University of Pavia

Ph.D. School in Electronics, Computer Science and Electrical Engineering

Cycle XXXII

Learning and Optimization Methods for the Efficient Management of HVAC Systems

Author:
Federica Acerbi

Supervisor:
Prof. Giuseppe De Nicolao

2016-2019

UNIVERSITY OF PAVIA

Abstract

Faculty of Engineering
Ph.D. School in Electronics, Computer Science and Electrical Engineering

Doctor of Philosophy

Learning and Optimization Methods for the Efficient Management of HVAC Systems

by Federica Acerbi

The thesis addresses the reduction of electricity demand in existing buildings by two different standpoints: the optimal management of multi-chiller systems and the development of thermal comfort models usable for the optimal calibration of indoor air temperature set-points. The former problem, involves both the identification of chiller efficiency and power consumption models and the design of optimized management strategies that minimize consumption while complying with operating constraints and the satisfaction of the cooling demand. In spite of the extensive literature devoted to chiller modeling, the advent of continuous monitoring, while opening the way to new machine learning methodologies, raises also new challenges in terms of model robustness. The thesis compares four existing models and proposes a new Gaussian Process approach, assessing their robustness on extensive field data collected over a 2-year period in an HVAC systems serving a large semiconductor plant. Concerning the optimal management of multi-chiller systems, in alternative to the variety of heuristic methods available in the literature, the thesis develops a new exact solution, that, by properly partitioning the solution space, decomposes the Optimal Chiller Problem into elementary Quadratic Programs, each of which is subject to only one inequality constraint. The exact solver is validated against the field data collected in the semiconductor plant, not only demonstrating its numerical efficiency, but also quantitatively assessing the potential energy saving achievable by its use as the core of an Optimal Chiller Sequencing strategy. Savings on the demand side can come also from the development of advanced thermal comfort models. In the thesis, statistical and machine learning methods are used to obtain a neat validation of the so called adaptive theory, according to which the outdoor temperature influences the neutral temperature regarded comfortable by the occupants. In this context, a multilogistic regression approach is first used to rigorously model the thermal comfort scores of the ASHRAE RP-884 dataset and a new definition of neutral temperature is proposed that, differently from the current one, explicitly accounts for the percentage of satisfied occupants.

Contents

Abstract iii

Contents iv

List of Figures vii

List of Tables xi

1 Introduction 1
- 1.1 Thesis Overview . . . 4
- 1.2 Collaborations and related pubblications . . . 7

2 HVAC Systems: basic notions 9
- 2.1 Water-cooled HVAC system . . . 9
- 2.2 Vapor-compression chiller . . . 10
- 2.3 Chiller efficiency . . . 12
- 2.4 Plant description . . . 12
  - 2.4.1 Dataset . . . 13

3 Learning chillers efficiency from data 15
- 3.1 Energy performance models . . . 16
  - 3.1.1 Gordon-Ng universal model . . . 17
  - 3.1.2 Bi-quadratic regression model . . . 18
  - 3.1.3 Multivariate polynomial regression model . . . 18
  - 3.1.4 Multilayer perceptron model . . . 18
  - 3.1.5 Gaussian process model . . . 19
- 3.2 The Covariate Shift Problem . . . 20
- 3.3 Field data . . . 22
- 3.4 Models comparison . . . 24
  - 3.4.1 Model comparison: global dataset . . . 24
  - 3.4.2 Models Robustness: covariate shift experiment . . . 25
- 3.5 Discussion . . . 29

4 Calibration strategies for the Gordon-Ng chiller model 31
- 4.1 The GNU model of chiller perfomance . . . 32
- 4.2 Calibrating the GNU model . . . 33

4.2.1 OLS estimation . . . 33
4.2.2 EIV estimation . . . 34
4.2.3 NLS estimation . . . 35
4.3 Data and methods . . . 36
4.3.1 Case Study 1 . . . 36
4.3.2 Case Study 2 . . . 38
4.4 Results . . . 39
4.4.1 Case Study 1 . . . 39
4.4.2 Case Study 2 . . . 44
4.5 Concluding remarks . . . 46

5 Chiller Plant Optimization 47
5.1 The Optimal Chiller Loading problem . . . 49
5.1.1 Quadratic power consumption model . . . 51
5.1.2 Partition of the solution space . . . 52
5.1.3 Divide and conquer strategy . . . 53
5.1.4 Reduction to equality-constrained problems . . . 55
5.1.5 Summary of the X-OCL Algorithm . . . 57
5.2 Test on Hsinchu benchmark model . . . 58
5.2.1 Hsinchu cooling plant model . . . 58
5.2.2 OCL benchmark: results . . . 59
5.3 The Optimal Chiller Sequencing problem . . . 63
5.4 A lower bound to the OCS problem . . . 63
5.5 X-OCS, a greedy approach to OCS . . . 64
5.5.1 OCS Simulated example . . . 65
5.5.2 OCS benchmark: results . . . 67
5.6 Validation on field data . . . 70
5.6.1 Chiller energy consumption models . . . 71
5.6.2 Real HVAC system: assessment of potential savings . . . 75
5.7 Execution time . . . 77
5.8 Discussion . . . 78

6 Thermal Comfort in Air-Conditioned Buildings 81
6.1 Neutral Temperature definition and estimation . . . 82
6.2 Thermal comfort survey data . . . 83
6.3 Exploratory analysis of thermal comfort data . . . 84
6.4 Direct Modeling of Thermal Comfort . . . 87
6.4.1 Polynomial regression model . . . 87
6.4.2 Multinomial logistic model . . . 88
6.5 Alternative definition of Neutral Temperature . . . 89
6.6 Comparison and discussion . . . 91

7 Conclusions 95

A Derivation of power consumption models from the COP ones 97

List of Figures

2.1 Water-cooled HVAC system scheme 10
2.2 Left: basic refrigeration cycle. Right: refrigerant circuit; P-H Diagram. 11
2.3 Primary circuit layout 13

3.1 Regression under covariate shift. Left: predictor trained via OLS on exponentially distributed data. Right: predictor trained on uniformly distributed data 22
3.2 Overall six-month dataset. Red: training data; Blue: testing data; green: kernel estimate of the joint distribution of the input covariates $Q_e(i)$ and $T_{ci}(i)$. 23
3.3 Monthly datasets. Red: training data; Blue: test data; green: kernel density estimate of the training distribution of the input covariates $Q_e(i)$ and $T_{ci}(i)$. . . . 23
3.4 Gordon-Ng Universal model under covariate shift. Surface: prediction model; crosses: test data. 26
3.5 Bi-quadratic model under covariate shift. Surface: prediction model; crosses: test data. 27
3.6 Multivariate Polynomial regression model under covariate shift. Surface: prediction model; crosses: test data. 27
3.7 Multilayer Neural Network model under covariate shift. Surface: prediction model; crosses: test data. 28
3.8 Gaussian Processes model under covariate shift. Surface: prediction model; crosses: test data. 28
3.9 The average RMSE values of the five models 29

4.1 Data set #1: three-dimensional representation of the COP as a function of measured chiller data. 36
4.2 Data set #2: three-dimensional representation of the COP as a function of measured chiller data. Blue dots: training data; Red dots: test data. 38
4.3 Comparison of OLS, EIV and NLS estimates of β_1 varying the noise magnitude of the measurements 40
4.4 Comparison of OLS, EIV and NLS estimates of β_2 varying the noise magnitude of the measurements 41
4.5 Comparison of OLS, EIV and NLS estimates of β_3 varying the noise magnitude of the measurements 42
4.6 RMSE% of the GNU model parameters estimated by OLS, EIV and NLS methods. From top to bottom: β_1, β_2, β_3. 43
4.7 Comparison of the medians of $RMSE^{OLS}$, $RMSE^{EIV}$ and $RMSE^{NLS}$ on test data 43
4.8 $RMSE^{OLS}$, $RMSE^{EIV}$ and $RMSE^{NLS}$ on test data. Solid line: median; Dash-dot line: variability limits 44

4.9 $RMSE^{OLS}$ and $RMSE^{NLS}$ probability density functions on training (top) and test data (bottom) . . . 45
4.10 Box plots of $RMSE^{OLS}$ and $RMSE^{NLS}$ on test data . . . 45

5.1 Hsinchu benchmark: P-PLR curves at $T = 24.5°C$ . . . 60
5.2 Estimated chillers' P-PLR curves on case study 1 (left) and case study 2 (right). 66
5.3 Cooling load demand profile for case study 1 (left) and case study 2 (right). . . . 67
5.4 Cooling load demand time series . . . 70
5.5 Three-dimensional representation of power consumption models. Red dots represent the experimental data. . . . 72
5.6 Goodness-of-Fit (GOF) plots of the quadratic power consumption models. Blue dots are the validation data. . . . 73
5.7 3-D histogram of the pairs (Q_{load}, T) . . . 74
5.8 Condenser inlet water temperature time series . . . 74
5.9 P-PLR curves of Field data benchmark at $T = 21.5\ °C$. . . . 75
5.10 Weekly power consumption time serie: comparison bewteen actual recorded data (red), the lower bound R-OCL (black) and the consumption associated with X-OCS (blue). . . . 76
5.11 Difference between the actual recorded power consumption and the one associated with X-OCS (hourly data) . . . 76
5.12 Average computation time per load . . . 78

6.1 Overall summer dataset. Coloured dots: occupants' thermal sensation votes on the ASHRAE scale [-3; +3] as a function of T_{in} and T_{out}. . . . 85
6.2 Left: blue dots represent the selected data corresponding to ASH = 0 while the red ones the remaining. Right: the blue dots, displayed in the three dimensional space in the left panel, are here shown in the plane Outdoor-Indoor air temperature. 85
6.3 Blue dots: data points (T_{out}, T_{in}) corresponding to $ASH = 0$; Black line: simple linear regression model; Blue line:Lowess regression; Red line: MLP Neural Network model. . . . 86
6.4 Polynomial model: Left, occupants' expected mean vote $\mathcal{E}(ASH \mid T_{out}, T_{in})$. Red curve: intersection of the surface with $\mathcal{E}(ASH) = 0$. When projected onto the plane (T_{out}, T_{in}), the red curve provides the neutral temperature model $\hat{T}_N = g(T_{out})$. Right, the estimated neutral temperature curve in the plane outdoor air temperature - indoor air temperature . . . 88
6.5 Left: Overall summer dataset. Right: Logistic cumulative distribution functions . 89
6.6 Ordered logit model. Left panel: occupants' expected mean vote $\mathcal{E}(ASH \mid T_{out}, T_{in})$. Red curve: intersection of the surface with $\mathcal{E}(ASH) = 0$. When projected onto the plane (T_{out}, T_{in}), the red curve provides the neutral temperature model $\hat{T}_N = g(T_{out})$. Right panel: the estimated neutral temperature curve in the plane outdoor air temperature - indoor air temperature . . . 90
6.7 Ordinal logit model. Surface: $\Pr(ASH = 0 \mid T_{out}, T_{in})$; The projection of the red curve onto the plane (T_{out}, T_{in}) represents the neutral temperature model $\hat{T}^*_N = g(T_{out})$. . . . 91
6.8 Binomial logit model. Surface: $\Pr(ASH = 0 \mid T_{out}, T_{in})$; The projection of the red curve onto the plane (T_{out}, T_{in}) represents the neutral temperature model $\hat{T}^*_N = g(T_{out})$. . . . 92
6.9 Comparison of neutral temperature models. . . . 93

A.1 Left: COP-PLR curves of chillers in case study n° 1. Red dots are the 50 samples used to identify the quadratic approximate power consumption model. Right: Identified quadratic P-PLR curves of chillers in case study n° 1. Red dots are the energy consumption data used for the training. 99
A.2 Left: COP-PLR curves of chillers in case study n° 2. Red dots are the efficiency sampled data. Right: P-PLR curves of chillers in case study n° 2. Red dots are the energy consumption data used for the training. 99

List of Tables

2.1 List of parameters 14
2.2 List of variables 14

3.1 Six months - RMSE on test data 25
3.2 Covariate shift experiment - 6-fold RMSE 26

5.1 Comparison of the results of IFA, DCSA, GAMS, TLBO and EIWO on Hsinchu benchmark. 50
5.2 Hsinchu benchmark: P-PLR curves coefficients 60
5.3 Comparison of the results of IFA, DCSA, GAMS, TLBO, EIWO and X-OCL on Hsinchu benchmark. 61
5.4 Published chillers' power consumptions and the xorresponding corrected values. . 62
5.5 Estimated parameters of chillers' P-PLR curves 66
5.6 Case study 1 68
5.7 Case study 3 69
5.8 Estimated parameters of the quadratic power consumption model (5.2) and corresponding $CV\%$ 71
5.9 Execution times for OCL and OSC benchmarks 77

6.1 Sources of raw data for the RP-884 databaset 84

A.1 Coefficients of chillers' COP-PLR curves 98
A.2 Estimated parameters of chillers' P-PLR curves 98

Chapter 1

Introduction

> *"Growing demand for air conditioning is one of the most critical blind spots in today's energy debate."*
>
> – Fatih Birol, *Executive Director IEA*

The use of Heating, Ventilation and Air-Conditioning (HVAC) systems is one of main responsible for the global electricity-demand growth. Such systems, being designed to fulfill civil and industrial environmental requirements, are installed and operated in different types of buildings such as industrial, commercial, residential and institutional buildings. Depending on the context, the complexity of an HVAC system can range from a small unitary system to a large centralized air conditioning system with primary and secondary distribution loops and central plant components.

The *International Energy Agency* (IEA) in its report "The Future of Cooling" explains that the world is facing a looming "cold crunch" [1]. The use of air conditioners in home and offices already accounts for about 20% of the total electricity used in buildings around the world and about 10% of the global electricity consumption. Suffice it to say that about a third of global households own an air conditioner. In countries such as the United States and Japan, more than 90% of households have air conditioning, compared to just 8% of the 2.8 billion people living in the hottest parts of the world. This trend is expected to grow as the world's economic and demographic growth unfolds also in the hotter countries.

In industrial applications - such as the clean rooms of semiconductor and pharmaceutical industries, medium/large scale shopping centers, chemical and manufacturing plants - HVAC systems use up to half of any given industrial building's total energy requirements [8]. In Japan about 60% to 70% of chillers are employed in industrial contexts, while the others are used for residential and commercial uses. The use of energy for cooling buildings has more than tripled between 1990 and 2016, and is foreseen to triple again by 2050.

Increased air conditioning loads entail not only an increase of the overall power demand, but also the need for generation and distribution capacity able to meet demand at peak times, placing further stress on the power system. In the longer term, the cooling demand could be reduced by either improving buildings' design or replacing existing HVAC equipments with more efficient ones. In the immediate, however, it is more sustainable and cost effective to improve the control algorithms so as to achieve higher operating efficiency. In a recent research it has been estimated that between 20% and 30% of building energy consumption could be saved through optimised operation and management without changing any component [2].

The development of advanced systems for building energy management is also one of the major topics within Industry 4.0 innovations programs. Modern energy and utility management systems are going to increasingly rely on data science techniques to extract knowledge from large volumes of data produced by the continuous monitoring and collection of field measurements. Possible applications include, but are not limited to, development of key performance indicators (KPIs), optimization strategies, fault detection, predictive maintenance and smart decisions through real-time communication with humans, machines and sensors.

This thesis addresses the reduction of the power consumption of existing air-conditioning systems through advanced modeling, performed by means of different learning methods, including classical statistical models, neural networks and kernel-based machine learning algorithms. The efficient management of HVAC systems is faced from two different perspectives, the production side and the load demand one. The former approach applies to large industrial and commercial buildings and calls for the design of energy-effective strategies for the optimal management of a multiple chillers system. The latter approach is applicable to any kind of building where the occupants' thermal comfort is the first target (homes, offices, gyms, shopping malls, etc.). The key is the estimation of thermal comfort models, usable for finding the air temperature set-point that optimizes the balance between thermal comfort and energy consumption. In this respect, accepting the suggestion from the so-called *adaptive theory* [3–8], the experimental relationship between thermal comfort, indoor air temperature and outdoor air temperature is investigated. Of course, in medium-large residential and commercial buildings with centralized HVAC systems, both approaches can be applied to boost the energy savings.

The first issue addressed in the thesis is the data-driven identification of efficiency models for water-cooled chillers, with special focus on the opportunities and challenges of a data-rich context. Chiller models are typically classified as either semi-empirical (grey-box), e.g. the Gordon-Ng Universal model [9] and its variants, or empirical (black-box). So far, the predictive capabilities of alternative models have typically been assessed and compared on data sets coming from laboratory tests or provided by chiller manufacturers as a result of the interpolation of few well-selected data points. The performances predicted by models identified on such ideal data could be overly optimistic with respect to field performances. On the other side, processing large

flows of field data by means of innovative computational intelligence methods has the potential to yield more accurate models, but also raises novel robustness issues.

