phoneme can define the expressive system of a language, but only a system with a subtle and organic web of interrelationships. It is the system as such that proclaims an opposition, a meaningful contrast with what is outside the system. The more complex the system, and the more far-reaching is the contrast, the more defining is the opposition.<sup>282</sup>

Like in the semiotic interpretation of ethnic signs, we can interpret the clusters and voids in our UMAP plot by observing their relations.

## The Memex as “Filmstrip”: Time-dependent Projection of JSTOR Datasets

Since locations in the UMAP space cannot be directly interpreted, a more interesting application of this procedure is the visualization of the development of a field of study over time. If the previous image were to be made in a time-dependent manner, encoding Jacobsen and Cohen’s writings as they were produced year after year, we would see that at the beginning

---

281 Giorgio Buccellati, “The Semiotics of Ethnicity: The Case of Hurrian Urkesh,” in Jeanette C. Fincke (ed.) 2010, *Festschrift für Gernot Wilhelm anlässlich seines 65. Geburtstages am 28. Januar 2010*, Dresden: ISLET, pp. 79-90

282 *Ibid.*, 80.

of their respective careers, all the dots would be purple. Over time, as the authors explored the space of their field, the dots of the representation would light up one after another until the overall picture seen above was fully developed.

We could further represent the development of time as a series of superpositions: by representing the map of keywords at each subsequent year, we would be able to see how an author's interests changed over time. Some keywords would remain illuminated for their entire career; other keywords would be lit for a season, and then the author's interest would move to another section of the map. The series of representations would be like stills from a filmstrip, and by playing them in order we would see some elements remain stable, and some changing over time.

A similar procedure could be applied to an entire domain: for instance, we could encode all the articles from a group of scholarly journals into our Memex. This time, instead of using keywords as a proxy for content, we could use the word vectors that encode phrases, paragraphs, or complete texts. By working in this manner we could encode more information about context within the mathematical representation of the text, and thereby gain greater precision in the interpretation of the meaning of clusters and voids within the plot.

Such a procedure would be most interesting for a set of journals that represent a specific field of research, from a variety of points of view. We could visualize the UMAP projection of the word embeddings as a function of time. This would show how the topics dealt with by those journals had changed over time, which clusters remained relevant, and which were transitory interests.

On the JSTOR website, is possible to request to download an entire set of articles, including their full text, in a format that is easy to use in a program.<sup>283</sup> I was granted access as a doctoral student, defined a dataset comprising all the full-text JSTOR articles from 1950 to 1980 containing the words “mesopotamia” and “religion”.<sup>284</sup> The dataset contains over 1700 articles. They are spread over a wide variety of journals, and include many articles that are not

---

<sup>283</sup>JSTOR offers its download in the jsonl format.

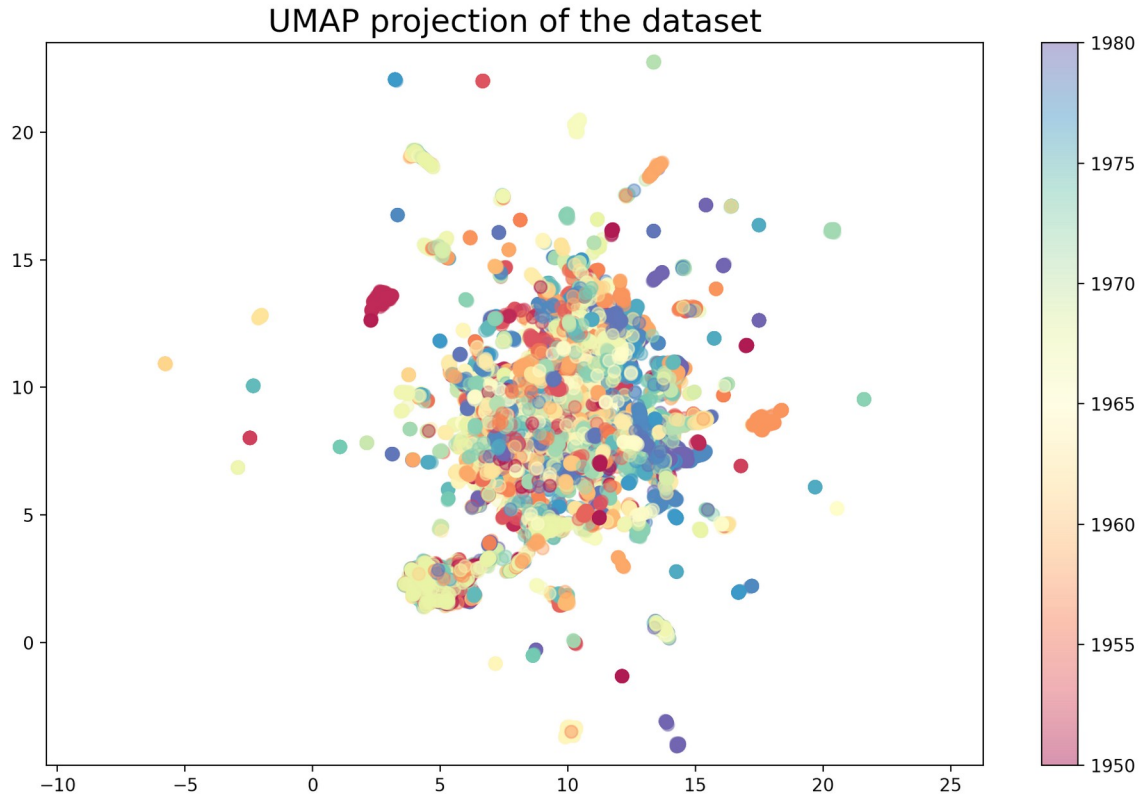
<sup>284</sup>This choice was made in hopes to be able to further contribute to the research project described in the last chapter. Unfortunately, JSTOR contains very few journals relevant to the field, so most of the articles I thus was able to access were not useful. However, it still offers an example of what is possible using my method.

in fact relevant to my topic, in journals that are not historical in nature. This fact seemed at first to be a good test of the usefulness of my system, since keyword searches on databases like JSTOR often return irrelevant articles which must later be manually excluded from a research project. If the automation of the search process can be improved, it would be a step in the direction I have outlined in this thesis. I encoded all the articles in a local database, using the SpaCy word embeddings to generate similarity values connecting the sentences of all articles. This process took some time, since the computer had to calculate millions of relations over the 1700 articles in the database.<sup>285</sup>

By applying the UMAP algorithm to the word embeddings in my Memex, I produced the following graph, which indicates the existence of structure within the data. Relative distance between points can be interpreted as a measure of the similarity between the semantics of each text. Thus, clusters represent groups of similar articles. Furthermore, I represented time as color, on a spectrum from red (1950) to violet (1980). In this way, the movement of scholarly interest in my dataset could be mapped in time:

---

<sup>285</sup>The sheer size of the files involved led to several technical problems. The raw jsonl file containing all the articles is as large as a movie file—700MB—even though it only contains text. At first, my computer (a modern MacBook Pro with 16GB of RAM) ran out of memory as the overall database crashed. After some research, I was able to make adjustments and successfully process the articles. But while I was processing the dataset for the first time, the American startup company Open AI released an updated version of its GPT3 language models, with a simple chat interface called ChatGPT and low cost API access to the word embeddings behind it. This product represents a significantly better form of word embeddings than the SpaCy models I had been using. I therefore wrote a new version of my software to download the word vectors from Open AI and saves it in the database of articles. From this information, I could again calculate similarity scores, and produce the same visualizations we have seen above. On my first attempt, in order to speed up computation, I processed the articles as 5000 character chunks instead of splitting at sentences.



Seeing thousands of articles represented as dots superimposed on a page is still an opaque representation, but it contains several features that promise to be useful. First, we have successfully reduced the dimension of the information from 1536 to 2. And in doing so, we see that the points do not map to an amorphous haze, but to specific regions of the page in identifiable clusters. The UMAP procedure has preserved some of the underlying structure in the ‘thought space’ represented by the articles. Second, colors too are not evenly distributed, but rather cluster in specific locations. This indicates that the interest in the field was concentrated in specific domains at specific times.

I produced a version of this same graph as a video, by looping through the years and watching where scholarly interest concentrated over time.<sup>286</sup> Such a visualization permits an overview of the historical development of the articles: some clusters remain stable over time, indicating topics that are of enduring interest to the scholars in the field. Other clusters appear and disappear, which might have been caused by sensational discoveries that sparked intense

---

<sup>286</sup>The video can be found here <https://youtu.be/DP51HS8zHd8>.



interest in a specific field for a year or two, and then pass out of style. I hypothesized that by examining the clusters that appear and disappear over time, and characterizing their contents through directly examining which articles map to that cluster (see, for instance, the large cluster located at coordinates [5,3], which contains articles from many different years) one could map the development of ideas and the interests in the field. But while tantalizing, both the static visualization and the video version fail to provide much usefulness for the scholar: it would be necessary to be able to directly explore the space and characterize the meaning of a given (x,y) coordinate, as well as view the time-dependent behavior of the clusters with more control and precision than the video version can afford.

Therefore, I built a more sophisticated UMAP viewer as part of my Memex. Using this viewer, large clusters can be examined more closely, revealing the layers of superposition of its texts. We can also map the time-dependent development of the field, by superimposing successive years in different colors and transparencies, and see where scholarly interest has concentrated. Furthermore, we can see exactly which papers give rise to which nodes, and find trails through the data that connect familiar papers with others that are less well known, surfaced by the algorithm.

François Furet writes that “Quantitative and serial history emerge as at once connected with and distinct from each other. But they share an elementary basis in that both substitute the series for the event, both make a construction from historical data in terms of probabilistic analysis. To the classic question “What is a historical fact?” they both give a new answer which transforms the historian's raw material—time.”<sup>287</sup> If time is the raw material of history, a structural approach will have to account for dynamism in time, as I have done by viewing the projection of the data in the Memex as a time-dependent process.<sup>288</sup>

---

287Furet, François. “Quantitative History”, in *Daedalus*, Winter 1971, Vol. 100, No. 1, pp. 151-167.

288See also the pedagogical ramifications of a distinction between “thinking” and “collective activity”, referenced by Kjeldsen and Blomhoj: “Sfard's theory of Thinking as Communicating begins with a claim made by Vygotsky that collectively implemented activities, historically established, are prior to the development of human skills. Thinking is such a skill, and Sfard (2008) defines it accordingly, as the “individualized version of interpersonal communication.” See Tinne Hoff Kjeldsen and Morten Blomheoj, “Beyond motivation: history as a method for learning meta-discursive rules in mathematics”, *Educational Studies in Mathematics* 80(3):1-23.

## Representation, encoding, and intelligence

Interpreting UMAP plots has brought us back to an issue that was raised at the outset of this thesis: “inner-referentiality”. In the first chapter, I discussed a paper by Johanna Drucker which explores several projects that attempted to construct models of the world without reference to external metrics or sources of truth. The most directly relevant of Drucker’s approaches involved attempting to represent time without reference to external measures, but only by constructing a representation from the data contained in works of fiction. Interestingly, the project was only a partial success—possibly because it contains an internal inconsistency that has been noted in artificial intelligence research. The analogy between her partial success and the characteristics of my own “inner-referential” model of historiography is worth exploring in order to better circumscribe my results.

Much research has supposed that by encoding data about the world into symbols that can be manipulated by a machine, a representation of the world is created within the machine itself. This approach may run aground, however, upon the circular logic it seems to imply. “Encodings can only transform, can only encode or recode, representations that already exist,” writes Bickhard (1995, 21). According to Bickhard, in order for an encoding to actually represent an element of the world, a meta-representation is required.

In my own view, the admission that a meta-representation is needed to justify a representation seems to imply a logically inconsistent infinite regress or an “unmoved mover” that is not itself inside the system.<sup>289</sup> Many researchers have assumed that the problem is not as

---

<sup>289</sup>This paradox was brilliantly explored by Douglas Hofstadter in his classic *Gödel, Escher, Bach: The Eternal Golden Braid*. If meta-representation is the only way for representations to actually be formed, emergence is impossible, and there can be no naturalistic explanation for intelligence and consciousness. Bickhard would like to avoid postulating a spiritual “outside”, a ghost in the machine, but finds himself in a conundrum. His resolution is to present emergence in almost tautologically simple terms: “The notion of emergence invoked here is nothing mysterious (though it can be conceptually complex: Bickhard, 1993a; Horgan, 1993; O’Conner, 1994). It simply refers to the fact that some sorts of things once did not exist, and now they do. At some point, they must have come into existence. If something that is of a different sort from what has existed before (even what has existed before locally, though the basic point can be made at the level of the whole universe) comes into existence, then that sort, or an instance of that sort, has emerged. Such a notion applies to molecules, galaxies, solar systems, patterns in self organizing systems, life, consciousness, and representation, among myriads of others. None of them existed at the Big Bang and they all do now. They have all emerged.” (Bickhard 1995, 22).

serious as this, and expect a large quantity of information to converge on a machine representation of experience from which further experiences can be predicted.<sup>290</sup>

Bickhard's research has focused instead for more than three decades (Bickhard 1997, 2005, 2008) on proposing an alternative way to understand representation as "interactivism" that has some of its intellectual antecedents in Gibson's *Ecological Theory of Perception*.<sup>291</sup> Bickhard points out that what he calls "encodingism"—the assumption that representations are constituted as correspondences—hides a critical problem about the nature of representation. The fact that a correspondence exists does not in itself announce the content of the correspondence:

so long as our modeling vocabulary is restricted to such factual correspondences, there is no way to provide (to an agent) knowledge of what the correspondences are with. It is crucial to realize that *knowing that* something is in correspondence and *knowing what* it corresponds to is precisely one version of the general problem of representation we are trying to solve! Thus, as an attempt at explaining representation, encodingism presupposes what it purports to explain.<sup>292</sup>

In my Memex, words are translated into mathematical representations, high dimensional vectors. These vectors can be thought of as "encoding" some degree of the semantic and syntactic content of the original text. Because they are mathematical objects, they can be compared with the variety of tools this thesis has explored. But the comparisons are only "inner-referential", because there is no ground of truth outside the system to which the words or the vectors can be compared. Or is there?

In fact, there *is* an external factor which must be considered, and which is the crucial feature that gives the Memex its value. The language model which is used to generate the word vectors in the first place is this feature. Language models are built by a procedure which from an information theoretical perspective could be likened to compressing the information content of a large amount of text. Therefore, word vectors generated by embedding historical articles

---

290 Doug Lenat's *Cyc* project is one of these. See <https://cyc.com/>

291 Bickhard, M. H. (2017). How to Operationalize a Person. *New Ideas in Psychology*, 44, 2-6. Bickhard, M. H. (2015). The Metaphysics of Emergence. *Kairos*, 12, 7-25. Mark H. Bickhard, "What Could Cognition Be If Not computation...Or Connectionism, or Dynamic Systems?," *Journal of Theoretical and Philosophical Psychology* 35, no. 1 (February 2015): 53–66, <https://doi.org/10.1037/a0038059>. Bickhard, M. H. (2013). The Emergent Ontology of Persons. In Martin, J., Bickhard, M. H. (Eds). *The Psychology of Personhood: Philosophical, Historical, Social-Developmental, and Narrative Perspectives*. (165-180). Cambridge: Cambridge University Press.

292 Bickhard 1995.

within the space of a language model contain information about the embedded text *as it exists in relation to all the text that was used to create the language model*. In themselves, word vectors are not inner-referential, because they are defined through comparison with all the text that generated the model. The overall system *is* inner-referential, in the sense that the entire body of text used to train the model plus the body of text embedded within it are not defined in reference to an external ground truth. Contemporary large language models contain both syntactical and semantical information about the languages contained in the texts used to train them, and by extension they can be used to summarize or encode other texts that were not used in the training set.<sup>293</sup> To return to the comment made about Drucker’s inner-referential model of literary time, perhaps it would have worked if the model had included two levels: one represented by the small group of fictional texts to be modeled, and a higher level containing a large body of text in which time was represented. At any rate, this is what has worked so far with language models.

As concerns the topic of this essay, it is important to clarify in what way the representation of historical information contained in my “Memex” does not incur Bickhard’s criticism. In order to build an automatic system that represents aspects of the world and draws inferences from its representation, it is necessary that the world-representation be a logically coherent mapping of the relations that obtain in the world. Thus the problem of representation must in some way be resolved within the logical system itself.<sup>294</sup> But in our case, these complexities (or contradictions!) are avoided because our aim is explicitly *not* to produce an automatic inference-producing machine.

---

293 See for instance the recent article by Steven Piantadosi criticizing Chomsky’s version of language sticks on the grounds that in fact large language models are able to do many of the tasks the Chomsky assumes are impossible for any except natural language speakers: large language models appear to be theories of language and world. Piantadosi, “Modern Language Models Refute Chomsky’s Approach to Language.” See also Piantadosi’s long list of other scholars (page 12) who concur that encodings do indeed escape Bickhard’s criticism.

294 For instance, the development of visual recognition systems such as those used in a Tesla car’s autopilot requires the construction of a valid representation of the immediate surroundings of the car—to put it mildly, a non-trivial problem to solve. Roboticians are currently exploring the usefulness, and some speak of the necessity, of having a body in order to perceive the world correctly. Yet, at the time of writing, both of these problems are being solved! Tesla has begun selling its “Full Self Driving” system, and Boston Dynamics is rapidly iterating more and more sophisticated robots that can run, jump, and even dance. See <https://www.youtube.com/watch?v=fn3KWM1kuAw>

My Memex retains the activities that a human can do much better than a machine squarely on the side of human action, whereas the activities that a human has more difficulty doing are to a certain extent outsourced to the machine. In particular, the Memex does not attempt to construct an independent model of language. It does not even attempt to account for negation within the automated part of the system, let alone more sophisticated nuances. Nor does it automatically create logical links between statements.

Instead, it creates a model of the interaction a particular group of texts has with a broader corpus. The choice of which texts to include is made by humans; and the logical relations between texts are encoded by humans. The machine remembers these connections, and is able to suggest other possibly relevant texts for review. But it does not presume to be intelligent, nor does it even claim to be without logical contradictions. It is just an efficient clerk.

These facts mean that the mystery of language and human inference remain unsolved. Or, to put it positively, they are preserved. My central axiom is that these activities of the human scholar are central to the project of humanities scholarship, and cannot be outsourced to a machine at this time. But at present it is not crucial to our project to unravel the linguistic complexities of the relation between words and sentences; it is enough that we can recall relevant phrases and evaluate them in the possibly unreflecting but certainly sophisticated way we have always done. Nor is it crucial that we build a logically consistent representation system such as Bickhard champions, for we are not expecting the representation in our database to be as complete as the one in our heads (which also is incomplete!).

The correspondences between symbolic elements in the graph, as we saw in the last two chapters, do not give the logical foundations for abduction. They are much simpler, less pretentious tools that aim to identify clusters and voids. They serve as a heuristic to reduce the amount of information the human scholar has to process, while reserving full judgment about the meaning of the details and full responsibility for creating a narrative to the scholar. It is the scholar who must establish the correspondences between elements of data and their narrative

meaning. To use Bickhard's computer science vocabulary, there is a necessity for a "provider of representation contents for the constitution of encodings" (Bickhard 1995, 17).<sup>295</sup>

Bickhard's "interactionist" approach to representation is "concerned with functionally realizable knowledge of the potentialities for action in, and interaction with, the world. Interactive representations do not represent what they are in factual correspondence within the world, but rather they represent potentialities of interaction between the agent and the world." (Bickhard, 1995). So, perhaps for a robot that interacts with the outside world, the problem of representation and meta-representation is solved by reference to experiences afforded by the body of the robot, its physical interaction with a world larger than its own representation. Its interaction offers the possibility for its models of reality to be tested against a meta-model, and improved.

In my method and in the software that instantiates it, the same holds true: the map contained in my Memex is not necessarily in factual correspondence with the world. It represents possibilities that the scholar must evaluate, include, or discard in his narrative about the world. Our maps truly are not the territory: the mapmaker is a crucial, irreplaceable part of the interpretation of the map. Interpreting the map is the activity that is most specifically

---

<sup>295</sup>The "provider of representation", in the domain of mathematical world-models, is often the person who defines the utility function used to process data and create a model. Bickhard points this out: "Reinforcement learning, for example, requires (among other things) a built-in utility function on the inputs — the system has to already know what inputs to seek, and usually must also have some loss function that is defined on errors. Such an approach can be practically useful in certain circumstances, but, as a general approach, it requires that critical and potentially complex knowledge be already built into the system before it can learn. That is, it requires already existing knowledge in order to learn knowledge. It is crucial to realize that this approach does not use these types of built-in knowledge just for convenience. Rather, the built-in knowledge is essential for the later learning of the system, and the model offers no account of how the initial knowledge can be learned. As a general approach, this immediately yields a vicious infinite regress — the regress of impossible emergence." (Bickhard 1995, 44). It is worth noting that since 1995, when Bickhard wrote those lines, so much has changed in the quantity and quality of computation and the practical realization of neural networks that his statement no longer rings true. In 1995, a self-driving car was pure fantasy. In 2023, it is a product that is available for purchase in some markets. This technology depends entirely on neural network reinforcement learning models of vision and reality. We may not be able to philosophically understand exactly how "true" the models are within the machine itself, but to some extent their truth can be evinced by their abilities. Driving ability, at least, can emerge from a model that has little prior knowledge of how the world works. I do not say "no prior knowledge", because these systems are probably fine-tuned to contain concepts of priority ("don't hit the bicyclist") and concepts that encode the laws of physics independently of the reinforcement learning process. However, recent statements by Tesla indicate that indeed the self-driving capability is entirely done with neural networks and no hard-coded rule following. Granted, the ability to drive a car is a great deal less "intelligent" than the ability to interpret history—and insofar as we are concerned with things like driving a car, Bickhard agrees.

human; remembering the map, computing its detail, and recalling it at opportune times, may also be the work of a mechanical clerk.

We have passed close to philosophical issues that are too large to be addressed here, in particular the question of origins: how did language emerge in the first place, if its complexity defies an inner-referential explanation? But like prudent Dedalus, we pass on by.

Our concern is the question of intelligence, specifically: by what method can the too-large amount of information contained in history books and articles be taken into account by a scholar wishing to produce a model of the world that is as accurate as possible? By what method can the scholar reach intelligence of the world? Howard Rheingold wrote that

There is simply too much information in the world to solve problems by checking every possible solution. The difference between brute-force calculation and human knowledge is the missing link (and holy grail) of hard-core AI research. [...] What a machine needs to know, practically before it can get started, is that the mysterious something that human chessmasters know that enables them to *rule out* all but a few possibilities when they look at a chessboard (or hear a chess situation described to them verbally). When a human contemplates a chess position, that person's brain accomplishes an information processing task of cosmic complexity. The human brain has obviously found a way to bypass the rules of exhaustive search — a way to beat the numbers involved in searching problem space. This is the vitally important trick that seems to have eluded artificial intelligence program designers from the beginning. [...] The point of expert-system building is not to outdo the brain but to *help* human reasoning by creating an intelligent buffer between brain processes and the complexities of the world — especially information-related complexities. A problem-pruning tool could be an important component of such an informational intermediary.<sup>296</sup>

Often, when scholars speak of the possibilities afforded by artificial forms of intelligence, the mind jumps to the dream of an independent agent in the machine, the veritable creation of life, reminiscent of Frankenstein and Golum dreams, the homunculus dreams of power-drunk madness. Sometimes such nightmares seem like they are on the threshold of becoming reality, and much of the AI discourse in the press and in popular literature revolves around the debate between enthusiasts for such an outcome and those who are worried about a future in which computers are independent agents, possibly more intelligent and powerful than humans.<sup>297</sup>

But together with these disquieting dreams, there is an older, simpler, humbler understanding of the relationship between humans and computers, in which the computer is

---

296 Howard Rheingold, *Tools for Thought*. MIT Press (2000), chapter 13.

<https://www.rheingold.com/texts/tft/13.html#Chap13>

297 A deeply moving rendition of this outcome is portrayed in Spike Jones' movie *Her* (2013).

seen as a tool, an extension of the mind, an “intellectual prosthetic” that has the explicit goal of co-production: cooperation, empowerment, augmentation, not replacement.

As an aside, it is worth noticing and thinking about the fact that humans seem to enjoy thinking about augmentation more than actually carrying it out. As Rheingold put it,

The "rule of two" (that computer power would double every two years) and the Engelbart-induced zeal of the augmentation team kept them fueled for an effort to bootstrap and continually adjust themselves to the capabilities of their upgraded tools — an effort that required extraordinary intensity. The bootstrapping and readjusting continued with unabated enthusiasm, at least until the early 1970s, when the idea of building a system that was meant to "transcend itself every six to eight months" to keep pace with hardware and software advances turned out to be more pleasant to contemplate than to carry out. It had been a challenging and exhilarating to build this new system for augmenting thought — but it wasn't as much fun having one's work habits augmented at a forced-march pace.<sup>298</sup>

The pace is not the point: it is not necessarily in order to work faster, that augmentation is important. The point is to take more information into useful account, and to this end it is necessary to augment human capabilities, as has already been done many times throughout history. The invention of language, writing, archives and libraries, printing presses, and electronic publication are all steps along the way to the information overload in which we now find ourselves.

The speed at which processing of data takes place is a less important feature in the academic setting, where many researchers have the luxury of guaranteed stipends and freedom to contemplate their topics in depth. What is essential in the academic setting, and which is sometimes pushed aside in the more competitive and fast-moving commercial sector, is attention to the entire problem at hand, and the interpretation of all the available information in order to come to a reasonable and defensible interpretation.

## **The Memex as “Microscope”: Visualizing the Atoms of Discourse**

A final metaphor to describe the Memex and its possibilities is as a microscope. One of the most obvious advantages of machines is that they can remember a virtually infinite amount of information and retrieve it on demand. This fact permits the Memex to present not only large-

---

298Rheingold, *Tools for Thought*, MIT Press (2000), chapter 11.  
<https://www.rheingold.com/texts/tft/11.html#Chap11>



scale dimension reduction visualizations that aid the comparison of bodies of text, but also to recall the granular details of the points where discourses intersect, agree, and disagree.

Historical discourse happens through debates carried out in a variety of settings: conference presentations, articles, and longer volumes are deployed as interventions in a sustained argument between alternative explanations of the facts. Some explanations emphasize cultural aspects; others derive their strength from archaeological artifacts; still others are instantiations of overarching theories about humans that find partial justification within historical events. To some degree, the participants in these activities are resistant to alternative explanations, which is another way of saying that they believe they have found the most convincing explanation. But well-intentioned historians will be interested to know if there is data that might convincingly be used to critique their theories and explanations, and move their position closer to an intelligent reading of the available facts.

In order to do so, it might be useful to employ a symbolic language to carry out some of the debate. Alan Kay noted that

The particular structure of a symbolic language is important because it provides a context in which some concepts are easier to think about than others. For example, mathematical notation first arose to abbreviate concepts that could be expressed only as ungainly circumlocutions in natural language. Gradually it was realized that the form of an expression and manipulation could be of a great help in the conception and manipulation of the meaning for which the expression stood. . . .The computer created new needs for language by inverting the traditional process of scientific investigation. It made new universes available that could be shaped by theories to produce simulated phenomena.<sup>299</sup>

This is a useful metaphor for my project. What symbolic language do historians think in? How could a symbolic language be developed that would help them to think more easily about the concepts that are most important in their work? One attempt at creating such a symbolic language is already present in the Memex, in particular in its ability to encode logical relations between affirmations. Because the Memex can be used by more than one scholar, it is possible (and likely) that these logical relations not be internally consistent, but rather express the multiplicity of points of view within the community of scholars.

In order to understand how the Memex could be used in this way, and what advantages its use might have, we can again refer to the architect Dedalus, who might serve as a guide


---

<sup>299</sup>Cited in Howard Rheingold, *Tools for Thought*, MIT Press (2000), chapter 11.  
<https://www.rheingold.com/texts/tft/11.html#Chap11>

through the labyrinth of practically infinite Graph Objects, webs of connections and relations, forests and subterranean networks of knowledge. If we could only find the architect of the network, we could successfully navigate its complexity. But there is no Dedalus, and the labyrinth is being designed and constructed as we speak! How could we successfully notice the overall structure of the network, as well as traverse each branching of the path correctly?

In essence, we need a way to develop a perception of the structure as a whole, even while that whole is a work in progress. As I have already recalled, my experience working with the Mes-Rel team made me wish for ways to record and recall the detailed richness of our conversations and the ‘conflict of interpretations’. My Memex is an attempt to build a map that helps us perceive where we are and decide where to go, with the ability to update the map as new information and relations are added.


As a simple example of how such a map could be created and used, consider a linear discourse such as a short article, represented using the method I outlined in Chapter 3 as a series of sentences connected by PRECEDES relations:



*Figure 5: A linear discourse*

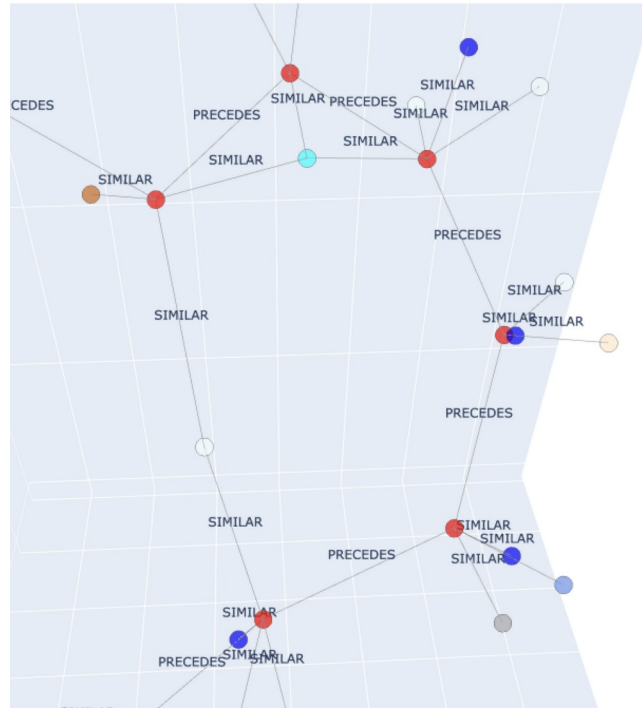
Each sentence follows the last; the discourse proceeds from start to finish. It can be experienced by the reader much in the same way it is represented, in a linear fashion. Now consider the same discourse, enriched by a few footnotes and references to other sources and texts:

Now the linear discourse in red is in relation to other discourses. Perhaps the texts referenced with other colored nodes are primary and secondary sources, as well as more tangentially related texts. Let us also note that the footnotes are points of contact to other linear discourses, as well as to “fragments” of information: if the blue dots were extended with their



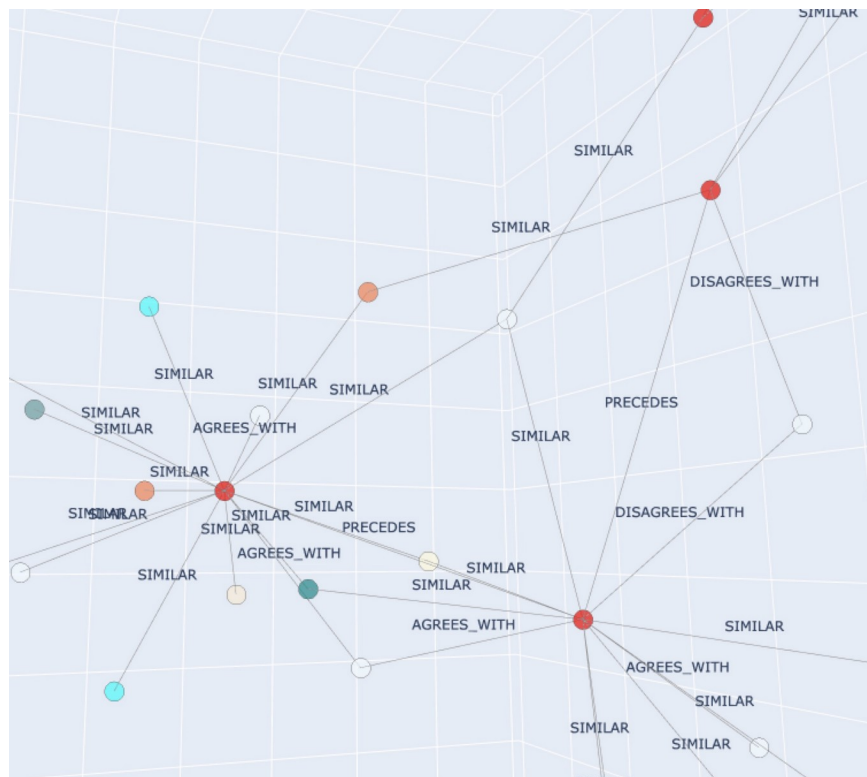
own PRECEDES relations, we would see that the point referred to by the red text is part of a larger discourse, with its own context, and presumably also with its own external points of reference.