In the machine learning literature, it is known that the statistical distribution of the covariates, i.e. the input variables of the model, plays a key role in determining its predictive capabilities. In industrial environments, characterized by fluctuations of plant operating conditions, the training dataset often covers just a portion of the input space, thus being only partially representative of test data statistics. Under this *covariate shift*, the identification procedure is challenged to find a model that provides satisfactory predictions also when covariates will fall in regions of the input space that are scarcely represented in the training data. In the thesis, the robustness properties of four literature models and a newly proposed Gaussian Process model are assessed on extensive field data. The results show that the best performances in terms of accuracy and robustness are provided by machine learning algorithms, i.e. the Gaussian Process model and the MLP artificial neural network, just followed by the semi-empirical Gordon-Ng Universal (GNU) model, that seems to take some advantage of its physical grounds.

The resilience of the GNU model against the covariate shift motivates a deeper look to the calibration methods for its parameters, looking for possible improvements. In the literature, the usual Ordinary Least Squares (OLS) approach has already been questioned due to the use of regressors affected by measurement noise, and the Errors in Variables (EIV) estimator proposed as more suited to yield unbiased estimates of the parameters. In the thesis, a third approach, based on Nonlinear Least Squares (NLS), is proposed and compared to OLS and EIV. Moreover, rather than on the unbiasedness of parameters, the attention is focused on the prediction of the chiller's coefficient of performance, a more significant target as long as energy efficiency is the final goal. Under a statistical viewpoint, NLS would be the most correct estimator, but on the considered case studies it turns out that the performance improvements over the simpler OLS calibration are not worth the increase of algorithmic complexity.

A crucial point for the efficient management of HVAC plants is the algorithmic optimization of the simultaneous operation of a multiple chillers system. The problem of determining the load fraction that each chiller has to deliver to minimize the overall system power consumption, subject to cooling load satisfaction, is known in the literature as Optimal Chiller Loading (OCL) problem. The OCL is a nonlinear, constrained, often non-convex optimization problem with both continuous and discrete variables. Its NP-completeness motivates the widespread use of heuristic methods which, though not guaranteeing optimality, strive to obtain a fair solution in a viable execution time. In the thesis, OCL is formulated as a Mixed Integer Quadratic Programming (MIQP) and, deviating from the heuristics used in the literature, an exact algorithm, named X-OCL, is worked out, based on a suitable partition of the solution space, assuming that chillers' power consumptions are quadratic in the part load ratio.

With reference to a widespread benchmark, the Hsinchu plant, the exact algorithm is used to say the final word on the relative performances of some leading heuristic methods. Moreover, X-OCL can be used to derive a lower bound (R-OCS) on the best achievable performance of the Optimal Chiller Sequencing (OCS) problem, where constraints on chillers' minimal up- and down-time constraints are added to the OCL problem. A notable feature of X-OCL is its complete parallelizability in elementary quadratic programming problems with equality constraints that admit a fast closed-form solution. It turns out that, even without fully exploiting parallelizability, the computational burden of the exact solver is affordable for the typical OCL and OCS benchmarks. Moreover, using X-OCL as the basis for a greedy solution of the OCS problem yields an algorithm, called X-OCS, whose performance is very close to the best achievable performance. The usability of the exact solver is further validated through the retrospective solution of the OCS problem over a 2-year period for the chiller system of a semiconductor plant.

The second approach to energy saving, i.e. the load demand one, hinges on the identification of occupants' thermal comfort models and the derivation of the so-called *neutral temperature curve*. In fact, HVAC energy demand is directly related to the indoor temperature setpoint, which results from a trade-off between thermal comfort and energy saving objectives. In the 70's the so called adaptive theory introduced the concept of physiological adaptation (acclimatization), suggesting that adapting the indoor temperatures setpoint to the outdoor temperature could help saving energy without compromising thermal comfort. In this view, statistical learning techniques have been applied to an experimental ASHRAE dataset in order to investigate and quantify the existence of the "adaptive effect", i.e. the higher is the outdoor temperature the higher is the indoor temperature that occupants consider comfortable. The results show a convergence between different statistical and machine learning approaches. Moreover an new definition of neutral temperature, that appears statistically sounder, is proposed and tested on the ASHRAE RP-884 dataset.

1.1 Thesis Overview

The thesis is organized in eight chapters (included the Introduction).

Chapter 2: HVAC systems: basic notions
The chapter provides an overview on centralized water-cooled HVAC systems, paying particular attention to chillers, which are the most energy-intensive components of such systems.

Chapter 3: Learning chillers efficiency from data
The optimal energy management of multiple chiller systems calls for the construction of mathematical models of chiller energy efficiency. The existing grey- or black-box models include

parameters that have to be estimated from experimental data. So far, the predictive capabilities of alternative models have been assessed and compared on data sets created by laboratory tests or provided by chiller manufacturers. In an Industry 4.0 context, the continuous monitoring and collection of field data discloses new opportunities but raises also robustness issues that are herein addressed. Herein, exploiting an extensive experimental dataset collected over a six-month period, four literature models and a new machine learning approach are compared. The second objective is assessing the robustness of the five models against covariate shifts, i.e. variations in the statistical distribution of the input variables that occur across different months. The grey-box Gordon-Ng model, though less accurate in nominal conditions than the Bi-quadratic and Multivariate polynomial models, proves however more robust against covariate shifts. The best performances, both in term of accuracy and robustness, are however provided by the two machine learning methods, with the Gaussian Process model performing better than the MLP artificial neural network.

Chapter 4: Calibration strategies for the Gordon-Ng chiller model

In this chapter, the calibration of the parameters of the Gordon-Ng Universal (GNU) chiller model is investigated. In its standard formulation, the GNU model is written as a linear-in-parameter model that can be calibrated by Ordinary Least Squares. It has been already observed elsewhere that, since the regressors are subject to measurement inaccuracies, the OLS approach is prone to yield biased estimates of the parameters. As a remedy, Andersen and Reddy proposed the adoption of an Errors in Variable (EIV) framework, showing that bias could be reduced or even eliminated by means of a corrected least squares algorithm. However, some questions remained open. Given that the EIV approach achieves bias reduction at the cost of increasing the variance, is it really preferable to OLS? If the final goal is not parameter estimation, but the prediction of the Coefficient of Performance (COP), how does OLS compare with EIV? And what is the most appropriate calibration method, under a statistical viewpoint? Finally, is the added complexity of a statistically rigorous approach employing Nonlinear Least Squares (NLS) really worth the potential improvements in COP prediction? In order to answer these questions, three estimation methods, OLS, EIV and NLS, are tested on two benchmarks: a public precise chiller performance dataset and an ASHRAE dataset. The results suggest, that OLS estimation, in spite of its suboptimality, may prove largely satisfactory both for parameter estimation and COP prediction, although it may be worth analyzing other more challenging COP prediction problem before the final word is said.

Chapter 5: Chiller plant optimization

For a multiple chiller system, Optimal Chiller Loading (OCL) deals with the problem of finding the configuration of partial load ratios that minimizes energy consumption while satisfying a prescribed total cooling demand. Due to the interplay of continuous and logical constraints,

OCL is an NP-hard problem, so that a variety of heuristic algorithms have been proposed in the literature. Herein, we develop an algorithm, named X-OCL, for its exact solution under the assumption that the chillers' power consumption curves are quadratic. The proposed method hinges on a decomposition of the solution space so that the overall OCL problem is reduced to a set of quadratic programming problems subject to only equality constraints that can be solved in closed form. The availability of an exact solver is exploited in order to assess and compare the performances of several literature algorithms, highlighting also some errors in the publihed results, due to incorrect extrapolations of the power consumption curves. Differently from the static nature of OCL, the Optimal Chiller Sequencing problem aims at real-time minimization of the energy consumption taking into account also minimal up- and down-time constraints that the chillers must satisfy. The new X-OCL solver is used for two purposes. First, X-OCL is used to compute a lower bound to the best OCS performance. Second, it used as the basis for the design of a greedy OCS solver, called X-OCS. The performances of X-OCS are tested on two literature benchmarks and on a Field data benchmark given by the HVAC systems of a semiconductor plant. In all cases, the performances of X-OCS are remarkably close to the lower bound, suggesting that there is no need for more sophisticated algorithms. The general applicability of X-OCS depend on its computational burden and on the quadratic assumption on the power consumption curve. In spite of the exponential growth of the computational burden with the number of chillers, the average computation time per load remains affordable even for a 9-chiller plant. Moreover, the X-OCL and X-OCS algorithms are amenable to massive parallelization, a feature that could be exploited to further move the boundary of treatable problems. Regarding the quadratic assumption, it is consistent with most of the literature benchmarks and it proved adequate also for the data-driven identification performed on the experimental data of the Field-data benchmark.

Chapter 7: Thermal Comfort in Air-Conditioned Buildings

An essential requirement for thermal comfort is maintaining an indoor air temperature considered satisfactory by the majority of occupants. According to the so-called adaptive theory, such a *neutral temperature* may change with the outdoor air temperature, a feature that can be exploited to achieve energy savings without reducing thermal comfort. In the literature, the characterization of this dependence relies on a two-step procedure. First, occupants' thermal sensation votes are processed at building level and outdoor temperatures are averaged to obtain a neutral and outdoor temperature pair for each building. The pairs are then used to fit the neutral temperature model. Herein, three approaches for estimating neutral temperature models without need for a preprocessing at building level are proposed and validated on the summer data of the ASHRAE RP-884 database: (i) regression of temperatures considered neutral by the occupants against the outdoor temperature; (ii) direct regression of the ASHRAE votes against indoor and outdoor temperature; (ii) estimation via logistic regression of the indoor temperature

that maximizes the percentage of satisfied users. Overall, seven neutral temperature models are successfully worked out. A first finding is that the the adaptive hypothesis can be ascertained also by models formulated at raw data level. When compared to each other, the seven models are in good agreement, especially those within the same approach, thus demonstrating the viability of neutral temperature modeling based on individual raw data.

Chapter 8: Conclusions
In the final chapter, the main findings are summarized and discussed.

1.2 Collaborations and related pubblications

The material of Chapter 3 partially appears in:

- F. Acerbi, G. De Nicolao. "Identification of the Gordon-Ng Chiller Model: Linear or Nonlinear Least Squares?." 2018 IEEE 23rd International Conference on Emerging Technologies and Factory Automation (ETFA). Vol. 1. IEEE, 2018.

The material of Chapter 4 and 5 was developed within the ECSEL European project "Power Semiconductor and Electronics Manufacturing 4.0" (SemI40).

- F. Acerbi, G. De Nicolao, et al. "Accuracy and Robustness Against Covariate Shift of Water Chiller Models." 2018 IEEE 14th International Conference on Automation Science and Engineering (CASE). IEEE, 2018.

The material of Chapter 6 was partially published in:

- F. Acerbi, G. De Nicolao, M. Rampazzo. "Thermal Comfort Control in Air-Conditioned Buildings: new data-driven approaches to Neutral Temperature estimation." 2019 18th European Control Conference (ECC). IEEE, 2019.

2

HVAC Systems: basic notions

Heating, Ventilation and Air-Conditioning (HVAC) systems are a combination of different equipments working together to provide heating and cooling air in buildings according to the designed requirements. Depending on the application, the complexity of an HVAC system can range from a small unitary system to a large centralized air conditioning system with primary and secondary distribution loops and central plant components. While an exhaustive taxonomy of HVAC systems and components is beyond the scope of this thesis, the interested reader may refer to [10]. The aim of this chapter is to provide an overview of centralized water-cooled HVAC systems with primary-secondary architecture, which will be the reference scheme along the thesis.

2.1 Water-cooled HVAC system

The scheme of a typical water-cooled HVAC system is illustrated in Fig. 2.1. It is mainly composed by three energy exchange loops:

- *Air distribution loop* Cold air is distributed by the *Air Handling Unit* (AHU) to maintain a certain indoor air temperature set-point in rooms within the building. When the indoor air temperature rises and exceeds the set-point it is returned to cooling coils where the chilled water extracts its excess heat, reducing both its temperature and humidity content, after which it is distributed once again to the rooms.

- *Chilled water loop* As the chilled water passes through the AHU's cooling coils in counter flow to the air loop, the heat extraction process increases the water temperature. Therefore, the chilled water leaving the cooling coil (chilled water return) will be warmer than the entering water (chilled water supply).

The equipment in charge of removing such excess heat from the chilled water loop is the chiller. The chiller exploits the vapor compression refrigeration cycle or the absorption refrigeration cycle to remove heat from the chilled water loop and reject it to the atmosphere by means of the condenser water loop.

The return chilled water enters the chiller where it is cooled to the desired temperature by transferring the heat to a primary refrigerant. This process is made possible by the work spent by the compressor for cooling the primary refrigerant.

- *Condenser water loop* The heat removed from the chilled water loop together with the compressor power input must be rejected to the atmosphere. This waste heat can be rejected by water-cooling, a process that uses water to collect the heat from the refrigerant and then reject it to the atmosphere.

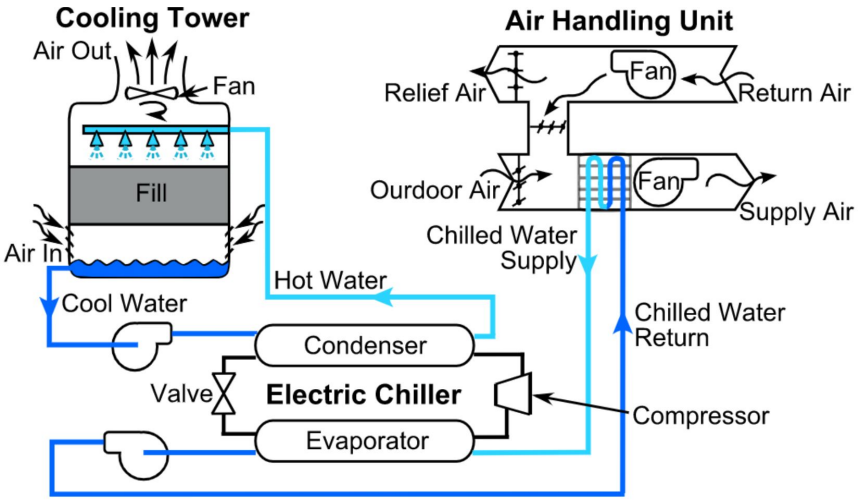

FIGURE 2.1: Water-cooled HVAC system scheme

2.2 Vapor-compression chiller

Vapor-compression chiller are widely employed in large commercial and industrial refrigeration systems. As shown in Fig. 2.2, the vapor-compression refrigeration system consists of four components that perform the four steps of the refrigeration cycle:

- *Evaporator* The evaporator in a centrifugal water cooled chiller is usually a shell and tube heat exchanger that removes heat from the entering chilled water, lowering its temperature in the process. The heat is used to boil the refrigerant changing it from a liquid to a gas.

- *Compressor* The compressor assembly is made up of a prime mover and a centrifugal compressor. The centrifugal compressor is a dynamic device similar to a centrifugal water

pump. It raises the pressure and temperature of the refrigerant by converting kinetic energy into pressure.

- *Condenser* Like the evaporator, the condenser is usually a shell and tube heat exchanger. It removes heat from the refrigerant gas causing it to condense to a liquid. The heat raises the temperature of the cooling water often referred to as condenser water. The condenser water then carries the heat to the cooling tower where the heat is rejected to atmosphere.

- *Expansion device* After the refrigerant condenses to a liquid, it passes through a pressure reducing device. This can be as simple as an orifice plate or as complicated as an electronic modulating expansion valve.

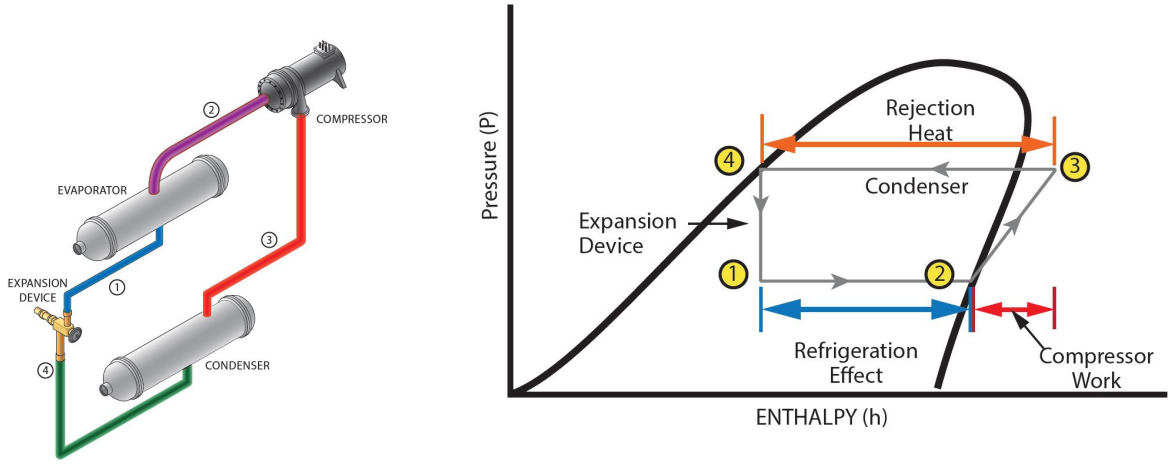

FIGURE 2.2: Left: basic refrigeration cycle. Right: refrigerant circuit; P-H Diagram.

Thermodynamically, the most common way of looking at the refrigeration cycle is the Pressure-Enthalpy (P-H) diagram, see fig. 2.2.

The evaporator process is from point 1 to point 2. As the refrigerant changes from a liquid to gas, the pressure stays constant. The heat is being absorbed as a phase change (latent energy). The refrigeration effect is the change in enthalpy from 1 to 2. The line from 2 to 3 represents the compression process. The work is the change in enthalpy from point 2 to point 3. Work of compression ends up as heat in the refrigerant. Then, there is the rise in refrigerant pressure from 2 to 3. The next process takes place in the condenser (from 3 to 4). The first section (outside the refrigerant dome) is the desuperheating process. Once the refrigerant is saturated, condensation occurs and the refrigerant changes from a gas to a liquid. The line from 3 to 4 is horizontal, indicating constant pressure. The final process is the expansion device, shown as a vertical line from 4 to 1, representative of the pressure drop that occurs as the refrigerant passes through the expansion valve.

2.3 Chiller efficiency

The chiller-operation efficiency is normally measured by the Coefficient of Performance (COP), defined as the ratio of the supplied cooling capacity Q_e, i.e. the heat extracted by the chiller, over the compressor electric power consumption P_{el}:

$$COP = \frac{Q_e}{P_{el}} \tag{2.1}$$

A centrifugal chiller is designed to operate at maximum efficiency at or near the rated point. However, as the chiller water temperature, cooling water temperature, and cooling capacity move away from the rated point, the losses in the refrigeration cycle increase and the chiller efficiency degrades.

Conventional design methods for chiller plants take into account both the annual cooling load distribution of buildings and their peak cooling loads, based on typical meteorological year data. Since the peak cooling load is reached only occasionally, this means that, most of the time, the chillers may operate at less than their full rated load, possibly far away from their optimal efficiency. Therefore, the characterization of chillers' part-load efficiency at actual working conditions is a key step in order to identify potential energy savings.

The final goal of the efficient management of a multi-chiller system is satisfying the cooling load demand while minimizing the system power consumption. The chiller power consumption can be easily derived by its COP. Indeed, if the chiller's efficiency is low (e.g. due to poor maintenance), then more electricity will be consumed to support the cooling capacity requirement. Alternatively, the characterization of the chillers could skip COP modeling and address directly the estimation of their power consumption curves, that provide all the information needed for optimization purposes. There exists an extensive literature on the identification of suitable models to describe either chiller efficiency or chiller power consumption. We will dwell on these topics in Chapters 3, 4 and 5.

2.4 Plant description

The experimental field data used in Chapters 3 and 5 were collected in a centralized HVAC plant installed in a large semiconductor factory located in Austria. The served buildings are almost totally devoted to the semiconductor production, but there are also offices and a canteen.

The HVAC plant has a primary-secondary architecture. The chilled water system is separated into primary and secondary loops by a decoupling bypass line. In the primary loop, see fig. 2.3, the water cooled chillers are connected in parallel. Each chiller has an associated constant primary pump. The secondary chilled water pumps are of variable speed type, controlled to

maintain a constant differential pressure across the building supply and return water. Two-way control valves are used to modulate the chilled water flow through individual AHU terminals. Variation of total building cooling load is reflected in the flow rate and temperature difference in the secondary loop. The number of working chillers (and primary pumps) will vary in order to follow the variations of the building cooling load, thus resulting in the on- and off-cycling of different chillers. In the primary loop, the distribution pumps can operate only according to an on/off behaviour. When switched on, they circulate a water flow rate whose nominal value is 464 m^3/h. In the other case (pumps switched off), the water flowing through the chiller is not cooled down and its amount depends on the pressure drop between the delivery and return.

A two-level control architecture is applied. The high-level supervisor specifies the modes of operation and the set-points for each chiller, while local controllers are in charge of maintaining the individual chiller's set-point. The system is controlled by means of a non-linear relay-based logic, that regulates the switching of chillers based on the temperature of the return water. The set point value for the outlet temperature of each chiller is 5°C ±1°C.

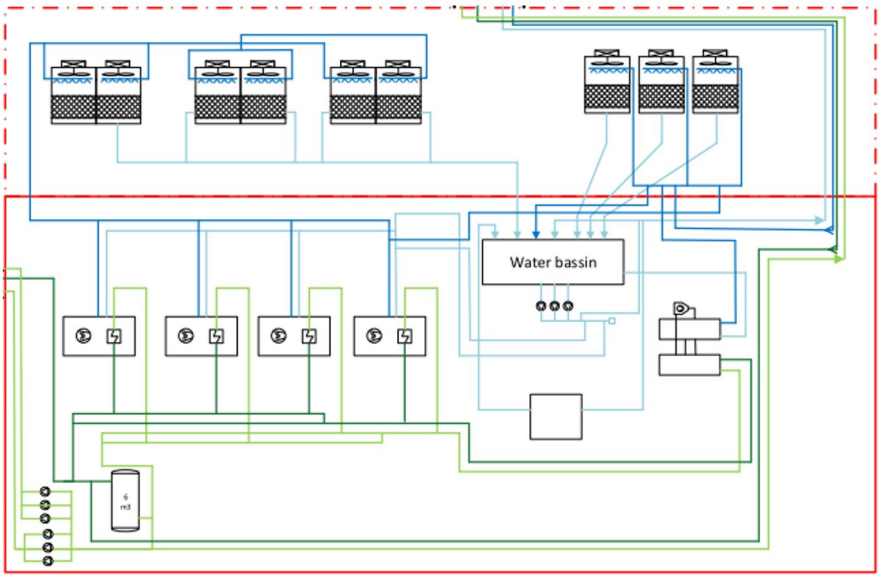

FIGURE 2.3: Primary circuit layout

2.4.1 Dataset

Several variables in the primary circuit were collectes over the period Jan 2016 - Jan 2019 with a sampling time $T_s = 60\ min$. Data were measured in closed-loop working conditions. Before starting the identification process, the data have been preprocessed removing outliers and infeasible records. The main parameters and variables of interest are listed in Tables 2.1 and 2.2, where t denotes the discrete time index.

TABLE 2.1: List of parameters

Symbol	Description	Value	Unit
T_s	Sampling time	1	h
$T_{set,i}$	Evaporator outlet water setpoint of $i-$th chiller	5	$°C$
$m_{evap,i}$	Nominal evaporator water flow of $i-$th chiller	464	m^3/h
$Q_{nom,i}$	Nominal cooling capacity of $i-$th chiller	2700	kW

TABLE 2.2: List of variables

Symbol	Description	Unit
$Q_{e,i}(t)$	Evaporator cooling capacity of $i-$th chiller	kW
$Q_{cond,i}(t)$	Condenser cooling capacity of $i-$th chiller	kW
$Q_{load}(t)$	Cooling load demand	kW
$m_{evap,i}(t)$	Evaporator water flow of $i-$th chiller	m^3/h
$m_{cond,i}(t)$	Condenser water flow of $i-$th chiller	m^3/h
$T_{ei,i}(t)$	Network evaporator inlet water temperature	$°C$
$T_{eo,i}(t)$	Evaporator outlet water temperature of $i-$th chiller	$°C$
$T_{ci,i}(t)$	Network condenser inlet water temperature	$°C$
$T_{co,i}(t)$	Condenser outlet water temperature of $i-$th chiller	°C
$P_{el,i}(t)$	Electric power consumption of $i-$th chiller	kW

3

Learning chillers efficiency from data

The combination of technology and scientific knowledge that is nowadays available has opened the way to the optimized management of large Heating, Ventilation, Air-Conditioning and Refrigeration (HVAC&R) systems, with significant environmental and financial fallouts. For example, in semiconductor manufacturing plants these facilities account for the greatest fraction of total electric energy consumption, amounting to approximately 25-30% of the fab total energy requirement. Between 40% and 60% of this requirement is due to chillers, which are therefore responsible for about 15% of the total energy use. Such energy intensiveness is a primary motivation for investigating the optimal management of multiple chiller systems [11–14].

Propaedeutic to any optimized management of a multiple chiller system is the knowledge of a mathematical model of their efficiency. Several models for predicting energy efficiency performances of water chillers have been proposed over last decades. They are typically classified as empirical (black-box) or semi-empirical (grey-box), depending on the amount of physical knowledge embedded in the model structure. Notwithstanding the effort required by the model building process, grey-box models maintain a certain appeal due the physical interpretability of the parameters, a feature that may prove worthy in certain industrial applications, e.g. fault detection [15–18]).

More than twenty years ago, Gordon and Ng [9] introduced a semi-empirical modeling approach for reciprocating chillers that has been later extended to several chiller types [19]. On the other hand, black-box models, in spite of their lack of physical insight, have been gaining an ever wider acceptance due to their prediction accuracy and generalization ability. Lee and Lian [20] compared eleven empirically-based performance models showing that, for many chillers types, the Bi-quadratic regression model [21] and the multivariate polynomial model [22] yielded the most accurate predictions. Swider [23] has expanded this comparison by including also neural network models that appeared able to further improve on the generalization abilities.

Despite a number of contributions, a comprehensive comparison based on extensive datasets is still lacking for industrial applications involving medium-large sized on-field operating refrigerators. Furthermore, model comparisons and validations were often based on chillers data set created by laboratory tests or provided by chiller manufacturers. According to [19, 24], after publication of the ARI standards [25] chillers manufacturers ceased to publish extensive datasets: most of the performance data available through the manufacturer sales representatives result from the interpolation of few (five to six) well-selected data points. A first objective of this chapter is to compare four models proposed in literature and a fifth new Gaussian-Process method on an extensive six-month dataset recorded on a water-cooled chiller which is part of a large HVAC system.

The predictive performance of models identified from data collected according to a properly designed experiment could be overly optimistic with respect to the actual performances of models identified from field data. In the machine learning literature, it is known that the statistical distribution of the covariates, i.e. the input variables of the model, plays a key role in determining the predictive capabilities of a model [26]. Indeed, a model trained on data concentrated in a small operating region may generalize poorly in regions that are scarcely sampled in the training dataset. In our case, the covariate distribution of field data depends on plant operating conditions dictated by a load demand that, as it will be seen in the experimental dataset, is subject to significant variations across different months. The second major goal of the chapter is therefore the assessment of the robustness properties of the five considered model in the face of covariate shifts occurring in real-world field data. For this purpose, a covariate shift experiment is performed and analyzed: for each of the six months the models are trained on one-month data and their predictive capability is tested on the remaining five months.

The chapter is organized as follows. In Section 3.1, four energy efficiency models are reviewed and a fifth machine learning model, based on the Gaussian Process framework is introduced. The covariate shift problem is concisely explained in Section 3.2, while in Section 3.3 the data are presented and some of their features discussed. The comparison of the five models identified from the global dataset and the results of the covariate shift experiment are presented in Section 3.4. Some concluding remarks (Section 3.5) end the chapter.

3.1 Energy performance models

In this section, five COP models for predicting the energy efficiency of water chillers are reviewed. The first four are taken from the literature, while the fifth is a machine learning method relying on a Gaussian Process approach.

3.1.1 Gordon-Ng universal model

The Gordon-Ng Universal (GNU) model is a grey-box (semi-empirical) model whose structure is based on first principles of thermodynamics and linearized heat losses. It relies on the following formula, linking together COP, Q_e, T_{ci}, T_{ei}:

$$\frac{T_{ei}}{T_{ci}}\left(1+\frac{1}{COP}\right)-1=\beta_1\frac{T_{ei}}{Q_e}+\beta_2\frac{T_{ci}-T_{ei}}{T_{ci}Q_e}+\beta_3\frac{Q_e}{T_{ci}}\left(1+\frac{1}{COP}\right) \tag{3.1}$$

All the three model parameters β_i admit a physical interpretation [19]:

- β_1: the total internal entropy production rate in the chiller due to irreversibilities;
- β_2: the rate of heat losses (or gains) from (or into) the chiller;
- β_3: the total heat exchanger thermal resistance.

As observed in [27], many manufacturer data-catalogs use the evaporator outlet water temperature T_{eo} instead of T_{ei}. Then, one can rewrite (4.1) with T_{ei} replaced by T_{eo} with the only caution that the physical interpretation of β_3 is slightly changed. The assumptions made in the model building suggest that the application of this model is best suited to chillers with inlet guide vanes capacity control. Since its introduction, a number of variants have been proposed in order to extend it to other types of chillers. In particular, Gordon and Ng discussed the possible introduction of an intercept term, at the cost of loosening the physical interpretation of the parameters [19, 27].

Let $\beta\in\mathbb{R}^3$ denote the unknown parameter vector and define

$$\begin{aligned} X &= \left[\frac{T_{ei}}{\dot{Q}_e} \quad \frac{T_{ci}-T_{ei}}{T_{ci}Q_e} \quad \frac{Q_e}{T_{ci}}\left(1+\frac{1}{COP}\right)\right] \\ Y &= \frac{T_{ei}}{T_{ci}}\left(1+\frac{1}{COP}\right)-1 \end{aligned}$$

Then, from eq. (4.1) the following linear-in-parameter formulation is immediately derived [19, 23, 27]:

$$Y=X\beta \tag{3.2}$$

Given experimental data $COP(i)$, $Q_e(i)$, $T_{ci}(i)$, $T_{ei}(i)$, $i=1,\dots,n$, the ordinary least square estimate β^{LS} is easily computed. It is worth noting that, in order to obtain a linear-in-parameter model, the identification step uses as output the variable Y that differs from COP, whose prediction remains however the final goal. The desired predictor is just obtained by solving

(4.1) for COP:

$$\hat{COP} = \frac{\frac{T_{ei}}{T_{ci}} - \beta_3^{LS}\frac{Q_e}{T_{ci}}}{1 - \frac{T_{ei}}{T_{ci}} + \beta_1^{LS}\frac{T_{ei}}{Q_e} + \beta_2^{LS}\frac{T_{ci}-T_{ei}}{T_{ci}Q_e} + \beta_3^{LS}\frac{Q_e}{T_{ci}}} \tag{3.3}$$

The statistical implications involved by the use of an output variable different from the one to be predicted will be discussed in depth in Chapter 3.

3.1.2 Bi-quadratic regression model

The Bi-Quadratic (BiQ) regression model is a black-box model featuring nine regressors that are rational functions of the covariates Q_e and T_{ci}:

$$\frac{1}{COP} = \beta_0 + \beta_1\frac{1}{Q_e} + \beta_2 Q_e + \beta_3\frac{T_{ci}}{Q_e} + \beta_4\frac{T_{ci}^2}{Q_e} + \beta_5 T_{ci} + \beta_6 Q_e T_{ci} + \beta_7 T_{ci}^2 + \beta_8 T_{ci}^2 \tag{3.4}$$

Similarly to the GNU model, the parameter vector $\beta \in \mathbb{R}^9$ can be estimated via ordinary least squares, provided that the model output is defined as $Y = 1/COP$.

3.1.3 Multivariate polynomial regression model

The Multivariate Polynomial Regression (MPR) model is a black-box model that describes COP as a second order polynomial in the three covariates Q_e, T_{ci}, T_{ei}:

$$COP = \beta_0 + \beta_1 Q_e + \beta_2 T_{ei} + \beta_3 T_{ci} + \beta_4 Q_e^2 + \beta_5 T_{ei}^2 + \beta_6 T_{ci}^2 + \beta_7 Q_e T_{ei} + \beta_8 Q_e T_{ci} + \beta_9 T_{ei} T_{ci} \tag{3.5}$$

Again, the parameter vector $\beta \in \mathbb{R}^{10}$ can be estimated via ordinary least squares.