We experience our movement through the labyrinth as a linear succession, but in their own essence, the ideas we encounter are interconnected in a multilinear way. A map helps us to perceive our sequential experience in its multiplanar context. Here is the map once more, enriched with a similarity relations as well as footnote references:



PRECEDES relations are automatically produced by the Memex at the time it ingests a new text and saves it to the database. SIMILAR relations are also automatically produced, by comparing the vector representations of each text. These two representations permit the Memex representation to show layers of interconnectivity: the linear text in red is now revealed, through the microscope of the Memex, to contain many hidden connections to other texts and ideas.

But this is not all: the representation can also contain logical connections, encoded as AGREES\_WITH and DISAGREES\_WITH relations encoded by the user. The Graph Object becomes more and more complex:



Logical relations are not automatically generated in my system, but they can be encoded manually. These relations could also be weighted according to the relative authoritativeness of the person who coded the relation, as a way to encode a fuzzy representation of trustworthiness within the relations, and within the topology of the network. Alternatively, logical relations could be aggregated to give a probability of agreement or disagreement between given statements. Or the logical layer could be retained in all its detail to serve as the memory of multiple points of view, such as the variety of points of view I encountered within the Mes-Rel research team.

I am not the first to think of such a machine. In his history of machines for thinking, *Tools for Thought*, Rheingold recalls the first attempts to create a learning mechanism within an expert system. He refers to Avron Barr, who acutely noted the limitations of human intelligence and pointed to machine augmentation as an important path toward greater human flourishing. Rheingold writes that

The key to taking advantage of these natural disagreements between experts, Barr realized, was to build in a mechanism for "remembering experiences," for keeping around old decisions, even if they were wrong, and

creating new rules from the outcome of disagreements. Taken far enough, this aspect of the system leads directly to one of the hottest issues in AI research — the question of whether programs can learn from experience. Barr was only interested in one specific aspect of this issue — the possibility of creating a means of tracking decisions and keeping track of instances where human experts disagree with each other.”<sup>300</sup>

By studying the places where two scholars differ, interesting information can be found. Schools of thought develop, disagree, and separate. New discoveries unite previously separate fields, and create new voids to be explored. Improving the state of the art is a more interesting goal than replacing one philosophy with another, or defending an older version of the facts against the incursions of new-fangled epistemologies. My aim, and the aim of the Memex, is to aid the researcher to navigate the sea of data efficiently and intelligently.<sup>301</sup>

In summary, graph representation emphasizes the fragments and their relations. It is like a microscope<sup>302</sup>, allowing us to see finely dimensioned details in the logical structure of discourse. It is also like a personal search engine, concentrated on the information that has been input by a researcher or team. Books and articles are written as linear successions of words, but in their essence, the ideas the reader encounters are interconnected. A map of the texts and their relations allows us to perceive our sequential experience in its multiplanar context and develop a perception of the structure as a whole, even when that structure is still a work in progress.

The Memex aims at creating such a map. The method I have already developed could support collaboration (such as in the Mes-Rel team) by permitting the detailed memory of a work in progress. Further, analysis of the geometry of the network may permits further insight: what information is abundant? What is missing? How are core ideas connected? Which narratives most completely account for the information?

## Simple Topological Features

It is possible to further specify the results by searching for topological features: for instance, “triangles”, where Buccellati and two other scholars are connected by relations that

---

300Howard Rheingold, *Tools for Thought*, MIT Press (2000), chapter 13.

<https://www.rheingold.com/texts/tft/13.html#Chap13>.

301The evolution of augmented systems depends in large part on the feedback between machines and their human “collaborators”. See 1.08.00 of <https://youtu.be/sG3PWet8fDk> for Englebart’s description in 1986.

302Clarify micronarrative

indicate agreement between all three. We can conclude that that affirmation has a particular importance. We can also consider “rhombi”, where four authors mutually agree, and so forth. A further refinement is to search for triangles or rhombi which have a connection strength above a threshold value.

While the search for “rhombi” may seem like a peregrine exploration in the context of historical research, recent work in other fields bolsters the hypothesis that the graph structure of relations between data points can reveal important characteristics of the data itself. For instance, a team of researchers recently looked for tetrahedra in the distribution of galaxies, and studied the relative proportion of shapes exhibiting a more strongly left- or right-handed structure.<sup>303</sup> If the galaxies are simply randomly distributed in the sky, the proportion of both shapes should be roughly the same. If instead some unknown dynamic affects the distribution of galaxies, one or the other should prevail. In fact, these researchers found a difference. This discovery points toward further research that might explain why the discrepancy was found; at the moment, the fact that there is a discrepancy indicates that there are more forces at play than have been identified so far.

In our case, we can move beyond the simple examples above. A more complex pattern will result in even tighter groups of authors agreeing on a topic. Instead of the ring a-b-c-d-a, we can also search for tightly connected rhombi such that the preceding holds, and also a-c and b-d are connected, as follows:



---

303See Oliver H. E. Philcox, “Probing Parity-Violation with the Four-Point Correlation Function of BOSS Galaxies,” preprint 2022, <https://doi.org/10.48550/ARXIV.2206.04227>.

With this pattern, we can infer that four different authors all have produced phrases that are similar. The number of such phrases is much smaller than before. Considering Buccellati's entire book on Mesopotamian religion, and all the documents available on the website at present, there are only five phrases in Buccellati that present this feature.<sup>304</sup>

The resulting five phrases by Buccellati are:

- 1 "The elements are therefore these different cultural articulations: on the one hand, the binomial fate/the gods in Mesopotamia, and on the other the identity between fate and god in the Bible: it is in this context that it becomes important to clarify my conception of polytheism as compared with monotheism."
- 2 "The history of religions presents us with a very wide range of typologies in which the various methods of conceiving the absolute and the relationship between the absolute and the tangible world are configured."
- 3 "There are innovations that mark particular historical moments, and that in some cases are then abandoned (as in the case of sacred marriage), while others produce profound changes in form and substance that are then acquired (as for the figure of the personal god or the themes developed by the wisdom tradition)."
- 4 "It is significant, in this regard, that Mesopotamian divination in no way refers to the great myths of creation, while for biblical prophets, creation is a central model for all subsequent interactions between God and his people, as for example in the various covenants that emerge as new forms of creation in founding new historical realities."
- 5 "The brief collection I give below (App. 12) is divided into themes."

The first four in the list are all significant phrases, relevant to the overall topic treated. It is also interesting that they deal with the overarching theme, as we might hope, and not with the smaller details of the internal chapters. This fact seems to indicate that our bibliography does in fact contain a variety of sources that support our core narrative. The fifth element in the list is not a significant phrase, as we also expect: our language model is only about 87% accurate.

---

304The query used to generate this list was: match a=(n)-[r1:SIMILAR]-(m)-[r2:SIMILAR]-(o)-[r3:SIMILAR]-(p)-[r4:SIMILAR]-(n) match b=(m)-[:SIMILAR]-(p) match c=(n)-[:SIMILAR]-(o) where n.author = "Giorgio Buccellati" and m.author <>"Giorgio Buccellati" and o.author <>"Giorgio Buccellati" and p.author <>"Giorgio Buccellati" return distinct nodes(a), nodes(b), nodes(c)



Another way to explore the crucial nodes of the text is to entrust the selection to Google’s PageRank algorithm.<sup>305</sup> By selecting the nodes written by Buccellati, and with a pageRank score over 3 (which indicates a relatively high “importance” within the network), we obtain the following table:

348	"There are innovations that mark particular historical moments, and that in some cases are then abandoned (as in the case of sacred marriage), while others produce profound changes in form and substance that are then acquired (as for the figure of the personal god or the themes developed by the wisdom tradition)."
507	"By biblical spirituality, therefore, I mean a system of apprehension of the divine recorded in a coherent and long-term way in a written canon. From a historical point of view, it is important to ask to whom the spirituality that was incarnate in religious institutions was addressed in these two cultural spheres of Mesopotamia and the biblical world, respectively."
1540	"It is significant, in this regard, that Mesopotamian divination in no way refers to the great myths of creation, while for biblical prophets, creation is a central model for all subsequent interactions between God and his people, as for example in the various covenants that emerge as new forms of creation in founding new historical realities."
2440	"The articulation of myths takes place in large compositions of high literary value, and the study of these texts in their complexity of form and content will take place elsewhere, in the second volume of this series (see below, 26.3), dedicated to Mesopotamian literature."
2599	"An extension of this activity involving not statues, but human beings, occurs in a short period of time (between the end of the third and the beginning of the second millennium) in southern Mesopotamia, when the participants in the ritual of sexual intercourse are not two statues, but two living beings, the king and a priestess: this is the so-called sacred marriage or hierós gamos, which is known to us mainly from poetic texts that describe the event with tones of great emotion and lyricism."
2631	"The list of kings suggests a specific sequence that covers many centuries and that, starting from a

---

305Google came to dominate the search engine market thanks to Sergey Brin and Larry Page’s intuition of PageRank, a centrality measure that considers the links between web pages as “votes” of the usefulness of a given page. The human decision to link to a page serves as a signal that the linked page is useful. Many histories of Google’s development exist. One good starting place is Nick Statt’s 2019 article in *The Verge* which outlines the rise and disappearance of Brin and Page. <https://www.theverge.com/2019/12/4/20994361/google-alphabet-larry-page-sergey-brin-sundar-pichai-co-founders-ceo-timeline> Accessed on 19 March 2022.

well-defined beginning (the descent of the institution from the heavens), leads us to an equally well-defined moment (the political situation at the time of writing) with the clear intention of showing a development in continuity."
---

Sacred marriage, which appears twice in this list, is not a central theme in Buccellati's book, although he does dedicate a chapter subheading to the topic. It is, however, a popular topic in much of the secondary literature which was included in the bibliography. The pageRank algorithm successfully locates nodes in Buccellati's text that are highly connected to other texts by other authors.

## Overview

At this point in our exploration, we may overview the many branches of science that we have encountered, draw our conclusions, and chart a course forward.

We began with the observation that the sheer quantity of information precludes exhaustive knowledge of any but the tiniest fields of study today. This is no less true for the historian than it is for other sciences, and perhaps it is even more important for a historian, because our study engages with the totality of human expression, and seeks to describe events and their causes and meanings. Artificially constricting our gaze to examine only a relatively few elements will lead to distortion; but how can perception, memory, and analytic ability be expanded beyond what it currently offers?

As we have seen, one solution to this problem is to engage technology to help process and organize information in the form of 'models'. Digital tools such as databases and search engines can assist in sorting through and finding relevant material, and machine learning algorithms can help identify patterns and connections that might be missed by human analysis. However, it is important to recognize that these tools have their own limitations and biases, and they should be used in conjunction with human expertise and critical thinking.

Another approach, which we have not addressed in detail, is to cultivate interdisciplinary knowledge and collaboration. By working with scholars and experts in related fields, the historian can gain a more comprehensive understanding of the contexts and implications of the

events they are studying. This can also help to challenge and expand the historian's own perspective and assumptions.

After posing this question, I surveyed a few ways that contemporary computational methods could aid historical research. Then I settled on a specific method and built a working version of it. At the end of the first iteration, it became clear that better interfaces were essential, so I built an improved version. Upon applying this method to successively large datasets, with successive improvements in the data processing infrastructure and in the visualizations, we arrived at the current point. At each turn in this story, different technologies have been implemented, and further improvement could be had by greater attention to those methods and techniques.

But in the context of this thesis, I have found it necessary to move on after surveying a representative area of the relevant technologies: the tech world moves very rapidly, and methods that were state-of-the-art when I began writing have been updated and rendered obsolete over the course of this project. In some cases, especially with the word embeddings recently made available by OpenAI<sup>306</sup>, I have preferred to rewrite my software to take advantage of new advancements rather than stick with my earlier versions. In many other cases, it seemed more reasonable to arrive at the bend in the road, observe the surroundings and perhaps opine about what shortcuts might lie off the broader road I end up taking—but then move on and aim to arrive at least at a glimpse of the destination I set out to reach, rather than losing my way in the infinite thickets of incrementally improving computer systems. Here, then, is an outline of our itinerary, together with some suggestions for further improvement and refinement.

**1. The source data we are attempting to consider is *text*.** This implies a series of problems that are not easy to solve. Much of historical data is printed in paper books, which must be digitized in order to be considered in a digital humanities project. This procedure is expensive and has not been undertaken for many, probably most, books on Mesopotamian history—my subject. Even where pdf versions exist, it is no simple matter to extract pdf text to

---

<sup>306</sup>The infuriatingly fast pace of technical development means that in December of 2023, the models I updated a year ago are already obsolete, and no longer available through OpenAI!

digital text with any degree of accuracy.<sup>307</sup> Worse still, articles on Mesopotamian history are replete with foreign languages, exotic character sets, specialized formatting to show poetic structures, and the like. Finally, even where digitized versions of relatively large corpora do exist (for example on JSTOR), the quality of the text is often poor, full of typographical errors and the results of automatic parsing.

All these difficulties compound, and mean that translating texts about history into vectors has to deal with a first, sometimes crucial, source of confusion. Further research into the methods considered in this thesis should concentrate significant attention to improving the accuracy of automated text extraction.

**2. Once the text data is available in a digital format, it must be parsed into “chunks” that can be handled by the vectorizer. How big should these chunks be?** Should they correspond to the entire document, and thus permit a web of connections to be built that preserves the document as its most basic “atom”? Or would a paragraph be a more logical unit in which to present a single idea? Or rather, perhaps the sentence is the best “atom” which can express a statement that other statements can agree or disagree with? In my case, I used sentences when it was feasible to do so, and crude 5000-character chunks of text where the number of nodes otherwise became computationally prohibitive.

Along with the philosophical and linguistic problem about what constitutes the atoms of a text, there are several other issues that bear on our question. One is simply technical: automatically breaking a text into sentences or paragraphs is not as simple as it may seem, since many sentences contain periods (like after “Dr.” or “e.g.” and it is not always easy to algorithmically decide where sentences end. Solving this problem ended up being more difficult than I expected, and even using state of the art language models I notice fairly frequent errors in sentence parsing. This is particularly true about the idiosyncratic style of some historians, as well as a result of the copious footnotes and tangential reflections that are often found in scholarly literature.

---

<sup>307</sup>The Adobe PDF format does an excellent job of preserving typesetting for human readers, but it does so by breaking apart a page of text into single characters that are located in the space of the page. Reverse-engineering this code to reproduce the original text is possible, but it is extremely error-prone when there are multiple columns of text, extensive footnotes, tables, and images.

Another issue is the variability of language and writing styles. Different authors, time periods, and genres of writing can all affect the structure and content of a text, making it difficult to apply a uniform approach to text analysis. This can be especially challenging for historical texts, which may use language and conventions that are no longer in common use. Furthermore, historical texts often contain a wide range of information, from factual details to personal opinions and interpretations. Disentangling these different types of information and determining their relevance to the research question can be a complex task.

Next, there is the question of context. Historical texts are often embedded in a specific cultural, political, and social context that can influence their meaning and significance. Understanding this context is crucial for accurately interpreting the text, but it can also be difficult to fully capture and account for in the analysis. The context of each affirmation changes its meaning, and this is more true the smaller the unit. Contemporary Large Language Models owe a great deal of their successes to the fact that they take context into account. But the reverse is not true (yet) in my implementation: the chunks of text I process, of whatever size, are not processed as part of their overall document. It is up to the reader to recover the document context by reading the nodes that are connected by PRECEDES relations.

Furthermore, there is an inverse relationship between the size of nodes and the length of time it takes to process the vectorization. For the 1700-article JSTOR dataset outlined above, it took 8 days of 24-hour a day processing to compute the relations between the nodes, and I was using clumsy 5000-character chunks in order to speed things up. Had I been using the single sentences I used in the smaller trial runs above, it would have taken many months to process on my relatively powerful MacBook Pro.

Of course, compute power can be purchased relatively cheaply. I could rent an AWS server or a Linode cluster and have access to much more power than my laptop. But doing this would require me to learn yet another system, with yet another series of technical difficulties, and I decided that it was beyond the scope of this thesis to do so: time was running out.

I expect that in the next iteration of my work it will be possible to expand my scope with the use of more powerful computers. There are many other ways to improve performance too,

through the use of better Python libraries<sup>308</sup>, through C code using “just in time compilation”, through attention to bottlenecks, and the careful writing of functions to reduce memory usage. These are more properly the subjects of a computer science dissertation; yet as they bear directly on the possibility of using computational tools for historical research, I have had to explore them to a limited degree in order to produce working software.

**3. Once the text is in the form of atoms, it must be made into a graph: connected by relations of various sorts, and saved in some form of database.** To create relations, I have depended almost entirely on SpaCy, and then more recently on the OpenAI word embeddings, combined with cosine similarity calculation. Other approaches could also be used: instead of word or sentence vectors, some researchers use the tried-and-true tf-idf statistic to measure similarity, to good effect, and there are other approaches being developed too.