3.1.4 Multilayer perceptron model

MultiLayer Perceptrons (MLP) are Artificial Neural Networks (ANN) models composed by elementary units, the neurons, which are linked together by weighted connections. The output of each neuron is a nonlinear transformation function, called activation function, where the inputs are either the m independent variables x_j or the outputs of neurons located in the previous layer. The model parameters are given by the weights. Herein, a one-hidden-layer MLP model made of fully connected neurons with activation function *tanh* is considered. The output of the network is a weighted linear combination of the outputs z_i of the n_h hidden-layer

neurons:

$$\begin{aligned} Y &= w_{00} + \sum_{i=1}^{n_h} w_{i0} z_i \\ z_i &= \tanh\left(\sum_{j=1}^{m} w_{ij} x_j + w_{0j}) \right) \end{aligned} \tag{3.6}$$

Here, $x \in \mathbb{R}^m$ denotes the vector of input variables, e.g. $m = 2$ and $x = \begin{bmatrix} Q_e & T_{ci} \end{bmatrix}'$.

Although ANNs are widely appreciated for their prediction and generalization capabilities, their complexity must be carefully tuned to avoid overfitting. When a fully connected structure is assumed, overparametrization may occur even for a relatively small number of neurons. For instance, a fully connected one-hidden layer MP with $m = 2$ inputs has $M = 1 + n_h(m + 2) = 1 + 4n_h$ weights. In order to prevent overfitting, a regularization penalty can be added to the conventional squared loss [28]:

$$J = \frac{1}{n} \sum_{i=k}^{n} e(k)^2 + \frac{\gamma}{M} \left(w_{00}^2 + \sum_{i=1}^{n_h} \left(w_{i0}^2 + \sum_{j=0}^{m} w_{ij}^2 \right) \right) \tag{3.7}$$

where $e(k) = \hat{Y}(k) - Y(k)$ is the prediction error, M is the total number of weights, and γ is the so-called regularization parameter, which can be tuned by either crossvalidation or maximization of the lilkelihood of an associated statistical model [28]. In order to maintain a complexity comparable to that of the BiQ and MPR models, $n_h = 3$ hidden neurons have been used, corresponding to 13 weights to be estimated. The minimization of the regularized loss (3.7) has been performed by means of the function *trainbr* of the Matlab Neural Network Toolbox.

3.1.5 Gaussian process model

Gaussian Process (GP) regression is a nonparametric statistical learning technique, relying on the assumption that the vector of output data $\begin{bmatrix} Y(1) & Y(2) & \dots & Y(n) \end{bmatrix}^T$ is a normally distributed vector, obtained by sampling the following Gaussian process (a random field, for the sake of precision, because $x \in \mathbb{R}^m$):

$$y(x) = f(x) + \varepsilon \tag{3.8}$$

Here, x is the covariate vector, ε is the noise term and $f(\cdot)$ is a zero mean Gaussian process, characterized by the (auto)covariance function

$$E[f(x)f(x'))] = k(x, x') \tag{3.9}$$

where $k(x, x')$ is the so-called kernel, that has to be properly chosen so as to encode all prior knowledge about the shape and structure of the unknown function $f(\cdot)$. An exhaustive overview of the many possible kernels is provided by Rasmussen [29]. Herein, we consider the combination of two widely used kernels, namely the linear and squared exponential one:

$$k(x, x') = \sigma_l^2 \left(1 + \sum_{i=1}^{m} x(i)x'(i)\right) + \sigma_s^2 e^{\left(-\frac{\|x-x'\|^2}{2l^2}\right)} \tag{3.10}$$

where the signal variances σ_l^2, σ_s^2, and the length-scale l^2 are called hyperparameters of the Gaussian Process.

The signal variances act as scaling factors. Small values are associated with functions that stay close to zero, while larger values allow more variation. If the signal variance is too large, the modeled function will be free to chase outliers. The length-scale instead describes the correlation between points and, consequently, the smoothness of the resulting function.

Hyperparameters can be tuned through the maximization of the so-called marginal loglikelihood:

$$\begin{aligned} logp(y|x) &= -\frac{n}{2}log2\pi - \frac{1}{2}log|K + \sigma_n^2 I| \\ &\quad -\frac{1}{2}y^T(K + \sigma_n^2 I)^{-1}y \\ K &= [K]_{ij} \in \mathbb{R}^{n\times n}, K_{ij} = k(x(i), x(j)) \end{aligned} \tag{3.11}$$

Once the hyperparameters have been tuned, it is possible to compute the prediction $\hat{f}(x^*)$ in correspondence of a new covariate vector x^* as

$$\hat{f}(x^*) = \sum_{i=1}^{n} \alpha_i k(x(i), x^*) \tag{3.12}$$

$$\alpha = (K + \sigma_n^2 I)^{-1}y \tag{3.13}$$

3.2 The Covariate Shift Problem

The aim of a prediction model is to identify a relationship that links the target output variable Y to the observed input vector x. For a given model structure, identification is performed using a training set made of input-output pairs $\{Y(i), x(i)\}, i = 1, \ldots, n$. The final goal is the accurate prediction of a test output Y^* associated to an unknown test input x^*.

Most techniques rely on the assumption that the test pair $\{Y^*, x^*\}$ will be drawn from the same joint distribution $p(x, Y)$ that generated the training data. Nevertheless, it may well happen that this hypothesis is violated, e.g. in industrial environments characterized by fluctuations of plant operating conditions. Indeed, the training dataset often covers just a portion of the

input space, thus being only partially representative of test data statistics. In these cases, the identification procedure is challenged to find a model that provides satisfactory predictions also when x^* falls in portions of the input space that are less represented in the training data.

The circumstance where the training and test inputs follow different distributions p^{train} and p^{test}, but the conditional distribution $p(Y|x)$ remains unchanged, has been studied in the machine learning literature under the name of *covariate shift problem*:

$$\begin{aligned} p^{\text{train}}(x,Y) &= p(Y|x)p^{\text{train}}(x) \\ p^{\text{test}}(x,Y) &= p(Y|x)p^{\text{test}}(x) \\ p^{\text{train}}(x) &\neq p^{\text{test}}(x) \end{aligned}$$

Covariate shifts have the potential of exacerbating biasedness of predictions performed on test data. Remedies that are effective when the sample size of the training dataset tends to infinity have been worked out [26], but in the finite sample case, the covariate shift remains a challenging problem whose impact is largely unexplored in application contexts.

In order to illustrate the possible effects of a covariate shift on the quality of output predictions, let us consider a toy regression example. Assume that $p(Y|x) \sim N(xe^{-x}, 0, 01)$. We want to train a predictor given by the third-order polynomial

$$\hat{Y} = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 x^3$$

whose parameters are estimated via ordinary least squares on a training dataset whose inputs are 50 samples $x(i)$ drawn from the distribution $p^{\text{train}}(x)$. Two alternative choices for p^{train} are considered, uniform and exponential:

$$\begin{aligned} p^u(x) &= \begin{cases} \frac{1}{b-a} & a \leq x \leq b \\ 0 & \text{otherwise} \end{cases} \\ p^e(x) &= \begin{cases} e^{-\lambda x} & x \geq 0 \\ 0 & x < 0 \end{cases} \end{aligned}$$

where $\lambda = 0.6$, $a = 0.5$, $b = 3$.

In Figure (3.1), the training dataset and the obtained predictor $\hat{Y}(x)$ are plotted for the two cases. The predictor in the left panel, trained on exponentially distributed input data, performs well for small values of x, at the cost of worse performances for large values of x. The predictor in the right panel, trained on uniformly distributed data, optimizes the average prediction error for test data sampled from the uniform distribution at the cost of delivering worse predictions for small values of x.

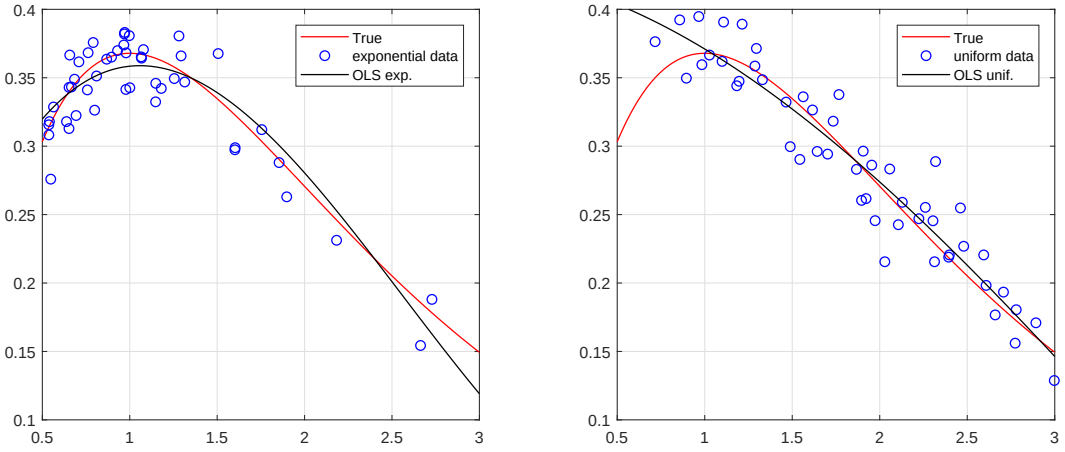

FIGURE 3.1: Regression under covariate shift. Left: predictor trained via OLS on exponentially distributed data. Right: predictor trained on uniformly distributed data

There is an obvious problem, however, if we need to optimize the predictor performance on exponential test data but only uniform training data are available, because the predictor in the right panel will inevitably lack accuracy for small values of x.

The example demonstrates that the predictive performance in some regions can be affected by shifts of the covariate distribution of training data, especially in presence of a model mismatch, which is rather common when complex physical processes are involved.

3.3 Field data

This study uses experimental data from a water-cooled chiller which is part of a large HVAC system that supplies cooling water to the Air Handling Unit (AHU) heat exchangers of a semiconductor fab. The dataset is composed by over $4,000$ data points collected at different working conditions from June until October 2016. The chiller is a Sulzer Unitop 22A–5106 with R134a refrigerant and nominal cooling capacity of 2.7 MW, installed in a Constant Primary Flow (CPF) chilled water system. It is equipped with a centrifugal compressor, whose capacity control is based on the regulation on the refrigerant gas flow by adjustable Inlet Guide Vanes (IGV). The chiller is regulated with constant evaporator outlet temperature T_{co}, whose setpoint is 5 °C.

The global dataset was randomly partitioned in two datasets: 70% for training and 30% for testing, respectively, see Fig. (3.2), where also the covariate distribution is displayed.

In order to assess to what extent the covariate distribution of field data is subject to shifts during the six-month period, in Fig. (3.3) the monthly datasets are displayed separately. The visual comparison of the covariate distributions, represented as green surfaces, shows that they

are subject to substantial changes across the semester. The most critical month is October, when the input data (red dots) are concentrated in a narrow region.

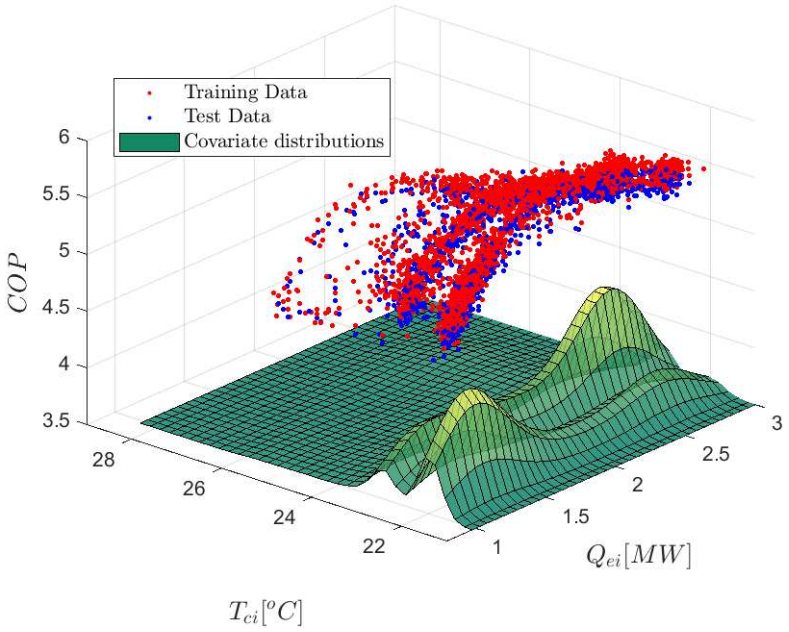

FIGURE 3.2: Overall six-month dataset. Red: training data; Blue: testing data; green: kernel estimate of the joint distribution of the input covariates $Q_e(i)$ and $T_{ci}(i)$.

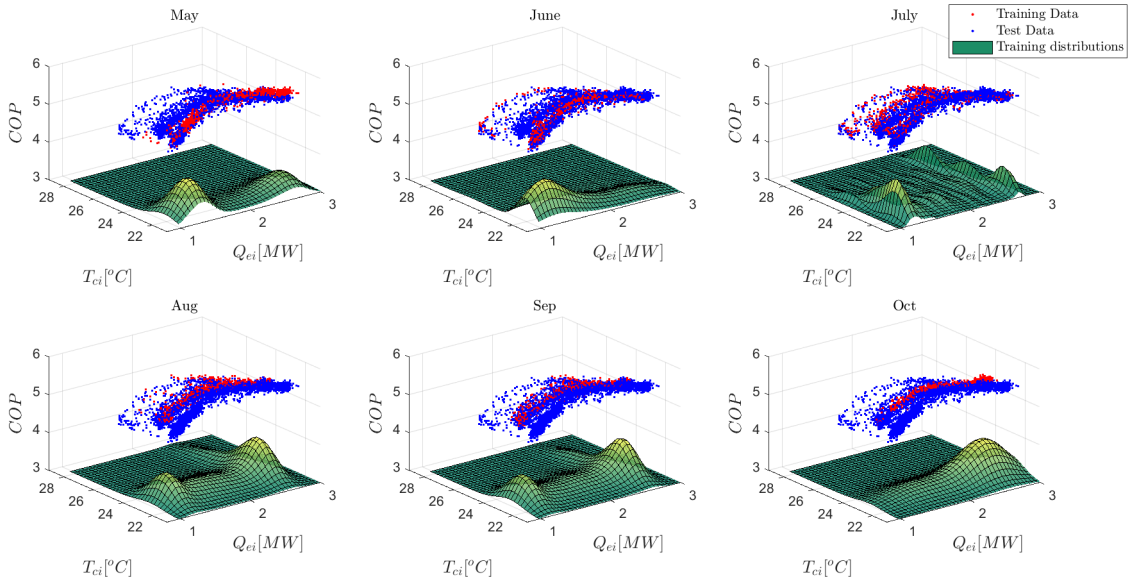

FIGURE 3.3: Monthly datasets. Red: training data; Blue: test data; green: kernel density estimate of the training distribution of the input covariates $Q_e(i)$ and $T_{ci}(i)$.

Not only data are concentrated it can happen to collect data which covers only a limited portion of the machine working space and having to extrapolate information far from these learned areas. This phenomenon, previously introduced as covariate shift problem is well represented in Fig. 3.3. The first picture show the complete available dataset, which was divided into 70% for

training and 30% for testing, and their joint probability distribution. In this case, since the training set and the test set are both numerous and informative enough they have the same density function. In contrast, the second one illustrates how the probability distribution of the training data significantly changes depending on the considered month. The red dots are the monthly training data while for each month the test set is the sum of the remaining five months data.

3.4 Models comparison

This section is divided in two parts. First, the five models presented in Section 3.1 are identified and compared on the global dataset. In the second subsection, the robustness of the models against the covariate shift is assessed using the monthly datasets. In all cases, the root mean squared error (RMSE) computed on the relevant test data was used as a performance indicator. Being an energy efficiency, the output is adimensional.

3.4.1 Model comparison: global dataset

The BiQ model depends on Q_e and T_{ci} alone. The GNU model conversely may depend on either (Q_e, T_{ci}, T_{ei}) or (Q_e, T_{ci}, T_{eo}). Since performances were comparable, the triplet (Q_e, T_{ci}, T_{eo}) was chosen, also because in the available data T_{eo} is regulated to a constant setpoint, implying that the efficiency model is completely captured by the surface $COP(Q_e, T_{ci})$. The MPR model depends on (Q_e, T_{ci}, T_{ei}). Actually, T_{ei} was found to affect the predicted COP only marginally. In view of this, when training MLP and GP models, just Q_e and T_{ci} were used as covariates.

The five models presented in Section 3.1 were identified on the complete six-month training dataset and used to predict the complete six-month test dataset. For the parametric models, i.e. GNU, BiQ and MPR, the Coefficient of Variation (CV) of the estimated parameters, defined as the ratio of the standard deviation of the parameter to its absolute value, was always less than 1%.

The training and test RMSEs are reported in Table 3.1, see also Fig. 3.9 for a graphical illustration. Note that in the MPR model COP depends on the triplet (Q_e, T_{ci}, T_{ei}). Given the weak dependence of T_{ei}, hereafter the 3-dimensional plots display the surface $COP(Q_e, T_{ci}, \bar{T}_{ei})$, where $\bar{T}_{ei}$ is the average value of T_{ei} computed on the overall dataset.

On this rich dataset, the worst predictive performance is that of the GNU model. The other four models guarantee comparable results, the GP being only marginally better than the others. The inferiority of the GNU model is explained by its limited flexibility - it has just three parameters - compared to the other models. Recall that BiQ and MPR models have nine and

ten parameters, respectively. Since the MLP ad GP models were trained via a regularization method, their complexity must be measured in terms of effective degrees of freedom [28], which result to be 12.7 and 19 respectively.

TABLE 3.1: Six months - RMSE on test data

	Train	Test
GNU	0.1955	**0.1995**
BiQ	0.0785	**0.0815**
MPR	0.0915	**0.0938**
MLP	0.0774	**0.0807**
GP	0.0744	**0.0781**

3.4.2 Models Robustness: covariate shift experiment

The monthly data were used to run an experiment whose goal is to assess the robustness of the considered models against covariate shifts. For a given month, e.g. May, all the five models were trained using only the training data collected during that month. The identified models were then used to predict the COP of all the test data collected from June to October, and the variance of the prediction error computed. The procedure was repeated for all the subsequent months: (i) train the models on one-month data and (ii) predict test data of the other five months.

For the GNU and BiQ models, the maximum CV of the estimated parameters was always below 15%, while higher CVs were observed for the MPR model, especially in the September and October datasets, the worst value being the September one (80%).

Table 3.2 reports the RMSEs both in training and test for each month. The three-dimensional plots are reported in Figures 3.4-3.8. The BiQ (Figure 3.5) and GP models (Figure 3.8) represent the worst and best performers, respectively.

It is interesting to note that the models behave much differently from each other. The GNU model, for instance, is often the worst performer, but it is relatively robust in the face of covariate shifts and its 6-fold RMSE (0.2164) is comparable to its global RMSE (0.1995).

The behavior of the BiQ (Figure 3.5) and MPR models is far less satisfactory: they perform so badly in October that their 6-fold RMSE (0.4415 and 0.3354, respectively) is definitely worse than GNU's one. When, as in October, the covariate shift renders the training data poorly informative, the increased flexibility guaranteed by the additional parameters turns into

a drawback, because the extra degrees of freedom are spent to overfit in a small region at the expense of generalization capabilities.

The MLP ANN proves acceptably robust: its 6-fold RMSE is 0.1347, the same order of magnitude of its global RMSE, equal to 0.0807. Finally, the GP model (Figure 3.8) appears rather robust with a 6-fold RMSE (0.0981) only marginally larger that its global RMSE (0.0781).

The results are suggestive of a possible beneficial effect of regularization (used in both the MLP and GP models) on the robustness against covariate shifts. As a matter of fact, it is known that regularization matches model flexibility to data informativeness thus helping to prevent overfitting [28].

TABLE 3.2: Covariate shift experiment - 6-fold RMSE

	May		**June**		**July**		**Aug**		**Sep**		**Oct**		**RMSE**
	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test	6-fold.
GNU	0.1640	0.2085	0.1760	0.2782	0.2083	0.1915	0.1851	0.2043	0.1891	0.2031	0.1432	0.2125	**0.2164**
BiQ	0.0800	0.0942	0.0867	0.0839	0.0720	0.0920	0.0735	0.0881	0.0664	0.0859	0.0612	2.2046	**0.4415**
MPR	0.0873	0.0992	0.0944	0.1096	0.0836	0.1025	0.0792	0.1237	0.0703	0.1123	0.0670	1.4651	**0.3354**
MLP	0.0799	0.0931	0.0818	0.0814	0.0705	0.0908	0.0722	0.0835	0.0654	0.0923	0.0599	0.3669	**0.1347**
GP	0.0792	0.0911	0.0818	0.0823	0.0703	0.0913	0.0720	0.0967	0.0659	0.0941	0.0621	0.1328	**0.0981**

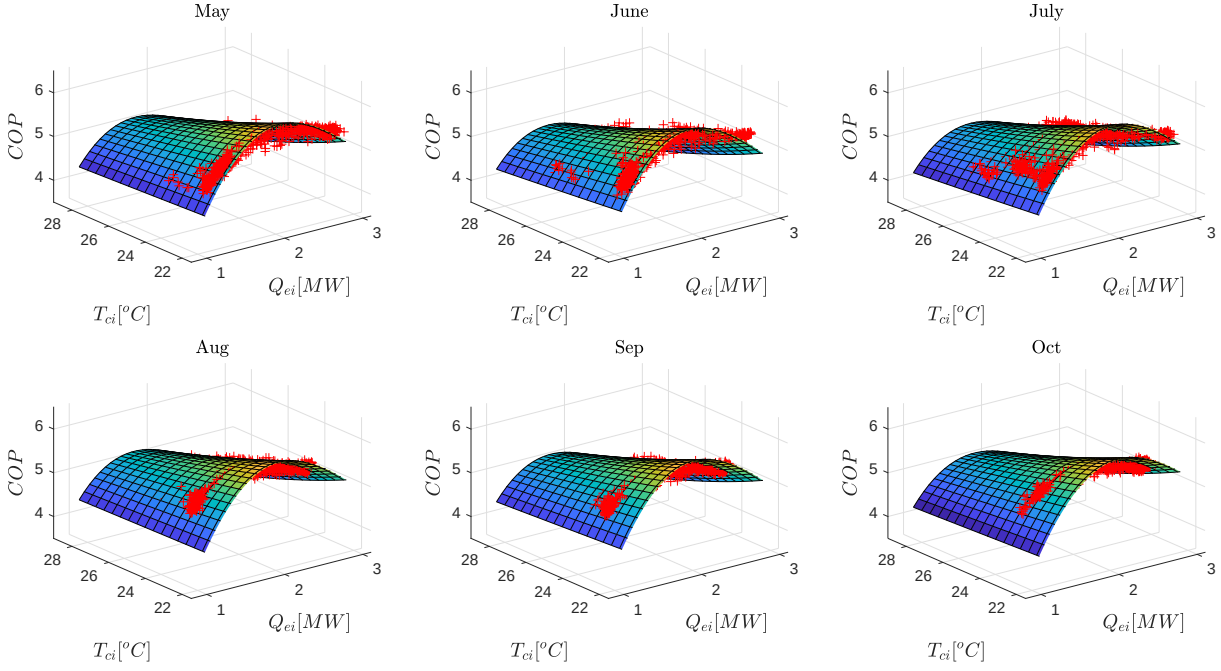

FIGURE 3.4: Gordon-Ng Universal model under covariate shift. Surface: prediction model; crosses: test data.

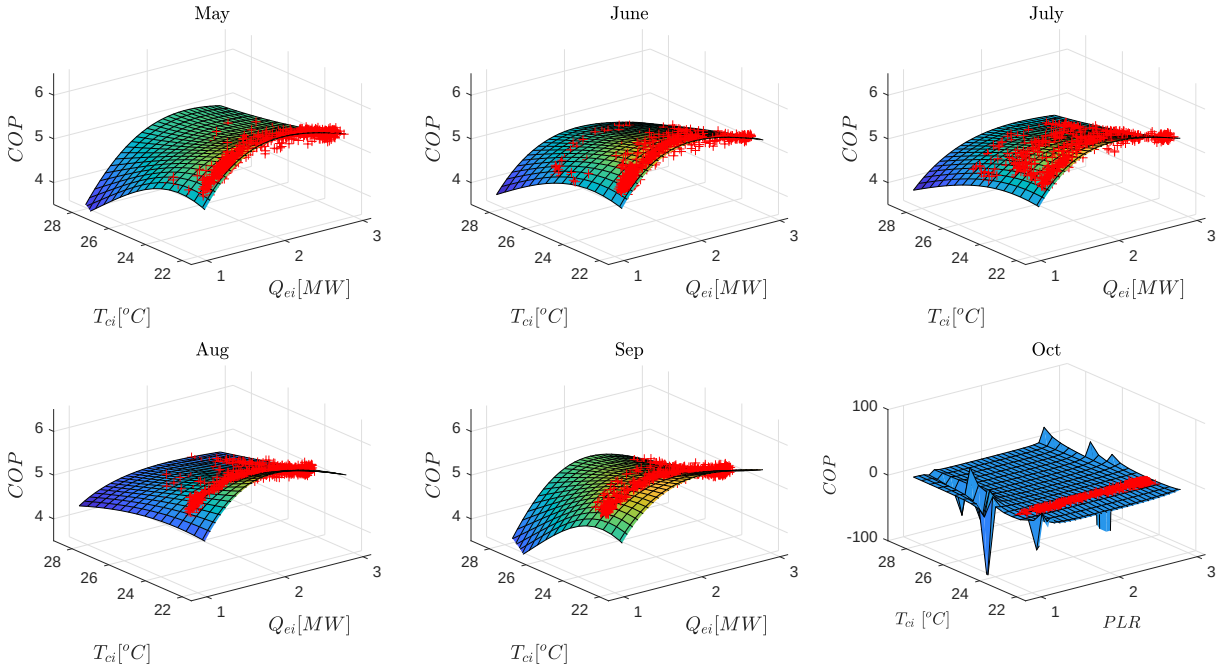

FIGURE 3.5: Bi-quadratic model under covariate shift. Surface: prediction model; crosses: test data.

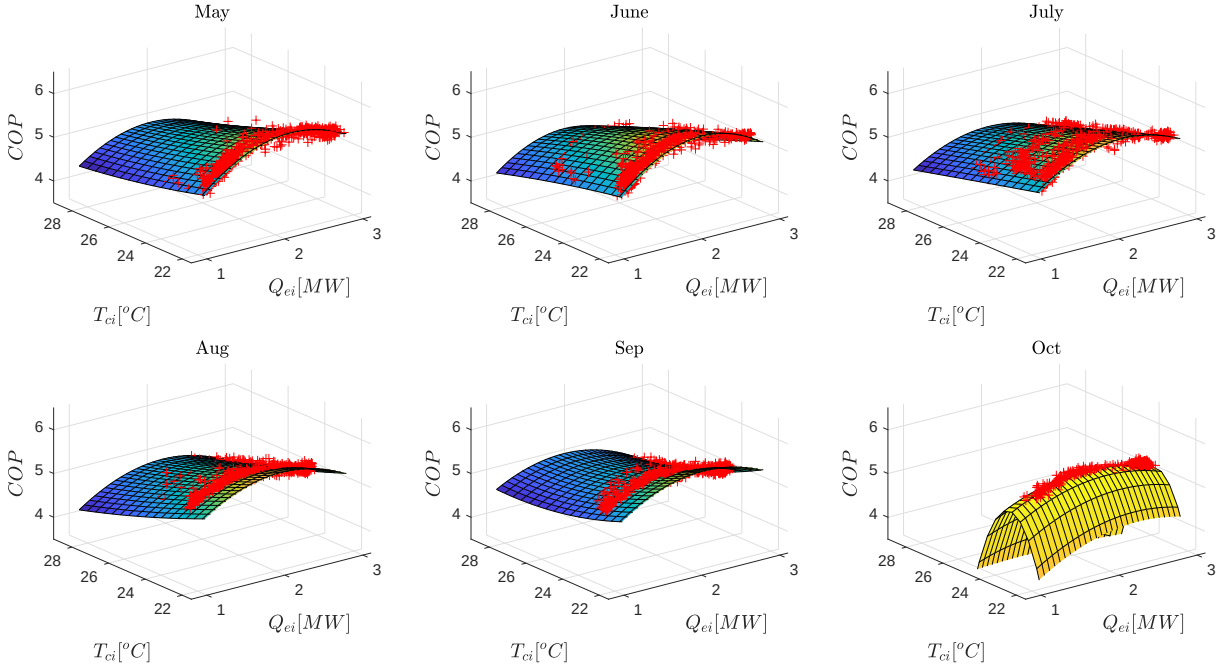

FIGURE 3.6: Multivariate Polynomial regression model under covariate shift. Surface: prediction model; crosses: test data.

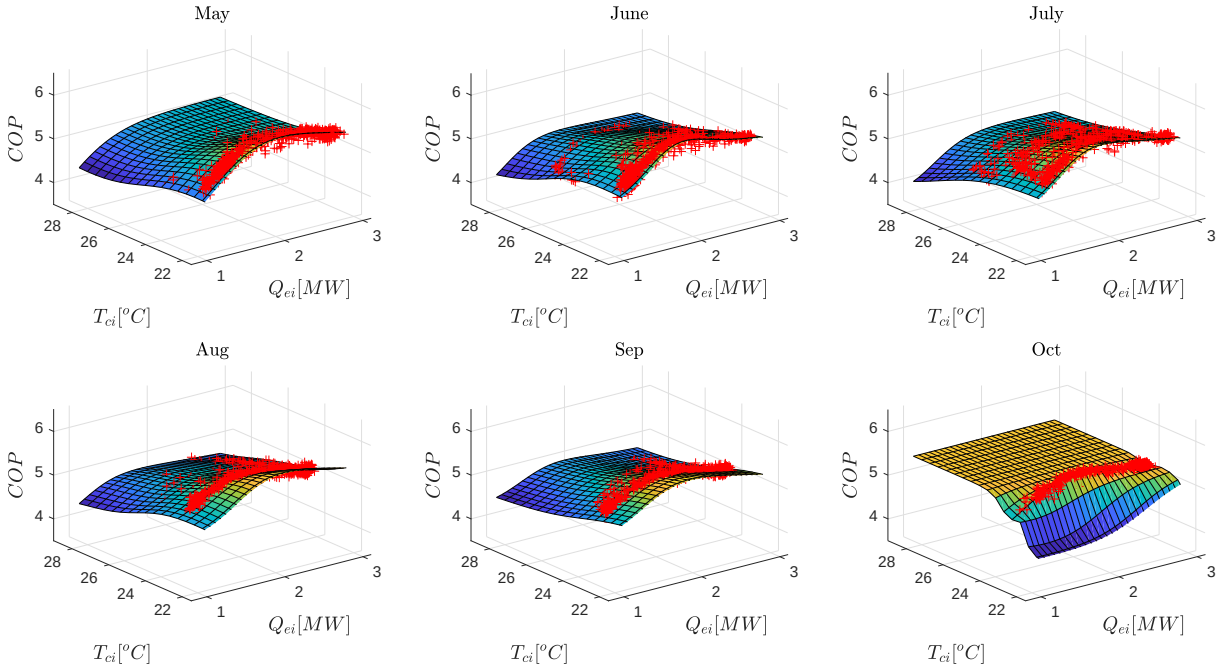

FIGURE 3.7: Multilayer Neural Network model under covariate shift. Surface: prediction model; crosses: test data.

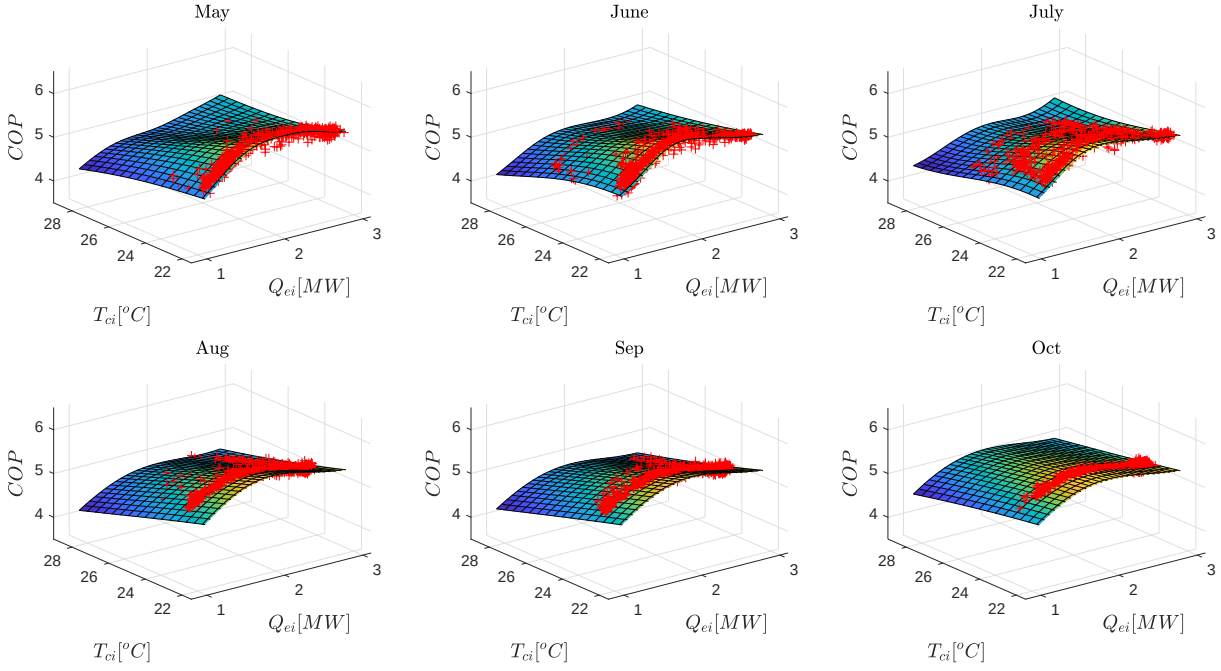

FIGURE 3.8: Gaussian Processes model under covariate shift. Surface: prediction model; crosses: test data.