The database format bears significantly on performance down the line. I used Neo4j and their Cypher programming language, but other options exist, and better performance can be had in this area than I currently am able to provide. However, my solution has the advantage of keeping several layers of the programming stack separate, and I solve many of the problems in maintaining a database by simply using an off-the-shelf solution that works well enough.

**4. The connected graph can be used for simple statistics and visualization, as shown in Chapter Four. It is also amenable to more sophisticated analysis, such as the approaches described and implemented in Chapter Five.** This is a field in rapid growth. Many areas of business and science are currently discovering Topological Data Analysis as a possible option to find useful information, and several software packages exist to implement TDA to find clusters and holes in a network.

Once clusters and holes have been located and their bounding nodes identified, the interpretive task of understanding what such a topology might mean for the historian is only beginning. Once vectorized texts have been connected through similarity (or other means) into a graph structure, what relationship do holes and clusters in that structure have with the original space they were generated in? How could reliable inferences be made from topological holes

---

<sup>308</sup>Pandas dataframes can be sped up by using Polars or by carefully thinking about how the underlying Numpy arrays are accessed and modified...

(and also: from holes in 1, 2, 3, ... n dimensions) to the meaning of the language model space that precedes them, and from that space back to the space of human discourse? The main intuition in this thesis is that there exists a relationship between these three alternate representations of the same information, but the ways in which it can be traversed in reverse, from vectors to text, remain largely unexplored.

Many questions remain: How can the vector space of the language model be interpreted in itself? In other words, what conclusions (if any) can be drawn from the specific place a text falls in the space? Does the space itself have any interpretable meaning? There exist some ways to explore these questions. For example, the language model could be trained to generate a fixed-length sequence of words based on the input vector, or it could be trained to generate an open-ended sequence of words until it reaches a stopping condition, such as a specific word or character.<sup>309</sup>

Another potential approach to understanding the relationship between holes and clusters in a graph structure and the original space from which they were generated is to consider the concept of dimensional reduction. Essentially, this involves mapping high-dimensional data onto a lower-dimensional space in such a way as to preserve certain important characteristics of the original data, as we have done above using the UMAP algorithm. Once this mapping has

---

309 This problem has already been solved fairly well for automated writing of descriptions of code blocks, such as has been implemented in ChatGPT. See Yuding Liang and Kenny Zhu, “Automatic Generation of Text Descriptive Comments for Code Blocks,” *Proceedings of the AAAI Conference on Artificial Intelligence* 32, no. 1 (April 27, 2018), <https://doi.org/10.1609/aaai.v32i1.11963>. The authors state that “We propose a framework to automatically generate descriptive comments for source code blocks. While this problem has been studied by many researchers previously, their methods are mostly based on fixed template and achieves poor results. Our framework does not rely on any template, but makes use of a new recursive neural network called Code-RNN to extract features from the source code and embed them into one vector. When this vector representation is input to a new recurrent neural network (Code-GRU), the overall framework generates text descriptions of the code with accuracy (Rouge-2 value) significantly higher than other learning-based approaches such as sequence-to-sequence model. The Code-RNN model can also be used in other scenario where the representation of code is required.” The specific problem I describe has also been attempted by several researchers. See for instance <https://aclanthology.org/D15-1188.pdf>, and [https://renom.jp/notebooks/tutorial/time\\_series/text\\_generation\\_using\\_pretrained\\_word\\_embedding\\_layer/notebook.html](https://renom.jp/notebooks/tutorial/time_series/text_generation_using_pretrained_word_embedding_layer/notebook.html). A simple way to solve an online word-guessing game (semantle.com) is posted here: <https://pastebin.com/FzjXGn8X>. The programmed solution depends on choosing words that have greater similarity to the previous guess, and can quickly triangulate the correct answer. However, this simple solution is only able to consider single words, and is not context-sensitive as would be needed for our problem.

been performed, it may be possible to identify patterns or structures in the lower-dimensional space that correspond to certain features or characteristics of the original data, as we have mentioned. For example, clusters in the lower-dimensional space may correspond to certain themes or topics present in the original texts, while holes or gaps in the lower-dimensional space may correspond to areas where there is a lack of coverage or information in the original texts.

However, making reliable inferences about the meaning of the language model space based on the topology of the graph structure is a more complex task. It may be necessary to develop new techniques and methods for analyzing the structure of the graph and relating it to the meaning of the language model space. Ultimately, the goal would be to develop a deeper understanding of the relationship between these three alternate representations of the same information, and to find ways to effectively traverse between them in order to gain new insights and knowledge. I find this work highly fascinating, and will pursue it in the near future, but am unable to produce results in time for inclusion in this dissertation.



## Chapter 6 – The Struggle for Synthesis

Throughout these pages, the main burden of my work has been the instantiation of a method of information processing and visualization that can help historians reach greater intelligence of their data. The goal of this method is to provide ways to account for a broader range of information and points of view, where quantity matters. The endpoint is the “struggle for synthesis”, the interpretation of the information that has been collected, organized, and visualized.

The first specific examples I gave of this process had to do with the comparison of Mesopotamian and Biblical spiritualities.<sup>310</sup> In Chapter 2, I mentioned several experiments I conducted to support the affirmation that there is a structural difference between the two spiritualities. If such a difference exists, there will also be different methods of understanding that are necessary and adequate to the different materials under examination. As we saw, the digital approaches I used lend support to such a reading. The strongest evidence was found in a graph representation of the relations between humans, gods, and the natural world in two representative bodies of text: the structures of the relations in the two corpora are markedly different.

In order to produce that representation, I fragmented the relevant texts beyond almost any measure. It would be reasonable to wonder whether such a procedure could be reconstituted into unity, yet in fact those fragments and their interrelations do indeed give reason to believe that there are structural differences between the spiritualities of the two peoples in question, and that those differences have the form that can be deduced by using other methods. In this case, the Memex appears as an extension and an addition to the tools already available. It gives us a new way to arrive at the same conclusions.

---

<sup>310</sup>Here as elsewhere I follow the terminology used by Buccellati.

In Chapter 3, I described the complete workings of the Memex and applied it to a simple historiographical question: to what extent do the reviews of a given volume adequately cover the richness of the contents of the volume, and what specific parts of the volume were given attention by the reviewers? Conversely, which parts of the volume were not discussed? The conclusions we can draw from the toy example in Chapter 3 are not particularly significant, but the method can be extended to a larger body of work and lead to interesting and actionable conclusions.

In Chapter 4, I did just that. I compared a book to a complete bibliography, in order to discover if the bibliography adequately supported the entire book or if there were important gaps in coverage. The view of the information provided in this way led me to find topics that required further attention in the bibliography, and also to characterize three studies of Mesopotamian religion in ways that confirm and specify the approaches their three authors took. The Memex again offered an additional way to view data. I also built a way to traverse the bibliography as a network of interrelated topics, not only as an alphabetical list. In these cases the Memex offers the abilities of a clerk, remembering details and bringing to light information that is not easily seen. Here, the value of the tool, more than directly interpretive, is to be found in its aid in discovering weaknesses in an argument, blind spots of the authors. However, this methodology segues directly into further methodologies, that make use of the same graph object but provide more important interpretable information.

Chapter 5 described three more ways of using the Memex to view the relationships between large amounts of text. First, using a dimension-reduction methodology, I mapped distributions of keywords and concepts to the flat space of the page, giving an intuitive way to compare and contrast the way a given author covers the intellectual space defined by the dataset. The same method was also applied to entire corpora. By representing the spatial distribution of topics as a function of time, I also showed how the Memex can be used to represent the development of a discipline<sup>311</sup> like a filmstrip of the topical distribution of

---

<sup>311</sup>This is one of the most interesting applications of the method, and unfortunately it brought me up to the limits of what was possible with the technical means at my disposal. I hope in the near future to develop this method further and lead to more interesting results.

documents over time. It is possible to characterize locations in the distribution by paying attention to their present contents, which will remain stable at future times as well.

While we cannot characterize space in absolute terms in the UMAP projection, on the basis of the texts encoded into the space, we can predict the contents of other objects projected into that same space. This means that we are able to see in a single birds-eye view where a field of study has concentrated and where it has not. This can give a very quick overview of what has been done in a given field and what has been not. The usefulness of this tool depends on the availability of easy to use clean datasets. Many are being produced, but they are not always easily available questions in ancient history. At present, the best source of workable datasets is the Constellate service provided on JSTOR. Not all important journals are available, unfortunately.

The final topic treated in Chapter 5 were the interpretive roads opened by the granular detail in the graph object (the “microscope” aspect), which gives us a powerful search engine to locate relevant affirmations across a dataset. By doing so we can be reminded, or learn for the first time, about ideas that are similar or different from our own. This can be very useful to develop a critical approach to an issue. It is easy to remain trapped within the ideas and interpretations that we have of a given issue, and not see the possibilities of other points of view.

In the Memex, other points of view are surfaced continually, automatically connected to the existing network of ideas. It acts as an external memory of the ideas encoded by the people who have been using the tool. Interpretation here is aided by the presence of a multiplicity of points of view, collected around the nodes where those points of view offer their most relevant and important contribution.

Even over the space of the three years I have spent writing this dissertation, which is not a very long time as a researcher, going back to my Memex I realize that there are many texts I do not recall in detail, and some that I have entirely forgotten. Of course, existing technologies like keyword searches can help when I have a digital document, but the Memex offers an additional advantage because it provides a single location where all this information is stored and is

quickly accessible. It stores a granular representation of the web of connections between concepts and can give them back to me at the moment of need.<sup>312</sup>

Thinking of interpretation as a process opposed to analysis begs an important question. Analysis aims at situating the available information as clearly and precisely as possible. An interpretation based on less complete information is of worse quality than one based on more complete data. Sometimes it seems possible to arrive at a good interpretation with scarce data, and to grasp a lot of information through intuition and other human abilities that are difficult to quantify. I marvel at the ability of scholars in the past to interpret fragmentary information in this way. But the fact that intuition is difficult to quantify does not mean it would not improve through a greater amount of data or greater quantification of the existing data. Without detriment to the truly great scholarly efforts of the past, one can feel justified frustration with conclusions that are presented as valid interpretations, when they clearly do not take into account a wide variety of relevant information. And herein lies the problem of too much information: even the most well-intentioned and honest scholar is unable to take into account all the relevant information. This is what the Memex aims to ameliorate.

Could machines substitute humans? The question is frequently encountered in the year 2023. Notwithstanding my attempts to confine this thesis to a reflection to history, and to a use of technology that can serve the historian, between the lines of many of the preceding pages there are allusions to this question.

The most prudent way to face this question is to avoid it. This approach is somewhat cowardly but fully justified when addressing an issue as difficult as this one. For this reason, I usually have stated that the method I have developed and the computer program I have written aims to be nothing more than “an efficient clerk”. I ring the bell for collaboration between humans and machines. I do not feel justified in affirming the immutable superiority of human intelligence, especially as its limits are coming so sharply into view. However, I am firmly convinced that at present, the best tool available to reach intelligence of history is indeed the

---

<sup>312</sup>More than once, while writing the dissertation and referring to the notes in the Memex, I realized that my previous reading had surfaced ideas that needed to be explored further, and this process modified my conclusions in the dissertation. An “extended mind” truly did augment my own powers of observation and reflection.

human mind. Indeed, as far as we can tell, the *only* consciousness in the universe that can reach such intelligence is the human mind. Along these lines, I intend this thesis to be a small, positive contribution to the ways historical information can be managed, archived, retrieved, and queried. But the larger question remains, so let me conclude this work with a short reflection on it.

The development and current popularity of digital humanism contains several different streams within it. One is a sort of glee, like a happy sigh of relief: at last humanists can be considered ‘scientific’, at last their production can be quantified and bear the prestige that inheres to the priesthood of Tech.<sup>313</sup> This first group might easily fall prey to a form of unreflective enthusiasm and laziness, since the first flush of applied statistics and the speed of computation can seem to substitute for the difficulty of thought and interpretation.

A second undercurrent one meets is fearful. Although it is necessary to give a show of using computers for one’s work in an academic environment that is increasingly governed by business interests, the second group is intimately convinced that real humanists do not need them for anything more than typesetting and email communication, and that the great majority of quality work has been done without their aid. The members of this group are dismayed at the immense investments made for technical faculties, and as their own funding is paltry by comparison, they wonder about the future of their disciplines.

There is also a third current, characterized by the attempt to re-imagine humanism, by both embracing technical developments and using them to their utmost, while simultaneously refining the human characteristics that govern and direct the use of tools and the interpretation of their results. It is difficult to maintain the equilibrium of spirit and the technical competence necessary to understand the rapidly changing world of technology. And for those who manage to keep up to date with technique, the spirit of technical power remains a siren, always threatening to pull the agile, and fragile, craft of humanism into the rocks.

I belong to this third category. I remain optimistic about the enduring importance of humanism, and I dedicate much of my time to understanding new technical developments that

---

<sup>313</sup>Some references to these ways of thinking are made in Daniel Allington, Sarah Brouillette, and David Columbia entitled “Neoliberal Tools (and Archives): A Political History of Digital Humanities”, Los Angeles Review of Books (2016).

bear on my studies. I also feel the pull of technique, its fascination as well as its unfortunate tendency to deaden and render superficial one's reflection. At times, my thesis reveals this inner tension, although I have for the most part decided to choose a single path and follow it to the end. Like other sailors before me, I could not resist listening to the siren's sing. Luckily, I was lashed to the mast by several captains more deeply versed than myself. I conclude, then, with a reflection on the nature of digital technique applied to historical research.

## What is Digital Humanism?

In 2016, the Los Angeles Review of Books published a critique of digital humanism penned by Daniel Allington, Sarah Brouillette, and David Golumbia entitled "Neoliberal Tools (and Archives): A Political History of Digital Humanities". The article offered a telling of the rise of digital humanities that emphasized its origins in the convergence of interests between "silicon valley venture capital" and "university administrators and funding bodies" [All citations in this paragraph, Allington et al. 2016]. According to the authors, this convergence aims at the "corporatist restructuring of the humanities", and constitutes a "neoliberal takeover of the university". The authors suggest that conservative, anti-interpretive ideas, especially as embodied in the University of Virginia, are at the heart of what DH is all about. Frustration, or "disdain, and at times outright contempt", as the authors would have it, with existing methods of scholarship led to the development of new procedures and standards for humanities research by DH enthusiasts, which amount to nothing less than a redefinition of the humanities. If Digital Humanities is a "reactionary force in literary studies, pushing the discipline toward post-interpretive, non-suspicious, technocratic, conservative, managerial, lab-based practice", we might do well to pay deep attention to what is happening and not limit ourselves to knocking down the straw man of corporate collusion.<sup>314</sup>

Some of the points made by Allington, Brouillette, and Golumbia are especially well taken, as when they say that "much of the more interesting side of Digital Humanities research has a

---

<sup>314</sup>If frustration with interpretation and with activist politics masquerading as research has generated a "postcritical" backlash, one should evaluate carefully the reasons for this frustration before ascribing them to a corporate conspiracy. Those reasons might include the philosophical and linguistic issues presented with wicked and profound satire in *The Square*, which won the Palme d'Or at the Cannes film festival in 2017.

tendency to resemble a slapdash form of computational linguistics adorned with theoretical claims that would never pass muster within computational linguistics itself” [Allington et al. 2016]. Alas, this can be true—in the first flush of applying computational methods to large corpora, a scholar may become inebriated. But infatuation and laziness are not confined to DH practitioners. Every intellectual fad generates its share of overly enthusiastic articles, some of which are embarrassing a few years later.

As the three authors repeatedly suggest, there is much to be said about the soft-power relation between funding sources and research conclusions. It is true that large private endowments and corporate grants come with strings attached to an underlying project that is usually more or less apparent. Research funded by a foundation with a clear mission statement may well seek to validate that preconceived idea, and research funded by a large corporation that deals in the the “new oil” of information, may well be the apparently innocent cover for the production of a “flexibly efficient” archive that can then be used for all sorts of occult postindustrial purposes.

The authors are however unfair in claiming that DH is “*not* about, despite its explicit claims, [...] the use of digital or quantitative methodologies to answer research questions in the humanities” [Allington et al. 2016]. This sort of hyperbole serves only to raise hackles. It does not aid in uncovering the importance, or lack thereof, to be ascribed to the current interest in DH. While the authors state that what is at stake is the “redefinition of technical expertise as a form (indeed, the superior form) of humanist knowledge”, they do not engage with this idea beyond denigrating it.

Brian Greenspan suggests that the positions expressed in the article might be resentful or nostalgic:

the LARB piece encapsulates a large undercurrent of *ressentiment* within academia that blames the digital humanities and neoliberalism alike for sapping both prestige and resources from the “pure” scholarly pursuits of merely thinking and writing, which allegedly require only books, pens, and paper; and need not involve any newer technologies at all, let alone teamwork, labs, or large operating grants. That attitude is, of course, hugely disingenuous: it perpetuates the monastic myth of the isolated (tenured) scholar as ideal, while ignoring how little anyone could get done today without the computers, email clients, catalogs and databases, e-journals, cloud storage, online book resellers, and social networks that keep us connected to the world of scholarship, not to mention online travel agents for booking passage to conferences and research archives. In today’s academy, we are all already digital.<sup>315</sup>

---

315 Greenspan, Brian. “The Scandal of Digital Humanities”, in Gold, M. and Klein, L. eds., *Debates in the Digital Humanities*, University of Minnesota Press 2019. Available at

In any case, the issues raised in the LARB article point further than its authors allow. DH is not merely a monkey-wrench that corporations are using in order to re-configure the academy for their own purposes. Even if corporations were to be chided or feared for legitimate motives, the tools and concepts that underlie DH as a practice and as a theoretical approach to knowledge are more profound than the profit motive.

I will address three aspects. I begin with what might be considered as a first-order understanding of what the DH approaches are: the “database problem” regarding the collection of data and its storage and dissemination in standardized and easily accessible formats. A second issue which points more directly at the core meaning of DH and its tools is the coordination of the center and the periphery, the particular and the universal, through the unprecedented abilities of computerized algorithms to render explicit the entire arc that leads from atomistic data to large-scale conclusions. (These capabilities are at the core of what I have presented in the previous chapters and instantiated in my Memex.) Never has the process of induction been so transparent. Third and finally, DH promises (or threatens) to redefine humanism itself.

## **The Database Problem**

Archives are undoubtedly an important part of the work of humanities research, and in particular of historical research. In this setting, does DH represent anything truly new? It is not easy at first glance to determine whether the exponentially larger quantities of information that are now available to the scholar, the corporation or the government are fundamentally different from older all-encompassing projects such as the encyclopedia.

Some hold the hope that in the power of databases and querying software, and of new logical structures and ways of synthesizing large quantities of information such as I have presented throughout this dissertation, new insights may come into being. Perhaps hidden within the mass of information that digital databases can contain, new patterns may be found that will unravel mysteries that have proven heretofore impenetrable. The patterns can be

---

<https://dhdebates.gc.cuny.edu/read/untitled-f2acf72c-a469-49d8-be35-67f9ac1e3a60/section/4b6be68c-802c-41f4-a2a5-284187ec0a5c>.



worked into narratives. The archive is not in itself a narrative, but it contains the raw material from which the historian can construct one.

Some databases are the digital equivalent of libraries or encyclopedias (Wikipedia), others are repositories for programming code that can be modified (GitHub), or digital photographs of the collections of prestigious museums (Museo.app). While these efforts are sourly criticized in the LARB article as serving “neo-liberalism”, it would seem fair to regard them instead as something larger and more interesting.

Christian Greco, the director of the Museo Egizio in Turin, has stated that the museum he directs is putting everything online with open copyright.<sup>316</sup> Greco believes that museums today should work together to try to construct a cultural landscape where objects are put in their context. Although objects and artifacts are scattered across the world, new technologies allow us to rebuild their original contexts, and so to accomplish a restitution of each object to its fuller identity. Greco expresses himself by saying that “each object is like the page of a book, and we must put together the book.”

While the idea of reconstructing the context of an object is fraught with epistemological difficulties, and Greco’s comment may overstate his case, it does show that the database problem is not essentially a question of reducing friction and “flattening” the world.<sup>317</sup> It is rather an issue of thinking beyond the boundaries of the local institution and considering the whole of the ancient society represented by the fragmentary artifacts dispersed around the world in museums. Further, it is also an issue of considering the whole of humanity as the rightful public which any museum serves.

Digital archives shared across the world seem at first glance to offer a mere difference in quantity, but perhaps they also offer a difference in kind: a way to think of scholarship as a global phenomenon, rooted in truly local institutions that play a necessary and interconnected role in the preservation and dissemination of information about the past.<sup>318</sup> Databases are

---

316*The Future of Memory*, a talk given at the University of Pavia on Feb. 9, 2021. There is a complex jurisprudence regarding free online archives, and in Italy it is not possible for state-funded museums to do as the Getty Museum has done. The Museo Egizio can, because it is a private foundation.

317As in Friedman’s bestselling essay. See Thomas Friedman, *The World is Flat*, 2005.

318In prof. Fabio Rugge’s address given in Pavia on 8 Feb. 2021, entitled “The State and Other Stories”, a related point was made. Prof. Rugge presented his observations on the idea that the state exists in order to “produce security” (*sicherheitprodukten*). The classic problems that the state solved in this sense, such as

foundational for the creation of knowledge. The more connected they are, presumably the better their contents can be interpreted by scholars, and the more nuanced and complete a field's conclusions will be—if tools exist to manage the immense complexity of the information.<sup>319</sup>

## From Data to Conclusions

Of a piece with the aspiration to create global archives is the dizzying idea that such a collection of information might permit scholars of the humanities to reach greater levels of certainty in their hypotheses and conclusions. I share this hope to some degree: as my explorations of the Memex in chapters Four and Five have shown, it is possible to reach broad characterizations of large amounts of text data that would escape the eye of even a careful scholar, simply because there is so much data.

In the academy today, there is massive enthusiasm regarding DH, in both funding sources and individual scholars scrambling to learn enough Python and NLP to apply some data science to their humanities research. In my view, this enthusiasm is part of an intellectual groundswell—which could be partially misguided, but which is not in essence the result of a conspiracy of big business, as the LARB article would have it.

---

bandits, are no longer particularly important. The problems of the future, in prof. Ruge's estimation, regard the equilibrium of the world itself, especially in terms of ecology and transnational governance. This fact requires a new understanding of what states do—one which is at once both more global and more local. The dynamic is similar to the observation that the construction of shared digital archives can generate a more global *and* more local presence in museums.

319 See Samuel Hays' 1974 review of four volumes dedicated to quantitative methods applied in historical research, "Historical Social Research: Concept, Method, and Technique". Hays writes that "The tragedy of conventional approaches to history which still dominate the profession is a preoccupation with dramatic personalities and events, top-level national happenings, uniformities in national culture, the ideas expressed by major "thinkers," and a potpourri of fractured segments of human endeavor which have constantly divorced historians from a systematic understanding of human life. Rarely capturing the imagination of historians are such questions as the patterns, shared by many people and persisting over time, of perception, mobility, family life and kinship ties, community and/or segmental relationships, religious perspectives within and across denominations, geographical ranges in human relationships, inequalities in economic condition and political power, or dominant-dependent relationships between geographical areas at different stages of economic development. *Ordered patterns in the structure of human relationships and in the processes of change in those relationships over the course of time have escaped historians as the basic context within which to think, ask questions, pose problems, and organize courses and textbooks.*" My emphasis. Since Hays wrote those lines, his desires have been well served by DH tools.

Three highly successful intellectual products, the logical formalization of certain segments of the world (Cycorp)<sup>320</sup>, large-scale data mining operations (Palantir)<sup>321</sup>, and a rules-based algorithmic interpretation of history to guide investment decisions (Bridgewater)<sup>322</sup>, apply data science to sensitive decisions by governments, large corporations, and hedge funds. While these ventures *are* a capitalistic leveraging of information, their products themselves resemble closely the project of studying history by using the tools of quantitative natural science.<sup>323</sup>

To what end might we apply mathematical methods to the study of history? Statistical methods have been widely applied to social science, although attempts to generalize from statistically significant measures to the enunciation of “laws” of human behavior have met with little practical success.<sup>324</sup> Perhaps the concept of “law” should be revisited in favor of a more

---

320Cycorp represents one of the oldest ongoing projects in Artificial Intelligence. It is a logical approach to representing and organizing knowledge in a machine-readable grammatical structure. New questions can be posed to the “inference engine” and leverage the existing “knowledge base” Cycorp has developed over the past 30 years. See <https://cyc.com/>

321A more recent outfit, Palantir employs about 2500 people and serves large governments and industries with data-mining software. It has a market capitalization at this writing of approximately 500 billion dollars. Compare that firepower with what the average university can deploy, and consider that they have little need to co-opt universities for their purposes. See <https://www.palantir.com/>

322Bridgewater founder Ray Dalio is one of the most successful investors alive. He has described his mechanistic philosophy for thinking about knowledge and acting upon it in his bestselling *Principles*. A recent application available on the App Store, the “Dot Collector”, is a form of his thought applied to any organization that wishes to collect and interpret granular data from its employees. The “Dot Collector” is yet another expression of the underlying project to organize and interpret data in order to increase the value of an organization. See <https://www.principles.com/> and Ray Dalio, *Principles: Life and Work*, Simon and Schuster (2017).

323These three businesses seem more advanced than anything the academy is currently doing in this field, not least because their “endowments” dwarf what any university can deploy. These businesses, and many more, are much larger than a scheme to cost-cut, as Allington et al. would have it, by reducing long-term employment of university staff: they represent a sea-change. To the list, we might add Google, which is “organizing the world’s information” according to its mission statement. Consider, too, books like Y. Harari’s *Homo Deus*, which makes a somewhat plausible case for an evolutionary leap in human capacity as a result of recent technological advances. Man has always wanted to become God, in Harari’s telling: and now this result is technically much closer to being realized than ever before.

324The “reproducibility crisis” is a widely-known feature of this failure. If “laws” were truly discovered in statistical social research, studies ought to be more reproducible than they actually are. The point was made by Clifford Geertz in his 1980 article “Blurred Genres: The Refiguration of Social Thought”: “many social scientists have turned away from a laws-and-instances ideal of explanation toward a cases-and-interpretations one, looking less for the sort of thing that connects planets and pendulums and more for the sort that connects chrysanthemums and swords.” See also Sewell’s point in *Logics of History*: “Temporal heterogeneity implies causal heterogeneity. It implies that the consequences of a given act are not intrinsic in the act but rather will depend on the nature of the social world within which it takes place. This assumption is quite contrary to the practices of mainstream social scientists, whose entire mode of operation is to discover and apply general causal laws, laws implicitly or explicitly assumed to be independent of time and

neutral and humble term, such as “pattern”. William McNeill put it with wonderful concision: “Pattern recognition of the sort historians engage in is the chef d'oeuvre of human intelligence. It is achieved by paying selective attention to the total input of stimuli that perpetually swarm in upon our consciousness.”<sup>325</sup> The language of patterns and selective attention fit well with what DH practitioners actually find in their research.

Historians tend to regard a statistical approach with caution. As Marc Trachtenberg puts it, “the covering-law approach was unacceptable because it failed to allow for human agency—for the role that individual human beings play in shaping the course of events”.<sup>326</sup> Although he does not use the term, Trachtenberg’s comment indicates an all-important question: freedom.<sup>327</sup> Historians tend to view their domain not only as the stochastic patterns of blindly careening billiard balls, but also as the patterns made by individuals and their decisions.

Throughout this dissertation, I have preferred a graph model of information in order to represent the relations between the sources in my database, instead of seeking causes directly. This choice was determined by the desire to consider each piece of information not only as an object that can be considered statistically, as in distributional analysis or the cruder tools of early digital humanities, but also as a part of whole: a whole discourse, a whole structure of reasoning. The distribution of elements and the distribution of their relations is an important aspect of the slippery, ever-crumbling domain of human language.

---

place.” p. 10.

325In William McNeill, “Mythistory, or Truth, Myth, History, and Historians,” *The American Historical Review*, February 1986, 2, <https://doi.org/10.1086/ahr/91.1.1>.

326Trachtenberg, p. 16.

327The postulate regarding whether or not human agency is a reality, heavily impacts any attempt to describe the “laws” of human behavior. On the face of it, a free human would be *ipso facto* not describable in terms of laws. Aggregate human activity could be described using the instruments of statistics, although if these tools were successful in modeling human behavior, it would seem that such behavior would have to be understood not as *free*, but as *random*. A huge bibliography exists on this topic. In 1963, Sartre shrugged his shoulders: “Each one is free to believe that physico-chemical laws express a dialectical reason or not to believe it.” *Search for a Method*, p. 12. For one accessible treatment, which concludes that freedom is an illusion, see Hofstadter, *I Am a Strange Loop*. See also Duncan, *How Intelligence Happens*. A good overview of research that leans in favor of a form of human freedom can be found in A. Lavazza, “Free Will and Neuroscience: From Explaining Freedom Away to New Ways of Operationalizing and Measuring It”. Some very recent research that suggests a sort of “freedom” even in relatively simple animals such as the cuttlefish. See Schnell et al., “Cuttlefish exert self-control in a delay of gratification task”.

This choice was also influenced by recently developed tools in DH that aid in graphical representation of complex data, such as Stanford's *Palladio*.<sup>328</sup> These tools have made it possible to visualize social networks, trade routes, the chronological growth and decline of word usage, migration patterns, and myriad other dynamics in an intuitive manner. Some interesting conclusions can be reached easily in this way, and the tools promise to open more and more fields to this form of visualization via the underlying logical structure, a "graph database", which represents data not as "spreadsheet" tables of information, but as nodes that are connected by relations.

An approach of this sort was expressed by William and John Robert McNeill in their *The Human Web*. The very word "web" as the overall metaphor for organizing the book and its interpretation of human history speaks to an idea that is similar to the graph database structure briefly referred to above. In the McNeills' view, the driving force in history is the "human ambition to alter one's condition to match one's hopes"<sup>329</sup>, which is coordinated and channeled by webs of social communication and cooperation. The overall arc of human history as seen by McNeill father and son moves from a time of "simple sameness to diversity toward complex sameness".<sup>330</sup> This is because "best practices spread; societies settled on a narrower range of traits, beliefs, institutions, all compatible with life inside far-ranging interactive webs". What kind of future can be hoped for in this view? The elder McNeill concludes his contribution to the book saying that

we also need face-to-face, primary communities for long-range survival: communities, like those our predecessors belonged to, within which shared meanings, shared values, and shared goals made life worth living for everyone, even the humblest and least fortunate. If so, perhaps the most critical question for the human future is how cell-like primary communities can survive and flourish within the global cosmopolitan flows that sustain our present numbers, wealth, and power, without being disrupted by those flows and without disrupting them. In other words, we need a new symbiosis all over again.<sup>331</sup>

Indeed: the coordination between the very small, local community, and the world-wide web of connections linking every industry, academy, and government is the decisive question for

---

328<http://hdlab.stanford.edu/palladio>

329John Robert McNeill and William Hardy McNeill, *The Human Web: A Bird's-Eye View of World History*, 1. ed (New York: Norton, 2003), 4.

330McNeill and McNeill, 322.

331McNeill and McNeill, 326.

politics in our time. The ecological question makes such coordination all the more pressing: we can no longer act as if the consequences of the actions of a community or state can be contained within the geographical boundaries of the same. We are all connected, whether we like it or not.

This affirmation leads us to consider that the question about Digital Humanism is not a sterile debate between luddites and progressives, or between socialists and capitalists, within the academy. Rather, it is one aspect of a much larger groundswell in contemporary society, that deeply affects government and industry as well as the academy: the desire to leverage digital tools in order to manage and preserve the connections between the very small and the very large. One must always be on guard against the logically impossible conceit of a “grand unified theory”, not least because such a project is self-referential and cannot avoid the paradoxes of infinite recursion, but the desire to organize and interpret a large database of information and an ever-more granular web of connections between things, people, and institutions, is as old as civilization itself.<sup>332</sup>

Conceiving of information as intrinsically connected, and formalizing those connections mathematically, promises a new level of control of the data and therefore a new level of ability to generate robust conclusions. Intuitively, one may expect that if the intermediate steps between pieces of data are explicit within a graph, it may become possible firmly to ground overarching intuitions in the details that are clearly known with a good degree of certainty, as I have attempted to intuitively show in using the Memex as microscope.

Philosophically, the aim behind such a project is the desire to interpret data without superimposing a pre-existing ideology upon the data. One may hope to avoid being a

---

<sup>332</sup>Michel Foucault, with characteristic verbosity, puts it thus: “What we have described, for example, as the Analysis of Wealth or General Grammar, thus according them what was perhaps a highly artificial autonomy, was it not, quite simply, political economy in an inchoate state, or a stage prior to the establishment of a truly rigorous science of language? Is it archaeology trying, by means of a retrograde movement whose legitimacy it would no doubt be difficult to establish, to regroup in an independent discursive practice all the heterogeneous and dispersed elements whose complicity will prove to be necessary to the establishment of a science? Again, the answer must be in the negative. What was analyzed under the name of Natural History does not embrace, in a single figure, everything that in the seventeenth and eighteenth centuries might validly constitute a prototype of the science of life, and figure in its legitimate genealogy.” Michel Foucault, *The Archaeology of Knowledge*, Routledge (2002), p. 180. What is happening in Digital Humanities may look like a refoundation, a total beginning, but its underlying idea has been proposed before more than once, as Foucault points out.