3.5 Discussion

The identification from field data of energy efficiency models for water-cooled centrifugal chiller has been addressed. Two main issues were considered: (i) model comparison on extensive field data of four models taken from the literature and a newly proposed technique based on the Gaussian Process framework; (ii) covariate shift experiment: assessing the robustness of the five models in the face of shifts of the covariate distribution that may occur when a chiller is operated in real-world conditions. Assessments and comparisons were conducted on six-month field data collected within the activities of the EU ECSEL Project SemI40. To the authors' knowledge, previous works relied on far less numerous datasets.

The model comparison confirms that black-box methods ensure a better fit compared to the semi-empirical GNU model: the price to pay for physical meaningfulness of GNU model parameters is an increased RMSE on test data with respect to more flexible black box techniques. In a rich-data context, were both the training and test data share the same covariate distribution, the BiQ, MP, MLP and GP models achieved comparable performances.

The picture changes, however, when the outcomes of the covariate shift experiment are considered, see fig. 3.9). Indeed, the availability of a six-month dataset allowed a rather unique experiment: using one-month data to train the model and then test its performance on the remaining five months. The results demonstrated that the covariate shift may be a critical issue for chiller models identified from field data. In fact, substantial variations of the model quality were observed across the months. In particular, the predictive performance of the GNU model remained stable, while the performances of empirical models like BiQ and MPR deteriorated to the point of becoming worse than GNU's ones.

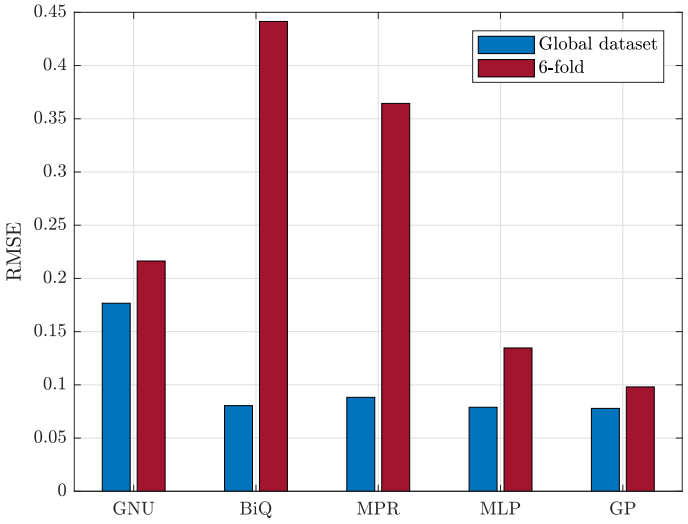

FIGURE 3.9: The average RMSE values of the five models

The moral of the story is that caution should be used before abandoning a parsimonious and physical-base model as the GNU one in favor of more flexible black box models, already proposed in the literature, such as BiQ, MP and MLP models. While in ideal data-rich contexts the GNU model is likely to be outperformed, the alternatives may lack the necessary robustness when trained on real-world data whose coverage of all possible operating conditions is far from being optimal. In these cases, the GNU model proved resilient, while the alternatives yielded poor generalization properties in scarcely sampled regions.

On the other hand, machine learning approaches equipped with regularization penalties, appear to be an appealing alternative to the GNU model. Not only they achieve smaller prediction errors but, especially the GP model, they are also robust to covariate shifts. This suggests that Gaussian Processes may represent an innovative and reliable approach to accurately model chiller efficiency in an Industry 4.0 setting, where models are directly obtained and updated from continuously monitored field data.

4

Calibration strategies for the Gordon-Ng chiller model

In the literature, chiller models are typically classified in two categories: empirical and semi-empirical. Semi-empirical models exploit the physical knowledge of the process in order to specify the functional structure of the mathematical relationships between the variables. It traces back to more than twenty years ago the publication of the Gordon and Ng model [9], a semi-empirical model for reciprocating chillers, that was later extended also to other chiller types [19, 27]. Based on first principles of thermodynamics and heat losses linearization, a low-order model is obtained whose unknown parameters admit a physical interpretation in terms of internal entropy production, heat losses (or gains), and heat exchanger thermal resistance.

While the simplicity of the GNU (Gordon-Ng Universal) model limits its flexibility, the need of calibrating only a few parameters that are physically meaningful has its advantages. First, the model is less prone to overparametrization issues that can arise when data are scarce or suffer from collinearity. Furthermore, in view of the physical interpretability of the parameters, the model is well suited to be used in the development of FDD (Fault Dectection and Diagnosis systems) [18, 30, 31]. As observed in [32], this kind of application calls for realistic estimates of the statistical distribution of the monitored parameters, especially when fault detection has to rely on control chart techniques. It was this observation that motivated Andersen and Reddy [32] to investigate whether the widely used OLS (Ordinary Least Squares) method was the most appropriate way to estimate the GNU parameters.

In fact, OLS provides unbiased parameter estimates and realistic confidence intervals only if the model regressors are error-free, a circumstance that hardly occurs for the GNU model, whose regressors depend on observed variables subject to measurement noise. Using a public precise chiller performance dataset [33], Andersen and Reddy not only showed that OLS produced

biased estimates but demonstrated that recasting the problem within the EIV (Errors in Variables) framework, the bias could be substantially reduced. They also observed that the bias issue could be of minor concern if the scope was chiller efficiency prediction, but they did not further investigate if and how the predictive performance of the model was influenced by the choice of either EIV or OLS as estimation procedure.

The present chapter focuses on the calibration of the GNU model. In particular, it addresses several unanswered questions. Given that the EIV approach achieves bias reduction at the cost of increasing the variance, is it really preferable to OLS? If the final goal is not parameter estimation, but the prediction of the Coefficient of Performance (COP), how does OLS compare with EIV? And what is the most appropriate calibration method, under a statistical viewpoint? Finally, is the added complexity of a statistically rigorous approach such as Nonlinear Least Squares (NLS), really worth the potential improvements in COP prediction? In order to answer these questions, three estimation methods, OLS, EIV and NLS, are tested on two benchmarks: a small-sized dataset already used by Andersen and Reddy [33], and an ASHRAE dataset [34], whose sample size is definitely greater.

The chapter is organized as follows. In Section 4.1, the Gordon-Ng Universal chiller model is concisely reviewed. In Section 4.2, its calibration from experimental data is discussed: two existing approaches, OLS and EIV, are reviewed and a third nonlinear least squares method is introduced. The two case studies are illustrated in Section 4.3 and the results discussed in Section 4.4. Some concluding remarks (Section 4.5) end the chapter.

Glossary

- Q_e: evaporator cooling capacity [kW]
- T_{ci}: condenser inlet water temperature [K]
- T_{ei}: evaporator inlet water temperature [K]
- P: electrical consumption [kW]
- $COP = Q_e/P$: Coefficient of Performance

4.1 The GNU model of chiller perfomance

The Gordon-Ng Universal (GNU) model is a semi-empirical model whose structure is based on first principles of thermodynamics and linearized heat losses. It is summarized by the following

relationship, linking together COP, Q_e, T_{ci}, T_{ei}:

$$\frac{T_{ei}}{T_{ci}}\left(1+\frac{1}{COP}\right)-1=\beta_1\frac{T_{ei}}{Q_e}+\beta_2\frac{T_{ci}-T_{ei}}{T_{ci}Q_e}+\beta_3\frac{Q_e}{T_{ci}}\left(1+\frac{1}{COP}\right) \tag{4.1}$$

All the three model parameters β_i admit a physical interpretation [19]:

- β_1: the total internal entropy production rate in the chiller due to irreversibilities;
- β_2: the rate of heat losses (or gains) from (or into) the chiller;
- β_3: the total heat exchanger thermal resistance.

4.2 Calibrating the GNU model

4.2.1 OLS estimation

Consider the GNU model (4.1). Let $\beta = \begin{bmatrix} \beta_1 & \beta_2 & \beta_3 \end{bmatrix}^T$ denote the vector of unknown parameters and define

$$x = \begin{bmatrix} \frac{T_{ei}}{Q_e} & \frac{T_{ci}-T_{ei}}{T_{ci}Q_e} & \frac{Q_e}{T_{ci}}\left(1+\frac{1}{COP}\right) \end{bmatrix} \in \mathbb{R}^{1\times 3} \tag{4.2}$$

$$y = \frac{T_{ei}}{T_{ci}}\left(1+\frac{1}{COP}\right)-1 \tag{4.3}$$

With this notation, eq. (4.1) admits the following linear-in-parameter formulation [19, 23, 27]:

$$y = x\beta \tag{4.4}$$

Given n experimental samples $COP(i)$, $Q_e(i)$, $T_{ci}(i)$, $T_{ei}(i)$, $i = 1,\dots,n$, we can form the vectors

$$Y = \begin{bmatrix} y(1) \\ y(2) \\ \vdots \\ y(n) \end{bmatrix}$$

$$X = \begin{bmatrix} x(1) \\ x(2) \\ \vdots \\ x(n) \end{bmatrix} = \begin{bmatrix} x_1(1) & x_2(1) & x_3(1) \\ x_1(2) & x_2(2) & x_3(2) \\ \cdots & \cdots & \cdots \\ x_1(n) & x_2(n) & x_3(n) \end{bmatrix}$$

so that

$$Y = X\beta$$

The ordinary least square estimate β^{OLS} is just given by

$$\beta^{OLS} = \left(X^T X\right)^{-1} X^T Y$$

In view of its simplicity, this represents the standard approach for calibrating the GNU model.

4.2.2 EIV estimation

In [32], it was observed that the choice of the estimation approach should depend on the prospective use of the model. In particular, Andersen and Reddy observed that in some contexts the accuracy of model parameters is particularly important and that this is most likely to happen when they have to be interpreted physically, a notable example being that of fault detection and diagnosis (FDD).

The OLS estimate β^{OLS} is a biased estimator of the true β, if the regressor matrix X is subject to statistical errors. This is exactly what happens for the GNU model, since the regressor vector $x(i)$ is a function of the measured $COP(i)$, $Q_e(i)$, $T_{ci}(i)$, $T_{ei}(i)$, that are obviously affected by measurement errors. Regression models that account for measurement errors in the independent variables go under the name of Errors in Variables (EIV) models [35–37]. For the GNU model, Andersen and Reddy suggested the use of the following corrected least squares estimator

$$\beta^{EIV} = \left(X^T X - S_{xx}^2\right)^{-1} \left(X^T Y - S_{xy}^2\right)$$

where $S_{xx}^2 \in \mathbb{R}^{3\times 3}$ is the covariance of the measurements error and $S_{xy}^2 \in \mathbb{R}^{3\times 1}$ is the covariance between the regression variables and the dependent variable. Three methods to estimate S_{xx}^2 and S_{xy}^2 are presented in [32]: propagation of errors, stochastic calculus, and stochastic simulation. In the present chapter, the third approach is adopted.

By injecting pseudorandom noise of variable amplitude into precise chiller data published by Ng et al. [33], Andersen and Reddy demonstrated that: (i) GNU parameters estimated via OLS are biased; (ii) the bias increases with larger measurement errors; (iii) EIV estimation yields unbiased estimates, suitable for FDD applications.

An important point is that the improvement on OLS estimates achieved by using EIV estimation is specific to the goal of estimating the physically interpretable parameters β_i. As recognized by Andersen and Reddy, "If the purpose of the model is purely predictive, then a model determined by OLS is probably as good as the EIV model for the mean predictive value". As far as one is interested in predicting COP values, the EIV approach may not be the most appropriate one. Next, we are going to show that neither OLS is a statistically correct solution when the purpose is "purely predictive'.

4.2.3 NLS estimation

Consider a target variable y and a vector of independent variables $\xi \in \mathbb{R}^m$. In order to predict y given ξ, we consider the model

$$\hat{y} = g(\xi, \beta) \tag{4.5}$$

where $\beta \in \mathbb{R}^q$ is a vector of model parameters. The predictive value of the model can be measured by the so called *risk*, for instance the expectation of a quadratic loss function:

$$R(\beta) = \mathbb{E}[L(\beta)], \qquad L(\beta) = (y - g(\xi, \beta))^2$$

where the pair (ξ, y) is randomly drawn from an unknown joint distribution $p(\xi, y)$. When n training pairs $(\xi(i), y(i))$, $i = 1, \ldots, n$, are available (we assume that $(\xi(i), y(i))$ are i.i.d. samples drawn from the joint distribution $p(\xi, y)$) the risk can be approximated by the *empirical risk*

$$R_{\text{emp}}(\beta) = \frac{1}{n} \sum_{i=1}^{n} (y(i) - g(\xi(i), \beta))^2$$

Let

$$\beta^* = \arg\min_{\beta} R_{\text{emp}}(\beta) \tag{4.6}$$

denote the parameter vector that minimizes the empirical risk. Note that β^* coincides with the maximum likelihood estimator under the statistical model

$$y(i) = g(\xi(i), \beta) + e(i), \qquad e(i) \sim \mathcal{N}(0, \sigma^2)$$

where the noises $e(i)$ are assumed indepedent of each other.

Coming back to the GNU model, observe that the OLS approach does not use COP as target variable, but y defined in (4.3). However, if the final goal is predicting the coefficient of performance, it is COP that should be used as target variable. Therefore, the first step is solving the GNU model (4.1) for COP:

$$COP = \frac{\frac{T_{ei}}{T_{ci}} - \beta_3 \frac{Q_e}{T_{ci}}}{1 - \frac{T_{ei}}{T_{ci}} + \beta_1 \frac{T_{ei}}{Q_e} + \beta_2 \frac{T_{ci} - T_{ei}}{T_{ci} Q_e} + \beta_3 \frac{Q_e}{T_{ci}}} \tag{4.7}$$

In order to reformulate in terms of eq. (4.5), define $y_{COP} = COP$, $\xi_1 = T_{ei}$, $\xi_2 = T_{ci}$, $\xi_3 = Q_e$. Then, we obtain:

$$\hat{y}_{COP} = g(\xi, \beta) = \frac{\frac{\xi_1}{\xi_2} - \beta_3 \frac{\xi_3}{\xi_2}}{1 - \frac{\xi_1}{\xi_2} + \beta_1 \frac{\xi_1}{\xi_3} + \beta_2 \frac{\xi_2 - \xi_1}{\xi_2 \xi_3} + \beta_3 \frac{\xi_3}{\xi_2}} \tag{4.8}$$

Therefore, if the purpose is prediction, the correct statistical procedure is finding the solution $\beta^{NLS} = \beta^*$ of the nonlinear least squares problem (4.6). Of course, compared to the OLS

method, the NLS approach is less straightforward, so that the question arises of whether the game is worth the candle.

4.3 Data and methods

In order to assess if and how much the choice of the parameter estimation method affects the predictive properties of the GNU model, two independent sets of experimental data were considered.

4.3.1 Case Study 1

The first data set, taken from [33], consists of 30 measurements from a reciprocating chiller, belonging to an experimental refrigerant circuit designed for laboratory tests. The chiller is water-cooled with R12 as refrigerant and nominal cooling capacity of 10.5 kW. Table 1 in [33] reports the experimental measurements of Q_e, T_{ei}, T_{ci}, P and their accuracies. Letting $COP = Q_e/P$, it is possible to plot the triplets $(T_{ci}(i), Q_e(i), COP(i))$, $i = 1, \ldots, 30$ as points in a three-dimensional space, see Fig. (4.1).