“marxian” or a “neoliberal” or any other categorical label, and instead simply be a historian: an artisan working in a precise domain of knowledge. The way this can be achieved is through attention to the inner-referential structure of the information.

One project developed along these lines, known as “Global Labor History” (GLH), attempts to deal with the extremely complex issues regarding labor and economy without belonging to an ideological camp.<sup>333</sup> One of its practitioners, Prof. Stefano Bellucci, presents GLH as a way of looking at labor history “with a postmodern twist”: as a new way of looking at complexity, with no overall guiding ideology. In his view, the role of the historian is to abandon a linear view of history, and accept the absence of a grand narrative, in order to avoid the pitfalls of justifying an *a priori* philosophical position with the historian’s tools.<sup>334</sup>

Shedding the teleologies and ideologies that distort theories and expositions of data is certainly a positive development. The project is initially appealing, but insofar as GLH purports to offer a humble, neutral language, capable of responding to and resolving the contrast

---

333One earlier attempt at the same goal is represented by Wittfogel. According to Donald Worster, “Wittfogel arrived at this position in attempting to solve a problem that earlier had baffled Marx: Why were the major civilizations of Asia so different from those of Europe, so lacking in capitalist development, and so unpromising for a Communist revolution? The answer, Marx had vaguely indicated, lay in the advanced water systems built by Asians to provide irrigation for their arid lands; from that base a distinct form of society had evolved in China, India, and the Near East. This much Marx realized, but he was at the same time reluctant to see in nature much more than a passive landscape in which human labor toiled and created. It was Wittfogel who took the argument over and insisted that the natural environment is not really passive but rather is a powerful determining force throughout history. People are forever struggling with the land in an ongoing ecological dialectic: there is the gist of the Wittfogel theory. The earth gets changed in the unfolding dialectic, but so do the people. For example, in the absence of ample rainfall Asian farmers in several places brought water to their fields. Eventually they created what Wittfogel called a “hydraulic society.” As their manipulation of water became more and more large-scale, they were forced to reorganize their social structures into elaborate hierarchies of power—into a chain of pharaohs, emperors, bureaucracies, and highly centralized states”. See “*History as Natural History: An Essay on Theory and Method*”, p. 5. In Worster’s view, “there is no Darwin in History”, and attempts like that of Marx have failed to predict how history actually goes. He concludes that it is time to let the ecologists give prediction a try.

334See Antoinette Burton’s statement that “at their critical best, histories of the global can and should act as ‘reorientation devices’, reminding us of the inheritance of empires past, present and future which presses down on them as well as their radical potential for directing our analytical presumptions and our methodological energies in different directions at once—and for enabling us to see not just the world and its historic fulcra, but the off-centre, the ex-centric, the polycentric. The anti-centric and the ‘remotely global’ as well.” In “Not Even Remotely Global? Method and Scale in World History”, p. 328. What DH methods could aid in the desire to account for the ex-centric as well as the centric, to preserve legitimate hierarchy along with the details of the data? An answer to this question would have to engage the fact that resources are scarce, and are usually allocated by someone who is (hierarchically) above the recipient of the resources, which inevitably colors the quantity of attention paid to the ex-centric.

between competing ideological interpretations of history, such a project necessarily contains, however implicitly, yet another teleology.<sup>335</sup>

A second project that aims at robustly connecting data to conclusions along the entire arc of inference is one I have repeatedly referred to throughout this dissertation: Giorgio Buccellati's website dedicated to his excavation of Urkesh.<sup>336</sup> The structure of Buccellati's archaeological work and publications reflects the kind of intrinsic link between "atomistic" data and conclusions that digital methods promise to render more complete and transparent. On his website dedicated to the data collected in Urkesh, he writes that a digital narrative in archaeology is characterized by three qualities:

1. Non-contiguity/capillarity. – With regard to the data, a nexus is seen where none is immediately apparent, allowing connections, in a complex referential tree, between the most minute supporting detail and the most generalized conclusion.
2. Discontinuity/reconfiguration. – In the construction of the narrative, elements of disparate nature and origin are reshaped into alternative sequences.
3. Non-linearity/multi-linearity. – Within the flow of an argument, explicit linkages connect multiple juxtaposed sequences.<sup>337</sup>

These qualities imply that the bedrock of archaeological reasoning is to be found in the fragments of data and in the connections between the fragments that emerge from an excavation. By representing these connections in a "complex referential tree", such as a graph database, the road that leads from specific elements to other elements can be reconstructed.<sup>338</sup> This reconstruction, or narrative, is not usually the only one possible. Narratives are usually "alternative sequences". Finally, as in the third level in Buccellati's outline, an argument may be constructed from the pieces of narrative that emerge from the patterns in the data.<sup>339</sup>

---

335The contrast between the marxist or marxian interpretation and that proper to GLH, which Bellucci mentioned in particular with respect to French scholars, may also be due to the lack of a shared "language" of concepts. In fact, in order to study a truly global phenomenon, universal categories of thought and speech are applied to the variety of experiences, with inevitable imprecisions.

336Urkesh.org. See also the ongoing development of a parallel system of websites dedicated to Mesopotamian culture in general, 4banks.net.

337<https://urkesh.org/main/main3a.htm>

338The reconstruction of these connections is the aim of the long and patient work of the archeologist. It is a process akin to the decipherment of ancient languages: first, the graphical signs are distinguished, then their syntax is deduced from the actual occurrences of each sign, and then (in many, but not yet all, cases) it becomes possible to link syntactically correct usages to extra-linguistic meanings, often through bi-lingual texts.

339As Buccellati explains elsewhere in his site, "The Urkesh Global Record is an argument driven global archive. It is an "archive" in that it contains and makes available in an ordered fashion the record of the



The patient work of reconstructing narrative sequences from fragments is an expression of the historian's act of faith: there is a "story" to be told, there is a narrative under the data, however fragmentary it first appears. Buccellati expresses the specific advantage the archaeologist holds in the contemporary academic scene thus:

We are trained to accept more readily the potential contiguity, as it were, of elements which are not in fact contiguous. We are undaunted in front of fragmentation, because it is by now our second nature to assume secret kinships (a term dear to semiotics, with a different meaning) among floating particles, whether or not we can detect the bonds that ground the kinships. We have developed an instinctive faith in the unbounded potential for reconstitution into unity of the most disparate elementary particles. And thus we relish the fragments, attracted as we are by the hidden dynamics they seem to display even when we do not know the target towards which that dynamism tends.<sup>340</sup>

The fragments themselves can be "relished", and there is no *a priori* teleology to be deconstructed. There is only an "instinctive faith" that the fragments *are* connected in some way. Their dynamics and connections with other fragments (the prodromes of "narrative sequences") are felt by the scholar to be attractive and interesting, even before the directionality of their dynamics can be discerned. Over the course of a life, so many fragments and connections can be discovered that a million-record website (for instance, [urkesh.org](http://urkesh.org)) is not sufficient to contain them. Such a project creates a fabric that can be woven continuously over a longer period than any one scholar has at his disposition, in a large-scale collaboration, and thus can push the light of understanding much farther into the darkness of the past.

## Redefining the Humanities?

The question that opens at this point is impossibly broad. Yet some answer to the query, "what is the end goal of humanities research?", underlies every position about what it would mean for the humanities to be "digital", and every possible usefulness of the Memex I have dedicated this dissertation to describing and developing. Buccellati states that there is "ample ground for a humanism that can remain true to its core values and be properly digital at the

---

excavations. It is "global" in that it incorporates the entirety of the data as observed during excavation. And it is "argument driven" in that there is an overarching narrative (in fact, several intersecting narratives) that subsume the structured data set." See <http://urkesh.org/main/main8.htm>  
340 <https://urkesh.org/main/main3a.htm>

same time. And my further claim is that archaeology is the best place where such digital humanism can take shape”.<sup>341</sup> His reason is that

to meet the challenge of ‘subjectivity, ambiguity, contingency, observer-dependent variables in the production of knowledge’ it is not sufficient to aim for better tools. These aim ultimately at eliminating "subjectivity" and all the rest, by breaking it all down into smaller and smaller components that are then susceptible of ever more refined quantitative analysis. ... archaeology is ideally suited for a constructive interaction between the social sciences and humanism in the full sense of the word. I argue that there is, on the one hand, a natively digital dimension to archaeology and, on the other, a commitment to deal with broken traditions in such a way as to recapture experience. Thus the notion of "digital humanism" is built into the very essence of archaeology, which can be seen as paradigmatic in illustrating it at its best.<sup>342</sup>

If in the discovery of a complex syntax one is ever to hope to make the leap to semantics, it can only happen through reference to something outside of the system of signs, nodes and relations that defines the underlying database. If humanism is to be understood as something more than quantitative analysis, it must dare to pose the thorny questions of causality and meaning.<sup>343</sup> Humanism aims higher than the quantitative manipulation of data. It entails an interpretive risk.

One important example of the possibility of “recapturing the experience” of an interrupted tradition that I have repeatedly referred to throughout these pages is represented by Buccellati’s studies of Mesopotamian religion. He uses the opposition between polytheism and monotheism to pinpoint the radical difference between these two approaches to the divine, and through that opposition, he distinguishes two approaches to knowledge. I quote one of his formulations of the issue:

(1) The point of departure: polytheism starts from fragmentation that seeks later to recompose unity, while monotheism assumes as its initial datum the unity of the absolute. (2) The goal: for polytheism it is a sum that one presumes can be recomposed, for monotheism it is a personal agent who one presumes can relate to us by their own initiative. (3) The modality of the process: for polytheism we must succeed in possessing the absolute as an ever-less secret object, while for monotheism we must discover how to open ourselves to the mystery.<sup>344</sup>

---

341 Giorgio Buccellati, “Review of Digital Humanities,” (2015) accessed March 6, 2021, <https://critique-of-ar.net/synopses/Burdick&al.htm>.

342 Ibid.

343 Sewell, in *Logics of History. Social Theory and Social Transformation*, University of Chicago Press (2005), points out that the semiotic model he proposes “implies that the social is a complex network, in which language games are nodes of articulation between various overlapping but differently shaped semiotic practices.” p. 356. In his insightful text he distinguishes clearly between mechanistic causality and semiotic explanation, which “applies only to humans”.

These two stances toward the divine in Buccellati's description would seem to also be the description of two activities of intelligence, analysis and synthesis, and point toward deduction and induction.<sup>345</sup> The two stances also point at two large-scale pitfalls. The fragmentary "polytheistic" approach which progresses toward a conclusion, hoping to discover the absolute at the end of a long road of unifying fragments, fails to recognize that an end-point of that progress is strictly speaking impossible. And the "monotheistic" approach which starts with a perception of the whole and uses it to guide the exploration and ordering of the fragments, risks question-begging and cooking the data to fit the theory.<sup>346</sup>

Buccellati calls his summary of different world-views a "structural comparison",<sup>347</sup> that is in a sense the unity of both induction and deduction. The structure is composed of fragments, which in the case of religion are the texts, rituals, myths, and sites that archaeology and history can study. The fragments are composed in a unity that can be seen repeated across many examples, and the unifying characteristic is what he calls the "structure". In the above instance, the underlying structure of polytheism has to do with the summation of fragments; in monotheism the structure has to do with the intuition of the whole.

---

344Giorgio Buccellati, "La Trinita' in Un'ottica Mesopotamica," *Rivista Di Filosofia Neo-Scolastica* 1 (2012): 29–48. Citation p. 36.

345Abduction might be a more precise indicator of the complement of deduction than is induction. See <https://plato.stanford.edu/entries/abduction/index.html> The rhythm of the heart, systole and diastole, is another classic way of putting it. Yet again: from Plato on, many great thinkers and pedagogues have proposed the need to "divide in order to unite".

346Mesopotamian polytheism seeks to arrive at control of destiny through knowledge, or "possession", in Buccellati's language, of the rules of regularity that govern the world [Buccellati, *Quando in alto nei cieli...* Jaca Book (2012), 39]. Something very like this search for regularity and control is a major feature of the scientific endeavor today. We are all Mesopotamian polytheists. The second term of the comparison, monotheism, also exhibits a positive desire at its root. In its definition of the absolute, monotheism affirms the existence of the innumerable, the unique. There is only one God; by extension, on earth, each person is unrepeatable. We also, at least in our universal declarations of human rights, are all monotheists.

347"Comparison" is a key topic for many historians, although often undervalued. For instance, A. A. Van den Braembussche regrets that "critical philosophy of history has devoted barely any attention to the comparative method." A.A. van den Braembussche, "Historical Explanation and Comparative Method: Towards a Theory of the History of Society", *History and Theory* 28(1): p. 10.

While embracing the power of digital humanism, it is important to avoid superficiality.<sup>348</sup> William Sewell points out that “while economics’ use of mathematics has increased its intellectual powers in important ways, its mathematical obsession has also dangerously narrowed its vision, defining as out of bounds the myriad crucial economic problems that do not readily yield up their secrets to mathematical methods”.<sup>349</sup>

If this is true of economics, it is more true of history. Our tools only partially define our workplace.<sup>350</sup> The power of mathematics should not blind us to its limits; the spaces of liminality in the human cultures and traditions we study must be cherished and respected, not hidden under the rug as embarrassing relics of irrationality from a pre-quantitative age. Not all has been discovered, not all has been quantified.

Sewell states that

admitting that social relations are culturally constituted would imply that the Newtonian grid of uniform space and time posited by the quantitative social sciences in fact crumpled and rent—that the world is too messy a place to be understood by a Newtonian social science. And so it is. [... O]ur messy and mutable world needs the conceptual tools that only a collaboration between interpretive social scientists and historians is likely to provide.<sup>351</sup>

The world is too complex to be described by simple laws that are valid in every place: the historian’s search ought instead to be for the conceptual tools that permit description and understanding of complexity as such, without reductionism.<sup>352</sup> Early DH tools had an intrinsic

348One excellent reason to “embrace operationalization”, as the author puts it, is to be found in R. C. Alvarado’s article “Digital Humanities and the Great Project: Why We Should Operationalize Everything—and Study Those Who Are Doing So Now”: “operationalization produces a rationalization effect, a disruption of tacit knowledge caused by the computer’s representational demand for explicit, discrete, and often reductive categories, which frequently requires one to refine held ideas into clear and distinct form.” Alvarado, R. C. “Digital Humanities and the Great Project: Why We Should Operationalize Everything—and Study Those Who Are Doing So Now”. In *Debates in the Digital Humanities* 2019, chapter 6. <https://dhdebates.gc.cuny.edu/read/4805e692-0823-4073-b431-5a684250a82d/section/f113a7d9-5860-4bca-8f85-e2ca4ee87a73> Accessed 10 March 2021.

349William Hamilton Sewell, *Logics of History: Social Theory and Social Transformation*, Reprint, Chicago Studies in the Practices of Meaning (Chicago London: University of Chicago Press, 2011), 354.

350The term “tools” begs some important questions. In particular it may seem to imply the neutrality of the “tools” themselves. This is a theme I have treated elsewhere. See Jonah Lynch, *The Scent of Lemons*, DLT 2012, chapter 6.

351Sewell, *Logics of History*, 17.

352McNeill writes that “The great and obvious difference between natural scientists and historians is the greater complexity of the behavior historians seek to understand. The principal source of historical complexity lies in the fact that human beings react both to the natural world and to one another chiefly through the mediation of symbols. This means, among other things, that any theory about human life, if widely believed, will alter actual behavior, usually by inducing people to act as if the theory were true.” See McNeill, William.

bias toward what Sewell calls the “Newtonian grid”, but new tools can and should be developed that exhibit less of this bias. The Memex is a sustained attempt to do just that: relations define the space of the data, rather than a predetermined grid.

A few times in this dissertation, reference has been made to the work of Johanna Drucker and her attempt to construct an entirely inner-referential representation of time. Some of the limitations of her approach were discussed, especially the fact that the inner-referential definitions had no extra-referentiality to correct them. The inner-referential dimension can be considered as analytical, while the extra-referential dimension corresponds to dialectics. These two dimensions of the Kantian project of establishing the conditions of possibility of a science are described by Buccellati in *A Critique of Archaeological Reason*:

analytics corresponds closely to what I call “grammar”; i.e., a purely inner-referential system. Dialectics, on the other hand, deals with the extra-referential dynamics, whereby elements are braced with concepts that are outside the grammatical system to which they belong. Here the critical approach aims at identifying the potential contradictions (antinomies) and resolving them. In this perspective, dialectics corresponds to hermeneutics: an interpretation that binds together spheres of meaning belonging to different referential systems.<sup>353</sup>

Most of my attention in this thesis has been focused on the first of these components: the analytical, “grammatical” approach to information that has been translated from text into mathematical objects. The reason for this has been stated a few times: my main goal has been to develop methods of managing large quantities of text, while permitting researchers to gain insight into overall patterns that they could not perceive without the use of my method or something like it. I have deployed visual metaphors to help the reader intuit in which ways this method can help a researcher. It is like an x-ray machine that reveals the hidden bone structure of a body; it is like the view from above that an aerial photograph would provide. It is like a special lens that projects high-dimensional information down to a flat page; like a video that permits one to see frame by frame how a complex system changes; and like a microscope that renders visible very small—but crucial—details.