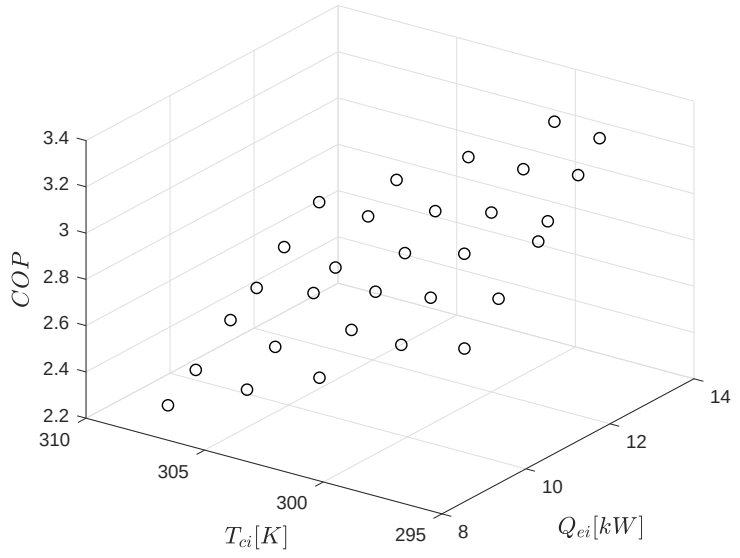

FIGURE 4.1: Data set #1: three-dimensional representation of the COP as a function of measured chiller data.

It is apparent that this high-precision dataset collected under a design-of-experiment protocol, is far from real-world field measurements which are typically affected by higher uncertainties. These data were used as starting point to simulate a range of more realistic scenarios in which the original (almost) "noise-free" measurements are blurred by normally distributed random

numbers with different standard deviations. Following [32], standard deviations were obtained as $\sigma_a(k) = w(k)m_a$ where $w(k) = 0.1k$, $k = 0, \ldots, 20$, and m_a, $a \in \{Q_e, T_{ei}, T_{ci}, P\}$ denotes the uncertainty magnitude specific to each variable, see Table 1 of [32]. Hence, $w(k) = 0$ would imply that no noise has been added to the observations, while $w(k) = 1.5$ would imply that for each of the four variables the noise σ_a is 1.5 times larger than the respective uncertainty m_a. The noise was added only to the four independent variables, while the blurred COP was just computed as $y = COP = Q_e/P$, where Q_e and P stand for the blurred values.

For each of the 21 error weights $w(k)$, the assessment was carried out on 1000 synthetic datasets, according to the following procedure:

1. For $i = 1, ..., 1000$:

   (a) *Creation of a new synthetic dataset*: a noisy dataset is created by blurring the original data with Gaussian noise such that $\sigma_a(k) = w(k)m_a$.

   (b) *Random Data splitting*: the dataset is randomly partitioned in two sets: 21 data for training and 9 for testing.

   (c) *Parameters estimation*: $\beta^{OLS}(i)$, $\beta^{EIV}(i)$, $\beta^{NLS}(i)$ are estimated from the training set. The NLS estimation has been performed by means of the function *nlinfit* of the Matlab Statistics and Machine Learning Toolbox.

   (d) *COP prediction*: on the test set, three alternative predictions $\hat{y}_{COP}(i)$, $i = 1, \ldots, 21$ are computed by plugging $\beta^{OLS}(i)$, $\beta^{EIV}(i)$, $\beta^{NLS}(i)$ into (7).

   (e) *RMSE Testing*: predictive performance is measured by the Root Mean Squared error computed on the test set:

   $$RMSE^{\alpha} = \sqrt{\frac{\sum_{i=1}^{9}(y(i) - \hat{y}_{COP}(i))^2}{N}}$$

   where $\alpha \in \{OLS, EIV, NLS\}$, depending on the set of estimated parameters used to compute $\hat{y}_{COP}$.

2. *Median and variability limits*: Based on the $1,000$ simulations, compute the median and also the 2.5th and the 97.5th percentiles of the parameters β_i^{α} and $RMSE^{\alpha}$, $\alpha \in \{OLS, EIV, NLS\}$.

This benchmark replicates the one used in [32] in order to compare the biasedness of the OLS and EIV estimators. The only difference is the splitting of the synthetic datasets into training and test sets, a choice dictated by the need of evaluating COP predictions, while Andersen and Reddy were concerned only with the accuracy of parameters estimates.

4.3.2 Case Study 2

The second data set comes from the ASHRAE research project 1043-RP [34] whose purpose was to study faults in chillers and generate data for Fault Detection and Diagnostic (FDD) applications. A centrifugal chiller with cooling capacity of 316 kW was tested at 27 different operating states under both normal and faulty conditions. Since we are not concerned with FDD, only the 432 measurements in steady-state condition with no faults in the system were used in our study. In Fig. 4.2 a three-dimensional plot of the data is displayed.

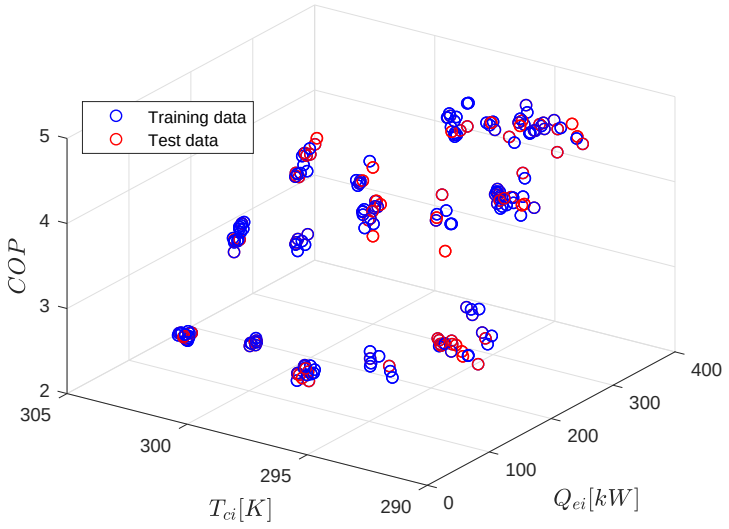

FIGURE 4.2: Data set #2: three-dimensional representation of the COP as a function of measured chiller data. Blue dots: training data; Red dots: test data.

Differently from the Data set #1, where only 30 samples were available, the data rich context dispenses from the need of injecting simulated noise in order to create several benchmark datasets. Rather, the distribution of the $RMSE$ is assessed by a bootstrap procedure [38, 39], according to the following steps

1. *Random Data splitting*: the whole dataset is randomly split in a training and test set consisting of 302 and 130 samples, respectively.

2. *Bootstrap*: For $i = 1$ to 1000

   (a) *Bootstrap extraction:* a bootstrapped set is created by randomly extracting 302 samples *with replacement* from the 302 samples of the training set.

   (b) *Parameters estimation*: $\beta^{OLS}(i)$ and $\beta^{NLS}(i)$ are estimated from the training set.

   (c) *COP prediction*: on the test set, two alternative predictions $\hat{y}_{COP}(i)$, $i = 1, \ldots, 21$ are computed by plugging $\beta^{OLS}(i)$ or $\beta^{NLS}(i)$ into (7).

(d) *Testing*: predictive performance is measured by $RMSE^{OLS}$ and $RMSE^{NLS}$ computed on both the training and test set.

3. *Distribution of RMSE*: kernel density estimation is used to obtain the probability density functions of $RMSE^{OLS}$ and $RMSE^{NLS}$ in training and validation.

The key point of the bootstrap procedure is that random sampling with replacement from 302 training samples is an effective surrogate of random sampling from the true joint distribution of the chiller data, that is obviously unknown.

An advantage with respect to Case Study #1 is that the assessment of the variability relies on the available data alone, without assuming detailed knowledge of the measurement uncertainty affecting each independent variable. Accordingly, we drop the EIV estimator whose implementation requires knowledge of the covariance matrices S_{xx} and S_{xy} that critically depend on measurement error statistics.

4.4 Results

4.4.1 Case Study 1

In order to assess the effect of the different estimators on the estimated parameters β_i, the median values and 2.5-th and 97.5-th percentiles are displayed as a function of the magnitude of the injected error, see Figs. (4.3), (4.4), (4.5).

For what concerns OLS and EIV estimation, the plots replicate rather faithfully Figs. (3)-(5) in [32]. The (small) differences are explained by our use of just 21 samples instead of 30. The stability of the results in spite of the sample size reduction confirms the limited complexity of the dataset. In accordance with [32], it is found that EIV estimation greatly reduce the bias. This is not without a cost, however, because, especially for large error magnitudes, the variance of β_i^{EIV} is substantially larger than that of the OLS and NLS estimates for all the three parameters. Conversely, no appreciable difference is observed between the distribution of β_i^{OLS} and β_i^{NLS}, $i = 1, \ldots, 3$.

In order, to ascertain whether the increased complexity of EIV regression is worth the candle, in Fig. (4.6) the percent RMSE for the three parameters and the three estimators is displayed as a function of the error magnitude. Although, the EIV estimate performs best, there is not a substantial improvement with respect to the OLS estimate that represents a very reasonable alternative, especially because it can do without the additional information needed to compute the covariance matrices S_{xx} and S_{xy}.

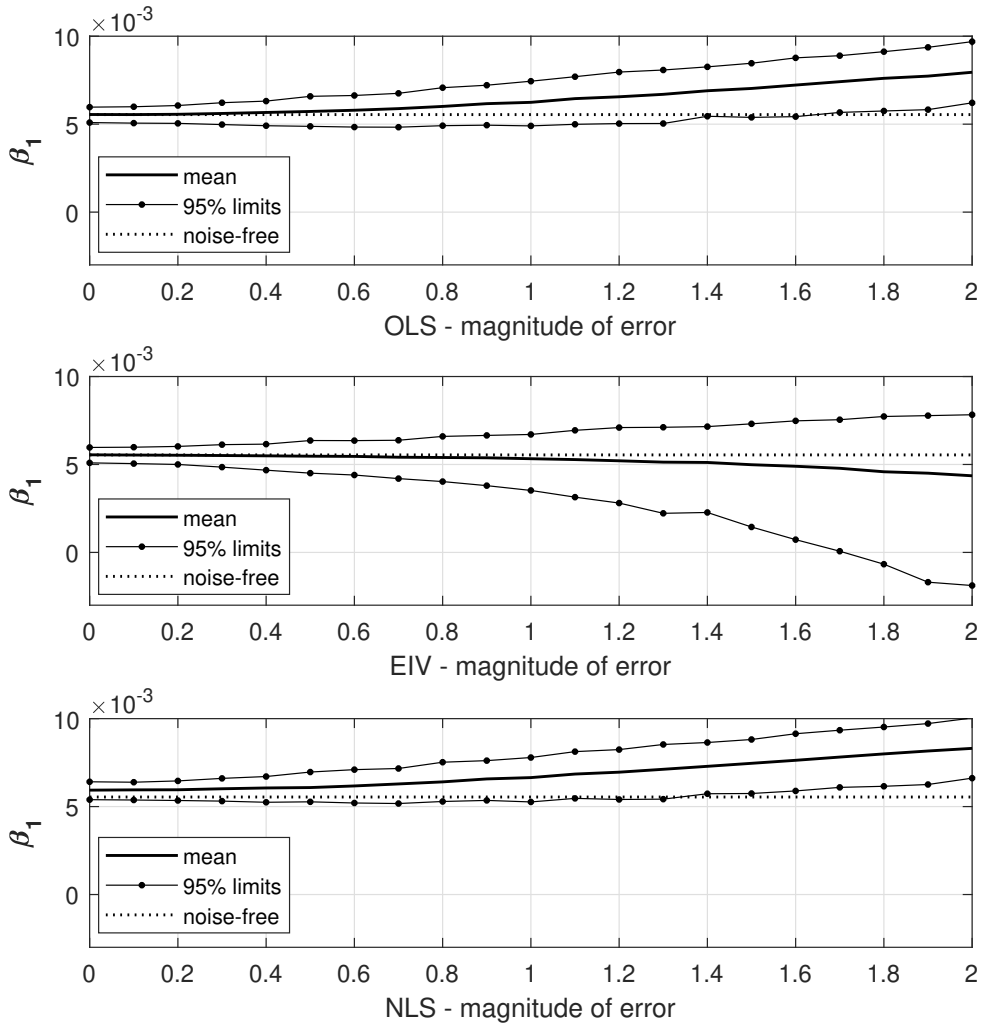

FIGURE 4.3: Comparison of OLS, EIV and NLS estimates of β_1 varying the noise magnitude of the measurements

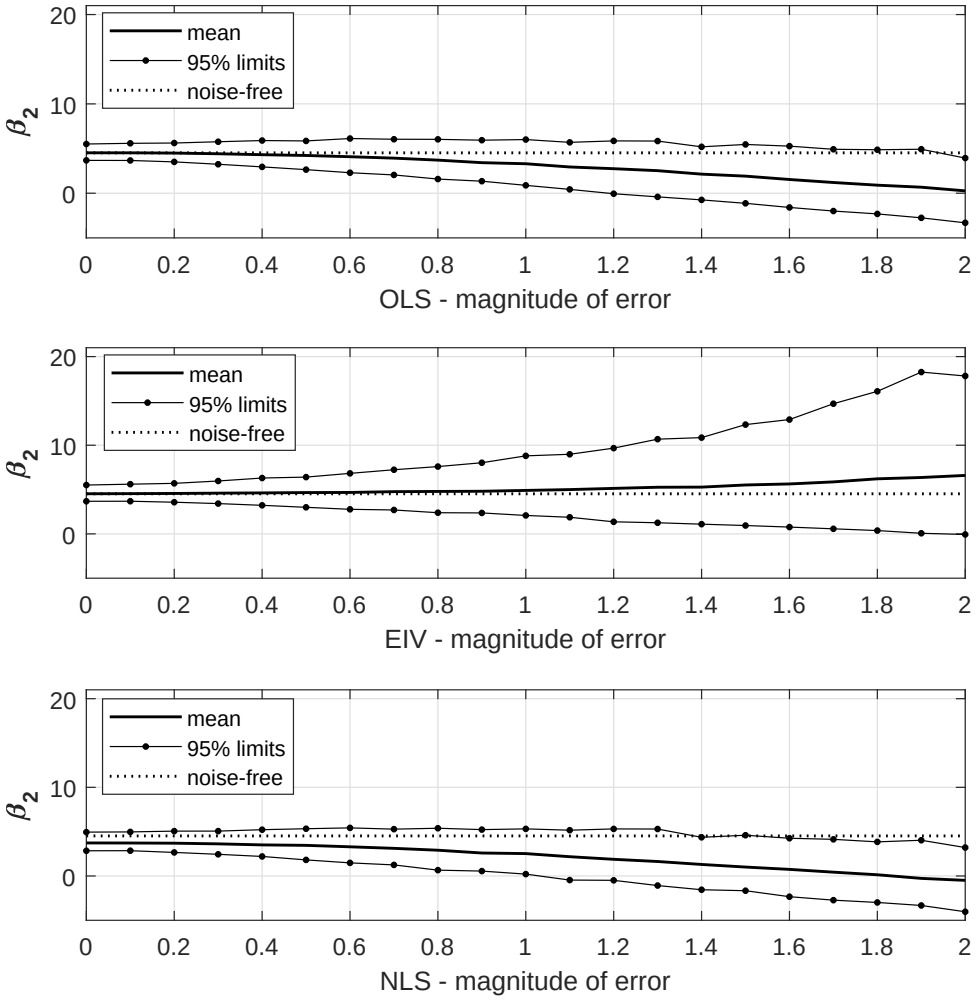

FIGURE 4.4: Comparison of OLS, EIV and NLS estimates of β_2 varying the noise magnitude of the measurements