The work of hermeneutics, “binding together spheres of meaning belonging to different referential systems”, is the completion of this project. How can the “grammatical” features of

---

“Mythistory, or Truth, Myth, History, and Historians”, in *The American Historical Review*, Feb. 1986, Vol. 91, No. 1, pp. 1-10. Citation p. 2.

353Giorgio Buccellati, *Critique of Archaeological Reason*, Cambridge University Press (2017), p. 17.

data, as rendered visible, recalled, and visualized in the variety of ways my method affords, be bound together into a coherent interpretation that leads a historian to greater intelligence of his subject?

When I began work on this dissertation, my hope was to achieve an interpretation of the Chicago school of historical research on Mesopotamia through the use of my program. Unfortunately, simply developing the program required far more time than I expected, and producing machine-readable text from the old paper and pdf versions of books and articles that are relevant to my topic also proved very time consuming. The archives that are available in an easily usable digital format, such as the JSTOR dataset I used in the last chapter, unfortunately contain very little material that is directly relevant to Mesopotamian history, which meant that the time-saving potential of my method in this domain was foiled by an *increase* in the time needed to process the text for use in the program. Furthermore, I had to spend a significant amount of time learning to optimize computer processing and produce graphical output, which was fascinating but did not immediately lead to a historical interpretation.

While all these limitations might seem to weaken the overall project, I hope my reader can distinguish between the contingencies of the specific case, and the opportunities my method affords in the general case. As more archives are digitized, as computers become more powerful, and as the overall amount of relevant information inevitably increases, historians will find it both more easy and more necessary to apply methods such as what I present here.

## **Conclusion**

In concluding this chapter, let me quote a passage almost at random from a classic work of literature, *The Once and Future King*, by T. H. White. It was one of my favorite books as a child, a thoughtful and wickedly funny look at King Arthur and the questions of power, law, truth, and goodness. The penultimate chapter opens with a description of a windy day. “Under the doors of the castle the piercing blasts tortured the flapping rushes of the floors. They boo’ed in the tubes of the corkscrew stairs, rattled the wooden shutters, whined shrilly through the shot windows, stirred the cold tapestries in frigid undulations, searched for backbones. The stone towers thrilled under them, trembling bodily like the bass strings of musical instruments. The

slates flew off and shattered themselves with desultory crashes.” Many other passages could have been chosen, since the novel contains many similarly complex and imaginative turns of phrase.

Let us suppose that these words were to be vectorized and compared to other texts, as we have been doing for the last two hundred pages. How would the language model render the literal meaning of the word “blasts”, which is the subject of several sentences? What does it mean that a blast “boo” in the stairs, and why are they tubes? Why are the windows “shot”--because they are leaded, or because they are broken? Like the stone towers, the reader can thrill along with the musical words that swivel and wind about, insinuating themselves into the corners of a cold castle and mimicking the sounds they describe (“flapping rushes”, “rattled shutters”, “desultory crashes”). Onomatopoeia and obliquity. What value would be found in the vectorized version of such a text? Could the computer render any of the poetry? What clusters would form, and what would they mean?

We have arrived at a boundary. We have explored several different ways to examine textual data using machine learning techniques, but the crucial problem of interpretation still remains front and center. Our intellectual telescope, the Memex, brings distant objects closer and like a microscope renders an invisible world visible. In doing so, it continuously opens new fields for discovery. Our curiosity quenched about what previously lay just out of sight, a hundred new questions arise.

Perhaps there is no end to our searching. As historians, we wish to know what happened, when, why how, and for what cause. We build models of cultures and civilizations, we explain and we muse. We conscript other scholars, and dig for clues, fill out libraries and digital databases with our fragments, and construct narratives. After all, “we shall not cease from exploration, and the end of all our exploring will be to arrive where we started, and know the place for the first time” (T.S. Eliot, *Four Quartets*).

# Appendix

This appendix contains some examples of the details of creating the encodings described throughout the dissertation. They are not exhaustive; further examples can be found on my GitHub account.<sup>354</sup>

## Encoding divinities in Cypher

Here is the encoding used for the example of comparing shuila prayers to psalms in Chapter 2:

```
//start inputting shuila prayer info below one prayer at a time
//prayer 393721 //////////////////////////////////////

//create gods
CREATE (enlil:God {name: 'Enlil', source: 'p393721'}),
(fate:God {name: 'Fate'}),

//create persons
(generic:Person {name: 'generic persons'}),

//create events
(eclipse:NaturalObject {name: 'lunar eclipse'}),

//create relations
(generic)-[:FEARS]->(eclipse),
(generic)-[:REQUESTS_LIFE]->(enlil),
(generic)-[:TRUSTS]->(enlil),
(generic)-[:REQUESTS_PROTECTION]->(enlil),
(fate)-[:CAUSES]->(eclipse),

// prayer 259033 //////////////////////////////////////
(sham:Person {name: 'Shamash-shum-ukin', source: 'p259033'}),

(marduk:God {name: 'Marduk', source: 'p259033'}),
(zarpanitu:God {name: 'Zarpanitu', source: 'p259033'}),

(sham)-[:FEARS]->(eclipse),
(sham)-[:REQUESTS_LIFE]->(marduk),
(sham)-[:REQUESTS_HEALTH]->(marduk),
```

---

<sup>354</sup><https://github.com/sibeliu>



// prayer 393771 //////////////////////////////////

(sin:God {name: 'Sin', source: 'p393771'}),  
(anu:God {name: 'Anu', source: 'p393771'}),  
(shamash:God {name: 'Shamash', source: 'p393771'}),  
(ishtar:God {name: 'Ishtar', source: 'p393771'}),  
(nabu:God {name: 'Nabu', source: 'p393771'}),

(generic)-[:OFFERS\_BEER]->(sin),  
(generic)-[:OFFERS\_FLOUR]->(sin),  
(generic)-[:KNEELS\_TO]->(sin),  
(generic)-[:REQUESTS\_JUST\_UTTERANCE]->(sin),  
(generic)-[:REQUESTS\_PEACE]->(sin),  
(generic)-[:REQUESTS\_FORGIVENESS]->(sin),  
(generic)-[:PRAISES]->(sin),  
(generic)-[:FEARS]->(eclipse),

// prayer 393796 //////////////////////////////////

(ninurta:God {name: 'Ninurta', source: 'p393796'}),  
(ashur:God {name: 'Ashur', source: 'p393796'}),  
(ashuritu:God {name: 'Ashuritu', source: 'p393796'}),  
(ash:Person {name: 'Ashurbanipal', source: 'p393796'}),

(ash)-[:OFFERS\_FLOUR]->(ninurta),  
(ash)-[:OFFERS\_BEER]->(ninurta),  
(ash)-[:OFFERS\_INCENSE]->(ninurta),  
(ash)-[:REQUESTS\_PROTECTION]->(ninurta),  
(ash)-[:REQUESTS\_FORGIVENESS]->(ninurta),  
(ash)-[:FEARS]->(eclipse),

// END OF SHUILA PRAYERS INPUT //////////////////////////////////

// NOW INPUT PSALMS

(lord:God {name: ['YHWH', 'Lord', 'God']}),  
(psalmist:Person {name: 'Psalmist'}),

// psalm 21

(king:Person {name: 'King', source: 'ps. 21'}),  
(enemies:Person {name: 'Enemies', source: 'ps.21'}),  
(lord)-[:GIVES\_LIFE {source: 'ps. 21'}]->(king),  
(lord)-[:ASKS\_LIFE {source: 'ps. 21'}]-(king),  
(king)-[:TRUSTS {source: 'ps.21'}]->(lord),  
(lord)-[:DESTROYS {source: 'ps. 21'}]->(enemies),

// psalm 33

```

(righteous:Person {name: 'Righteous'}),
(heavens:NaturalObject {name: 'Heavens'}),
(nations:Person {name: 'Nations'}),

(righteous)-[:SHOUT_FOR_JOY {source: 'ps. 33'}]->(lord),
(righteous)-[:PRAISES {source: 'ps. 33'}]->(lord),
(righteous)-[:THANKS {source: 'ps. 33'}]->(lord),
(righteous)-[:SING {source: 'ps. 33'}]->(lord),
(lord)-[:CREATE {source: 'ps. 33'}]->(heavens),
(lord)-[:FRUSTRATES_PLANS {source: 'ps. 33'}]->(nations),
(lord)-[:OBSERVES {source: 'ps. 33'}]->(nations),
(psalmist)-[:WAITS_ON {source: 'ps.33'}]->(lord),
(psalmist)-[:TRUSTS {source: 'ps.33'}]->(lord),
(psalmist)-[:HOPEES {source: 'ps.33'}]->(lord),

```

```
// psalm 71
```

```

(psalmist)-[:TAKES_REFUGE {source: 'ps.71'}]->(lord),
(psalmist)-[:REQUESTS_PROTECTION {source: 'ps.71'}]->(lord),
(psalmist)-[:REQUESTS_LIFE {source: 'ps.71'}]->(lord),
(psalmist)-[:PRAISES {source: 'ps.71'}]->(lord),
(enemies)-[:SLANDER {source: 'ps.71'}]->(psalmist),
(lord)-[:SHAMES {source: 'ps.71'}]->(enemies),
(lord)-[:TEACHES {source: 'ps.71'}]->(psalmist),
(lord)-[:REVIVES {source: 'ps.71'}]->(psalmist),

```

```
// psalm 96
```

```

(world:NaturalObject {name: 'World'}),
(lord)-[:JUDGES {source: 'ps 96'}]->(world),
(lord)-[:JUDGES {source: 'ps 96'}]->(nations)

```

## How are relations encoded?

In order to create the PRECEDES relations, I wrote a script<sup>355</sup> that receives a block of text, parses it into sentences, and then saves the sentences in a data structure that encodes the order of sentences as a relation between nodes. Sentence parsing was handled by the SpaCy package described in chapter 2, which is accurate, but can make mistakes when parsing sentences that contain many special characters, abbreviations with periods, and the like. In order to prepare the text, I first wrote a cleaning routine called `clean_text`. Here is the code:

---

<sup>355</sup>The main logic of the program is written in Python. Interaction with the database, which is a Neo4j graph database, is handled in another language, called Cypher. Display to the screen is currently handled with HTML and Javascript in a web browser. The first iteration of the user interface used the Python package Tkinter.

```

print ("cleaning text")
text = "\n".join([ll.rstrip() for ll in text.splitlines() if ll.strip()])
text = " ".join(text.split())
text = text.replace('[', '(')
text = text.replace("]", ")")
text = text.replace('"', "'")
return text

```

This routine first removes empty lines, then removes whitespace characters and replaces them with a simple space, and finally replaces square brackets and double quotes with parentheses and single quotes. Brackets cause problems later on because they look to the computer like programming commands, and not simple text. Once the text is cleaned up, it is ready for processing. The full code is in the appendix; the relevant section is below. First, the cleaned text is processed with SpaCy to encode part of speech and word vector data for the text as a new object named `nlp_text`:

```

text = clean_text(original_text)
nlp_text = nlp(text)

```

Then, document codes are established in order to be able to easily find the complete document to which all its sentences belong. In order to establish a unique code for each one, we simply count the nodes currently in the system, and use that number as the `doc_code`:

```

#make a unique document code: the current count of nodes is the first available node
cypher_query = """MATCH (n) RETURN COUNT (n)"""
doc_code = graph.run(cypher_query).data()
doc_code = doc_code[0]['COUNT (n)']

```

We also create a node that represents the entire document, for future use. Although whole documents are not connected directly, being able to reference them without having to calculate the path between all the sentences, and thereby reconstruct the original text, will lead to

performance improvements down the line. The document node contains the full text of the paragraph(s) that were saved by the user and can be recalled simply by knowing the document code, which is saved in each of the sentences it is composed of.

Then, we loop through the sentences in the document and create one node for each sentence:

```
for index, value in enumerate(nlp_text.sents):
```

```
temp_name = "s"+str(index)
```

```
content = str(value)
```

```
cypher_query += 'CREATE ('+temp_name+':Sentence{author_ln: "'+author_ln+'", author_fn: "'+author_fn+'", title: "'+title+'", file: "'+filename+'", saved_by: "'+user_name+'", doc_code: "'+str(doc_code)+'", content: "'+content+''}) \n'
```

Each sentence is encoded along with the author name, the title, its document code, and the name of the person who saved this document. Other information could also be saved, such as the time and place in which it was saved. Such chronological and geographical data could be useful later on.

Finally, sentences are connected with PRECEDES relations in order to encode the structure of the paragraph within the graph itself:

```
for index, value in enumerate(nlp_text.sents):
```

```
if index > 0: #do nothing for the first node
```

```
temp_name = "s" + str(index)
```

```
temp_name2 = "s" + str(index-1)
```

```
cypher_query += """
```

```
MERGE ("""+temp_name2+""")-[:PRECEDES]->("""+temp_name+""") \n
```

```
"""
```

The Cypher query is concatenated until all sentences have been processed, and is then run through the Neo4j database to establish the linear order of sentences as PRECEDES relations.

## Logical relations

Logical connections can also be created by the user. These relations offer an alternative to the relations that can be automatically generated by word embeddings and similarity calculations, much as in the previous chapter we used two alternate methods of studying Mesopotamian religion. The first method used computers to automatically characterize text; the second method used the computer to visually represent relations that had been encoded manually by the historian.

The code to do this is quite straightforward. Our routine for creating an AGREES\_WITH relation is as follows. The only noteworthy detail is that we first check to make sure that the two phrases that agree are actually distinct phrases: it would be possible for a user to make a mistake and set a phrase in agreement with itself, which would be tautological and might lead to errors downstream:

```
def agree(graph, source_id, target_id, username):

    if source_id != target_id:
        cypher_query = """
MATCH (t), (s) \n
WHERE ID(t)="\" + str(target_id) + "\"" AND ID(s)="\" + str(source_id) + "\"" \n
MERGE (s)-[:AGREES_WITH {username: '\"\" + str(username) + '\"\"}]->(t) \n
"""

        graph.run(cypher_query)
    return
```

The same procedure may be used to create a DISAGREES\_WITH relation. A feature of the system as described thus far is that while PRECEDES relations are automatically created from the text, logical connections must be created by the user. The relevant routines are called when the user himself decides that two phrases agree or disagree, and not before.



## Bibliography

- Adams, Frederick Ray, and Kenneth Aizawa. *The Bounds of Cognition*. Malden, Ma: Blackwell Publishing, 2008.
- Alstola, Tero, Shana Zaia, Aleksi Sahala, Heidi Jauhiainen, Saana Svärd, and Krister Lindén. “Aššur and His Friends: A Statistical Analysis of Neo-Assyrian Texts.” *Journal of Cuneiform Studies* 71 (January 2019): 159–80. <https://doi.org/10.1086/703859>.
- Andor, Jozsef. “Cognitive Grammar.” *Acta Linguistica Hungarica* 52, no. 4 (2005): 341–66.
- Andreski, Stanislav. *Social Sciences as Sorcery*. 3. impr. London: Deutsch, 1974.
- Armstrong, Robert L. “The Midpoint on a Five-Point Likert-Type Scale.” *Perceptual and Motor Skills* 64, no. 2 (April 1987): 359–62. <https://doi.org/10.2466/pms.1987.64.2.359>.
- Balter, Michael. “Science.” *World’s Oldest Stone Tools Discovered in Kenya*, April 14, 2015. <https://www.science.org/content/article/world-s-oldest-stone-tools-discovered-kenya>.
- Banks, David. “The Extent to Which the Passive Voice Is Used in the Scientific Journal Article, 1985–2015.” *Functional Linguistics* 4, no. 1 (December 2017): 12. <https://doi.org/10.1186/s40554-017-0045-5>.
- Barfield, Owen, and Wystan H. Auden. *History in English Words*. Gt. Barrington, MA: Lindisfarne Books, 2002.
- Berners-Lee, Tim, Dan Connolly, Lalana Kagal, Yosi Scharf, and Jim Hendler. “N3Logic: A Logical Framework for the World Wide Web.” *Theory and Practice of Logic Programming* 8, no. 3 (May 2008): 249–69. <https://doi.org/10.1017/S1471068407003213>.
- Bezaroova, Galyna. “The Problem of the Subject and Object in the Methodology of History.” *Current Issues of Social Sciences and History of Medicine*, no. 2 (February 23, 2017): 62–67. <https://doi.org/10.24061/2411-6181.2.2017.41>.
- Bickhard, Mark H. “What Could Cognition Be If Not computation...Or Connectionism, or Dynamic Systems?” *Journal of Theoretical and Philosophical Psychology* 35, no. 1 (February 2015): 53–66. <https://doi.org/10.1037/a0038059>.
- Borghini, Anna M., Laura Barca, Ferdinand Binkofski, and Luca Tummolini. “Varieties of Abstract Concepts: Development, Use and Representation in the Brain.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 373, no. 1752 (August 5, 2018): 20170121. <https://doi.org/10.1098/rstb.2017.0121>.
- Börner, Katy. *Atlas of Science: Visualizing What We Know*. Cambridge, Mass: MIT press, 2010.
- Brann, Eva. “Plato’s Theory of Ideas.” *The St. John’s Review* 32, no. 1 (n.d.): 29–37.
- Brückner, Martin, Sandy Isenstadt, and Sarah Wasserman, eds. *Modelwork: The Material Culture of Making and Knowing*. Minneapolis London: University of Minnesota Press, 2021.
- Buccellati, Giorgio. *Critique of Archaeological Reason*. Cambridge University Press, 2017.
- . *Dal profondo del tempo: all’origine della comunicazione e della comunità nell’antica Siria*. Firenze: Società editrice fiorentina, 2014.
- . “La Trinita’ in Un’ottica Mesopotamica.” *Rivista Di Filosofia Neo-Scolastica* 1 (2012): 29–48.
- . *Quando in Alto i Cieli...* Jaca Book, 2012.
- . “Review of Digital Humanities.” Accessed March 6, 2021. <https://critique-of-ar.net/synopses/Burdick&al.htm>.
- . “The Semiotics of Ethnicity: The Case of Hurrian Urkesh,” n.d., 20.

- . “*When on High the Heavens...*”: *Mesopotamian Religion and Spirituality with Reference to the Biblical World*. Translated by Jonah Lynch. London New York: Routledge, 2024.
- Burdick, Anne, ed. *Digital Humanities*. Cambridge, MA: MIT Press, 2012.
- Casati, Roberto, and Achille C. Varzi. *Holes and Other Superficialities*. Cambridge, Mass: MIT Press, 1994.
- Caucheteux, Charlotte, and Jean-Rémi King. “Brains and Algorithms Partially Converge in Natural Language Processing.” *Communications Biology* 5, no. 1 (February 16, 2022): 134. <https://doi.org/10.1038/s42003-022-03036-1>.
- Chomsky, Noam. *Syntactic Structures*. Mouton, 1957.
- Churchland, Patricia Smith, and Terrence J. Sejnowski. “Neural Representation and Neural Computation.” *Philosophical Perspectives* 4 (1990): 343. <https://doi.org/10.2307/2214198>.
- Clark, A., and D. Chalmers. “The Extended Mind.” *Analysis* 58, no. 1 (January 1, 1998): 7–19. <https://doi.org/10.1093/analys/58.1.7>.
- Collar, Anna, Fiona Coward, Tom Brughmans, and Barbara J. Mills. “Networks in Archaeology: Phenomena, Abstraction, Representation.” *Journal of Archaeological Method and Theory* 22, no. 1 (March 2015): 1–32. <https://doi.org/10.1007/s10816-014-9235-6>.
- De Bolla, Peter, ed. *Explorations in the Digital History of Ideas: New Methods and Computational Approaches*. Cambridge New York: Cambridge University Press, 2024. <https://doi.org/10.1017/9781009263610>.
- Dempsey, Corinne G. *Bringing the Sacred down to Earth: Adventures in Comparative Religion*. New York: Oxford University Press, 2012.
- Eliot, T.S. *The Four Quartets*, n.d.
- Eve, Martin Paul. “Distance and History.” In *The Digital Humanities and Literary Studies*, by Martin Paul Eve, 129–50, 1st ed. Oxford University Press Oxford, 2022. <https://doi.org/10.1093/oso/9780198850489.003.0005>.
- Fernandino, Leonardo, Jia-Qing Tong, Lisa L. Conant, Colin J. Humphries, and Jeffrey R. Binder. “Decoding the Information Structure Underlying the Neural Representation of Concepts.” *Proceedings of the National Academy of Sciences* 119, no. 6 (February 8, 2022): e2108091119. <https://doi.org/10.1073/pnas.2108091119>.
- Franzosi, Roberto. “A Third Road to the Past? Historical Scholarship in the Age of Big Data.” *Historical Methods: A Journal of Quantitative and Interdisciplinary History* 50, no. 4 (October 2, 2017): 227–44. <https://doi.org/10.1080/01615440.2017.1361879>.
- Furet, Francois. “Quantitive History.” *Daedalus: Journal of the American Academy of Arts and Sciences* 100 (1971): 151–67.
- Glanzberg, Michael. “Truth.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2021. Metaphysics Research Lab, Stanford University, 2021. <https://plato.stanford.edu/archives/sum2021/entries/truth/>.
- González-Márquez, Rita, Luca Schmidt, Benjamin M. Schmidt, Philipp Berens, and Dmitry Kobak. “The Landscape of Biomedical Research.” Preprint. Scientific Communication and Education, April 11, 2023. <https://doi.org/10.1101/2023.04.10.536208>.
- Grafton, Anthony. “The Footnote from De Thou to Ranke.” *History and Theory* 33, no. 4 (December 1994): 53. <https://doi.org/10.2307/2505502>.
- Griffin, Gabriele, and Matt Steven Hayler. “Collaboration in Digital Humanities Research – Persisting Silences,” n.d.
- Hale, Kenneth L., and Samuel Jay Keyser. *Prolegomenon to a Theory of Argument Structure*. Linguistic Inquiry Monographs 39. Cambridge, Mass: MIT Press, 2002.



- Halliday, M. A. K., and Christian M. I. M. Matthiessen. *Halliday's Introduction to Functional Grammar*. Fourth Edition. Milton Park, Abingdon, Oxon: Routledge, 2014.
- Hofstadter, Douglas R. *I Am a Strange Loop*. New York, NY: Basic Books, 2007.
- Hosseini, Kasra, Kaspar Beelen, Giovanni Colavizza, and Mariona Coll Ardanuy. "Neural Language Models for Nineteenth-Century English." arXiv, May 24, 2021. <http://arxiv.org/abs/2105.11321>.
- Hutchins, Edwin. *Cognition in the Wild*. 8. pr. A Bradford Book. Cambridge, Mass.: MIT Press, 2006.
- Jacobsen, Thorkild. *Formative Tendencies in Sumerian Religion*. Doubleday, 1961.
- . *The Treasures of Darkness: A History of Mesopotamian Religion*. New Haven London: Yale University Press, 1976.
- . "Toward the Image of Tammuz." *History of Religions* 1, no. 2 (n.d.).
- Kaiser, Daniel, Arthur M. Jacobs, and Radoslaw M. Cichy. "Modelling Brain Representations of Abstract Concepts." Edited by Leyla Isik. *PLOS Computational Biology* 18, no. 2 (February 4, 2022): e1009837. <https://doi.org/10.1371/journal.pcbi.1009837>.
- Kim, Dorothy. *Disrupting the Digital Humanities*. 1st edition. Santa Barbara, CA: Punctum Books, 2018.
- Lenat, Doug, and Gary Marcus. "Getting from Generative AI to Trustworthy AI: What LLMs Might Learn from Cyc." arXiv, July 31, 2023. <http://arxiv.org/abs/2308.04445>.
- Lenneberg, Eric H. *Biological Foundations of Language*. John Wiley and Sons, 1967.
- Liang, Yuding, and Kenny Zhu. "Automatic Generation of Text Descriptive Comments for Code Blocks." *Proceedings of the AAI Conference on Artificial Intelligence* 32, no. 1 (April 27, 2018). <https://doi.org/10.1609/aaai.v32i1.11963>.
- Lordkipanidze, David, Abesalom Vekua, Reid Ferring, G. Philip Rightmire, Christoph P.E. Zollikofer, Marcia S. Ponce de León, Jordi Agustí, et al. "A Fourth Hominin Skull from Dmanisi, Georgia." *The Anatomical Record Part A: Discoveries in Molecular, Cellular, and Evolutionary Biology* 288A, no. 11 (November 2006): 1146–57. <https://doi.org/10.1002/ar.a.20379>.
- McCarty, Willard. "Making and Studying Notes: Towards a Cognitive Ecology of Annotation." In *Annotating Scholarly Editions and Research: Functions, Differentiation, Systematization*, 271–97. Berlin: Walter de Gruyter, 2020.
- McCullagh, C. Behan. "Historical Realism." *Philosophy and Phenomenological Research* 40, no. 3 (March 1980): 420. <https://doi.org/10.2307/2106406>.
- . "Language and the Truth of History." *History and Theory* 44, no. 3 (October 2005): 441–55. <https://doi.org/10.1111/j.1468-2303.2005.00335.x>.
- McInnes, Leland, John Healy, and James Melville. "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction." arXiv, September 17, 2020. <http://arxiv.org/abs/1802.03426>.
- McNeill, John Robert, and William Hardy McNeill. *The Human Web: A Bird's-Eye View of World History*. 1. ed. New York: Norton, 2003.
- McNeill, William. "Mythistory, or Truth, Myth, History, and Historians." *The American Historical Review*, February 1986. <https://doi.org/10.1086/ahr/91.1.1>.
- Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. "Efficient Estimation of Word Representations in Vector Space." arXiv, September 6, 2013. <http://arxiv.org/abs/1301.3781>.
- Millet, Juliette, Charlotte Caucheteux, Pierre Orhan, Yves Boubenec, Alexandre Gramfort, Ewan Dunbar, Christophe Pallier, and Jean-Remi King. "Toward a Realistic Model of Speech Processing in the Brain with Self-Supervised Learning." arXiv, March 20, 2023. <http://arxiv.org/abs/2206.01685>.

- the Meanings of Nouns.” *Science* 320, no. 5880 (May 30, 2008): 1191–95.  
<https://doi.org/10.1126/science.1152876>.
- Moreton, Simon. “Rethinking ‘Knowledge Exchange’: New Approaches to Collaborative Work in the Arts and Humanities.” *International Journal of Cultural Policy* 22, no. 1 (January 1, 2016): 100–115. <https://doi.org/10.1080/10286632.2015.1101081>.
- Moro, Andrea. *The Boundaries of Babel*. 2nd ed. MIT Press, 2015.
- Nefdt, Ryan M. “Infinity and the Foundations of Linguistics.” *Synthese* 196, no. 5 (May 2019): 1671–1711. <https://doi.org/10.1007/s11229-017-1574-x>.
- Nieto, Nicolás, Victoria Peterson, Hugo Leonardo Rufiner, Juan Esteban Kamienkowski, and Ruben Spies. “Thinking out Loud, an Open-Access EEG-Based BCI Dataset for Inner Speech Recognition.” *Scientific Data* 9, no. 1 (February 14, 2022): 52. <https://doi.org/10.1038/s41597-022-01147-2>.
- Oppenheim, Adolf Leo, Robert H. Brill, Dan Barag, Axel von Saldern, and Corning Museum of Glass, eds. *Glass and Glassmaking in Ancient Mesopotamia: An Edition of the Cuneiform Texts with Contain Instructions for Glassmakers with a Catalogue of Surviving Objects*. Repr. The Corning Museum of Glass Monographs 3. Corning: Corning Museum of Glass Pr. [u.a.], 1988.
- Oppenheim, Adolf Leo, and Erica Reiner. *Ancient Mesopotamia: Portrait of a Dead Civilization*. Rev. ed., 11. impr. Chicago: Univ. of Chicago Press, 1998.
- Paccanaro, Alberto, and Geoffrey E. Hinton. “Learning Distributed Representations of Relational Data Using Linear Relational Embedding.” In *Neural Nets WIRN Vietri-01*, edited by Roberto Tagliaferri and Maria Marinaro, 134–43. Perspectives in Neural Computing. London: Springer London, 2002. [https://doi.org/10.1007/978-1-4471-0219-9\\_12](https://doi.org/10.1007/978-1-4471-0219-9_12).
- Peterson, Kevin J., and Hongfang Liu. “An Examination of the Statistical Laws of Semantic Change in Clinical Notes.” *AMIA Joint Summits on Translational Science Proceedings. AMIA Joint Summits on Translational Science 2021* (2021): 515–24.
- Philcox, Oliver H. E. “Probing Parity-Violation with the Four-Point Correlation Function of BOSS Galaxies,” 2022. <https://doi.org/10.48550/ARXIV.2206.04227>.
- Piantadosi, Steven T. “Modern Language Models Refute Chomsky’s Approach to Language,” n.d.
- Poole, Fitz John Porter. “Metaphors and Maps: Toward Comparison in the Anthropology of Religion.” *Journal of the American Academy of Religion* LIV, no. 3 (1986): 411–98.  
<https://doi.org/10.1093/jaarel/LIV.3.411>.
- Portenoy, Jason, Jessica Hullman, and Jevin D. West. “Leveraging Citation Networks to Visualize Scholarly Influence Over Time.” *Frontiers in Research Metrics and Analytics* 2 (November 27, 2017): 8. <https://doi.org/10.3389/frma.2017.00008>.
- Qiu, Wenjun, and Yang Xu. “HistBERT: A Pre-Trained Language Model for Diachronic Lexical Semantic Analysis,” 2022. <https://doi.org/10.13140/RG.2.2.14905.44649>.
- Rheingold, Howard. *Tools for Thought: The History and Future of Mind-Expanding Technology*. 1st MIT Press ed. Cambridge, Mass: MIT Press, 2000.
- Robertson, Stephen. “Understanding Inverse Document Frequency: On Theoretical Arguments for IDF.” *Journal of Documentation* 60, no. 5 (October 2004): 503–20.  
<https://doi.org/10.1108/00220410410560582>.
- Rodriguez-Pereyra, Gonzalo. “Nominalism in Metaphysics.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2019. Metaphysics Research Lab, Stanford University, 2019. <https://plato.stanford.edu/archives/sum2019/entries/nominalism-metaphysics/>.
- Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. “Learning Representations by Back-Propagating Errors.” *Nature* 323, no. 6088 (October 1986): 533–36.  
<https://doi.org/10.1038/323533a0>.

- Schrimpf, Martin, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A. Hosseini, Nancy Kanwisher, Joshua B. Tenenbaum, and Evelina Fedorenko. "The Neural Architecture of Language: Integrative Modeling Converges on Predictive Processing." *Proceedings of the National Academy of Sciences* 118, no. 45 (November 9, 2021): e2105646118. <https://doi.org/10.1073/pnas.2105646118>.
- Sewell, William Hamilton. *Logics of History: Social Theory and Social Transformation*. Reprint. Chicago Studies in the Practices of Meaning. Chicago London: University of Chicago Press, 2011.
- Smith, Mark S. *The Origins of Biblical Monotheism: Israel's Polytheistic Background and the Ugaritic Texts*. 1. issued as an Oxford Univ. Press paperback. Oxford: Oxford Univ. Press, 2003.
- Stephenson, Neal. "Communication Prosthetics: Threat, or Menace?" *Whole Earth*, Summer 2001.
- Tahmasebi, Nina, Lars Borin, and Adam Jatowt. "Survey of Computational Approaches to Lexical Semantic Change." arXiv, March 13, 2019. <http://arxiv.org/abs/1811.06278>.
- Takagi, Yu, and Shinji Nishimoto. "High-Resolution Image Reconstruction with Latent Diffusion Models from Human Brain Activity." Preprint. Neuroscience, November 21, 2022. <https://doi.org/10.1101/2022.11.18.517004>.
- Wang, Yixiao, Zied Bouraoui, Luis Espinosa Anke, and Steven Schockaert. "Deriving Word Vectors from Contextualized Language Models Using Topic-Aware Mention Selection." arXiv, June 15, 2021. <http://arxiv.org/abs/2106.07947>.
- Wattenberg, Martin, Fernanda Viégas, and Ian Johnson. "How to Use T-SNE Effectively." *Distill* 1, no. 10 (October 13, 2016): 10.23915/distill.00002. <https://doi.org/10.23915/distill.00002>.
- Zampieri, Marcos, Shervin Malmasi, and Mark Dras. "Modeling Language Change in Historical Corpora: The Case of Portuguese." arXiv, September 30, 2016. <http://arxiv.org/abs/1610.00030>.
- Zhu, Wanzheng, and Suma Bhat. "Euphemistic Phrase Detection by Masked Language Model." arXiv, September 10, 2021. <http://arxiv.org/abs/2109.04666>.