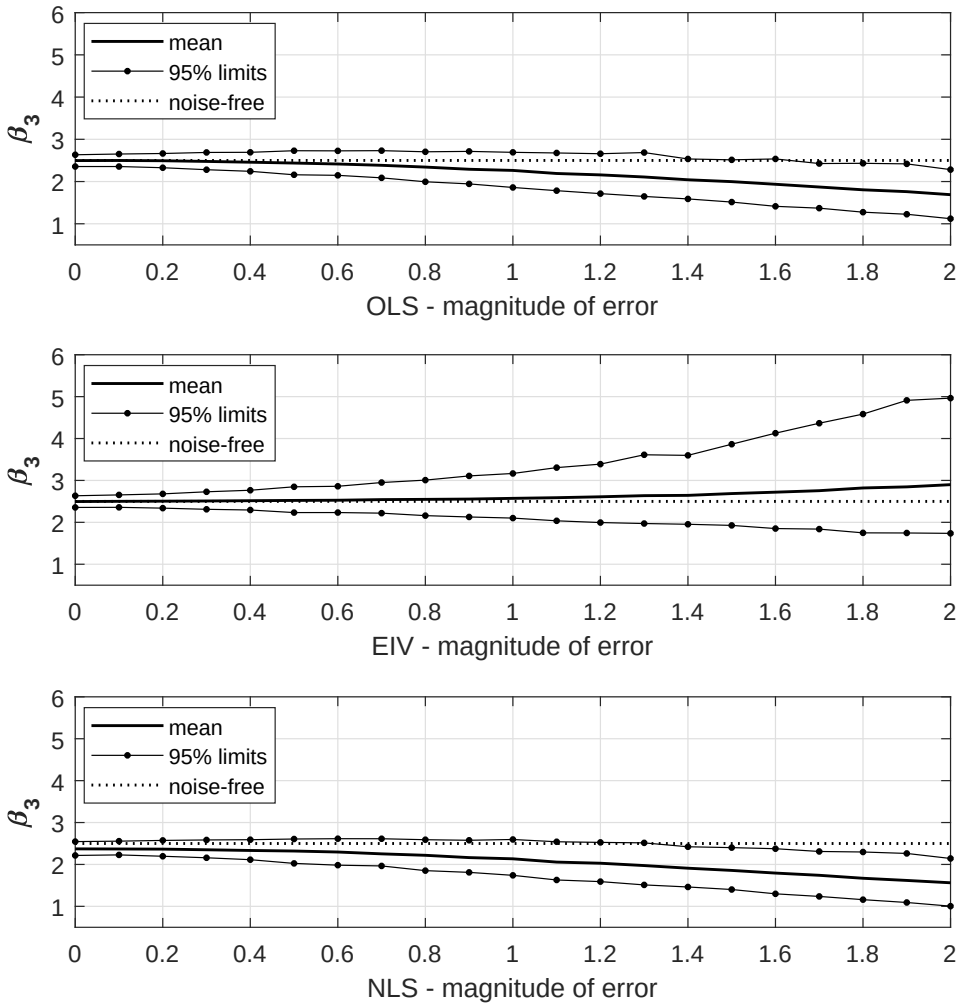

FIGURE 4.5: Comparison of OLS, EIV and NLS estimates of β_3 varying the noise magnitude of the measurements

In Fig. (4.7) the predictive performance of the three methods is compared in terms of median RMSE. The superiority of OLS and NLS over EIV is apparent. Again, no appreciable difference is observed between OLS and NLS. The percentiles of the distributions, are displayed in Fig. (4.8): it appears that EIV is worse not only in terms of median performance but suffers also from a larger variability, especially for large error magnitudes.

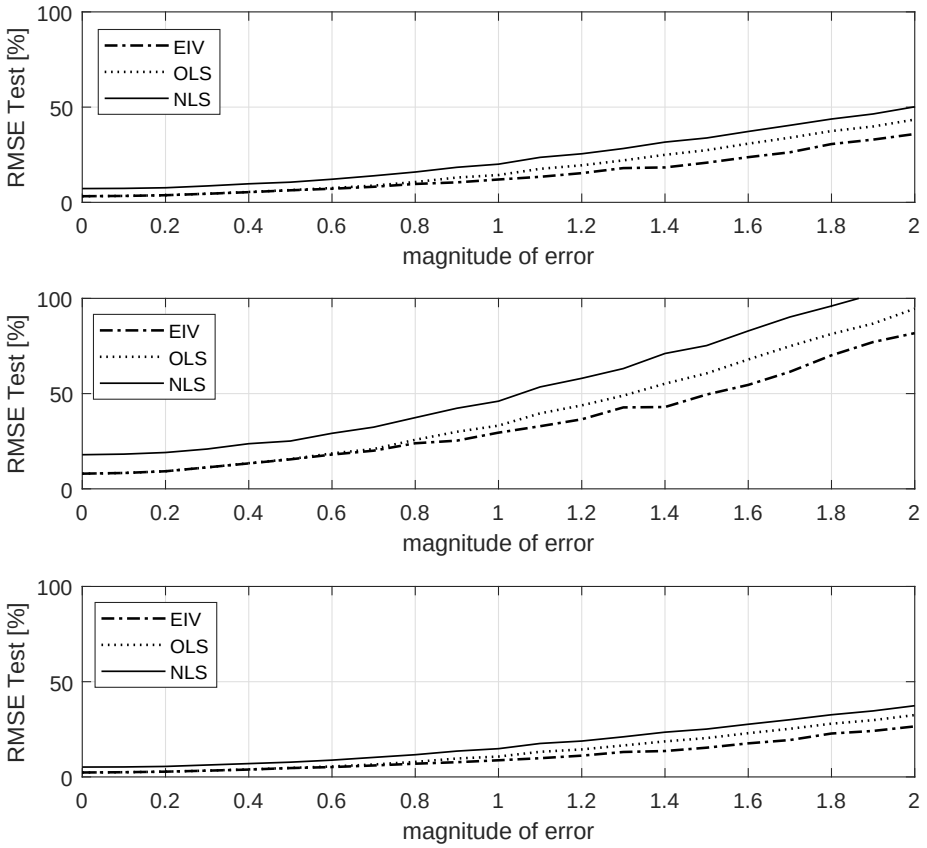

FIGURE 4.6: RMSE% of the GNU model parameters estimated by OLS, EIV and NLS methods. From top to bottom: β_1, β_2, β_3.

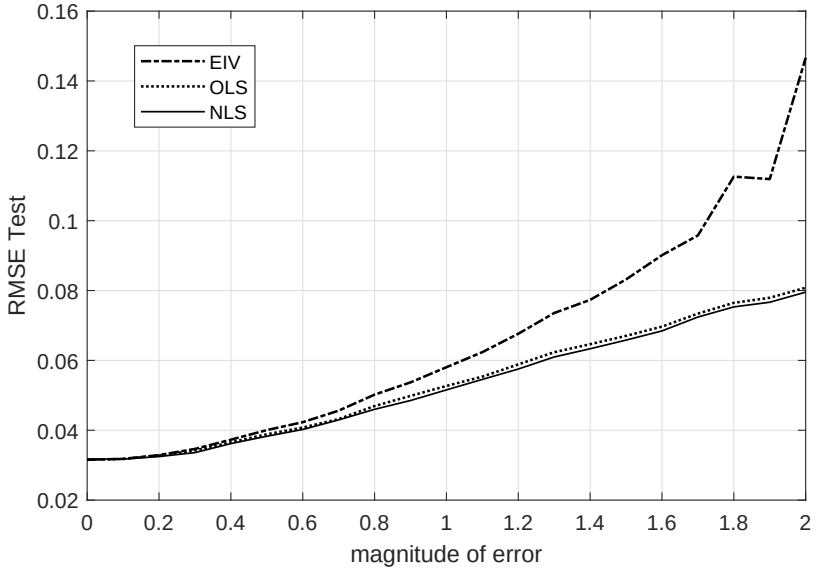

FIGURE 4.7: Comparison of the medians of $RMSE^{OLS}$, $RMSE^{EIV}$ and $RMSE^{NLS}$ on test data

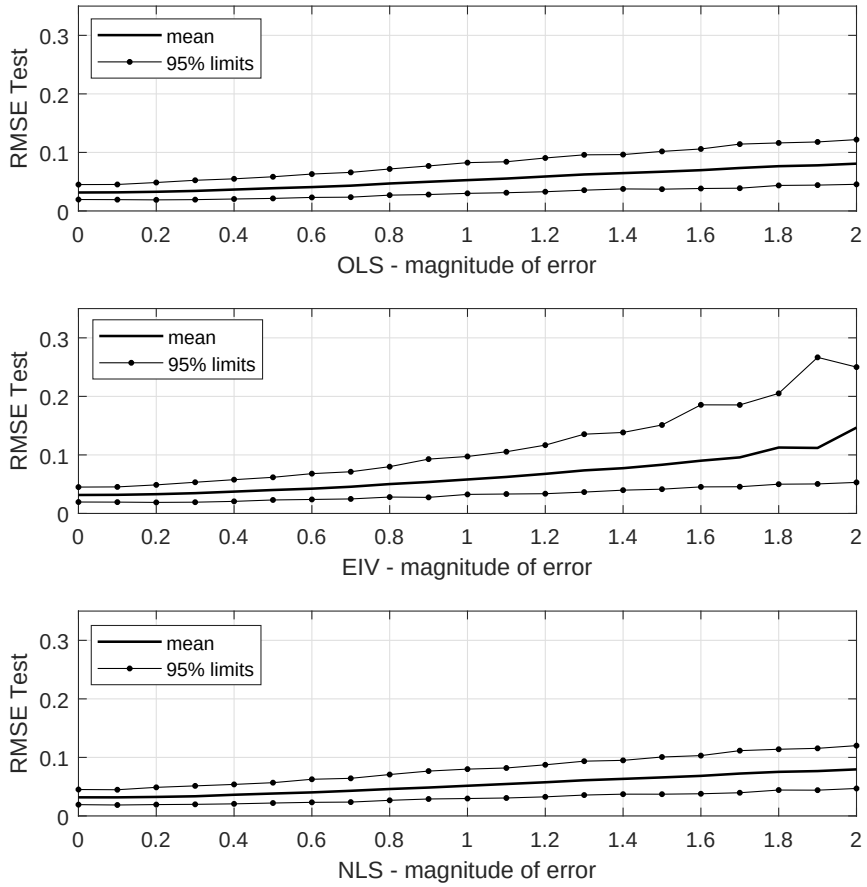

FIGURE 4.8: $RMSE^{OLS}$, $RMSE^{EIV}$ and $RMSE^{NLS}$ on test data. Solid line: median; Dash-dot line: variability limits

4.4.2 Case Study 2

The estimated densities of the training and test RMSEs are displayed in Fig. (4.9) for both the OLS and NLS estimators. A further visual comparison of the predictive performance on the test data is provided in Fig. (4.10) in terms of boxplots.

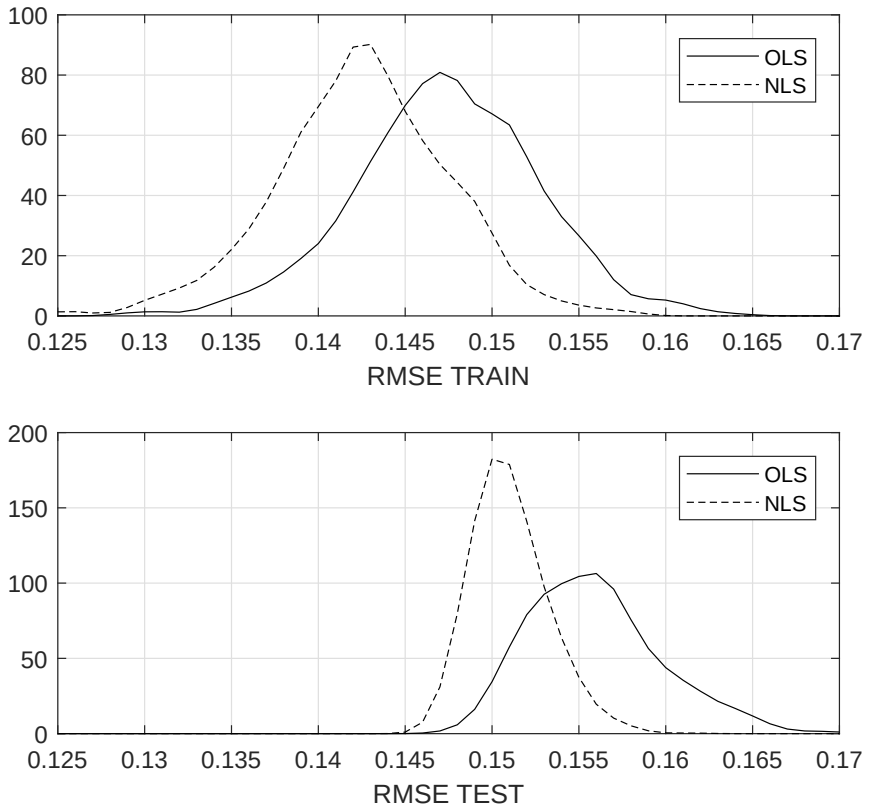

FIGURE 4.9: $RMSE^{OLS}$ and $RMSE^{NLS}$ probability density functions on training (top) and test data (bottom)

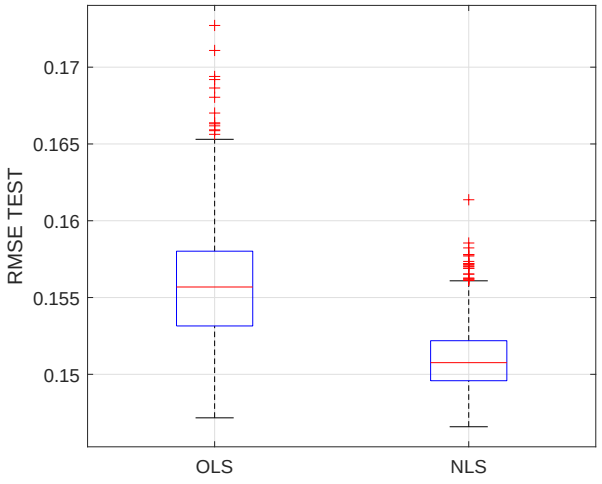

FIGURE 4.10: Box plots of $RMSE^{OLS}$ and $RMSE^{NLS}$ on test data

It appears that NLS yields smaller RMSEs both on training and test data. From Fig. (4.10) it is seen that the third quartile of $RMSE^{NLS}$ is below the first quartile of $RMSE^{OLS}$. Although

this might suggest that the difference is neat, its practical significance should not be overstated. Indeed, it suffices to compare the magnitude of the RMSEs with that of the observed COPs, displayed in Fig. (4.2), to draw the conclusion that the superiority of NLS estimation can hardly be of any practical significance as far as the goal is predicting the performance of the ASHRAE benchmark.

4.5 Concluding remarks

Although OLS represents the standard way to calibrate the Gordon-Ng Universal chiller model, it is known that it yields biased estimates of the parameters because it neglects the measurement noise affecting the independent variables. This deficiency had motivated the investigation of Errors-in-Variable (EIV) regression as an alternative to OLS when unbiased estimates of physically meaningful parameters are needed. Some questions, however, had remained unanswered: Even if EIV regression is less biased than the OLS one, is it really preferable, given its larger variance? How does EIV regression compares to OLS regression in terms of COP prediction? What is the statistically rigorous approach to COP prediction? Is the achievable improvement worth the cost of adopting a rigorous approach? This issues have been addressed with the help of two benchmark data sets: a high-precision dataset taken from [33] and the ASHRAE project 1043-RP dataset [34].

The analysis of the first case study has confirmed that the EIV approach reduces the bias of parameter estimates. At the same time, EIV seems to offer only a marginal improvement on OLS, also considering the need of estimating some covariance matrices whose calculation requires detailed information on the sensor noises that may not be routinely available. A second finding is that, as far as COP prediction is concerned, EIV regression is indeed inferior to OLS regression. Furthermore, it has been shown that a statistically rigorous approach to COP prediction can be formulated as a Nonlinear Least Squares problem. The predictive performance of the NLS approach was found to be indistinguishable from that of OLS in the first case study. In the second benchmark there is a statistically significant improvement that, however, does not appear to be of practical significance.

Although the present study suggests that a rigorous NLS approach may be of little use in practice, one should be aware of the limitations inherent to the two benchmarks. In fact, as seen from the three-dimensional plots, see Fig. (4.1) and (4.2), the COP data appear to be easily predictable by means of a low-complexity model. Hence, before a final word can be said, it may be worth probing the estimators on further benchmarks where COP prediction is more challenging.

5

Chiller Plant Optimization

The efficiency of HVAC systems is closely related to the efficiency of the chillers unit. Since a multiple chiller system typically consists of chillers having different capacities and performances, there are usually several chillers' part loads combinations able to satisfy the load demand. The problem of determining the load fraction that each chiller has to deliver in order to minimize the system power consumption and meanwhile satisfy the cooling load demand, is known as Optimal Chiller Loading (OCL) problem. In the last decade, several methods have been proposed to solve the OCL problem [40–56]. Chang et al. [40] assume that the chillers' efficiency is a quadratic function of the partialization and maximize the sum of the chillers efficiencies (i.e. the Coefficient of Performance) using the Lagrangian method. Of course, this is not equivalent to the "canonical" OCL problem, where system power consumption is instead minimized. In a second paper Chang [41] has addressed the canonical OCL problem again by the Langrangian multipliers methods, assuming that all the chiller power consumption curves are convex and cubic in the part load ratio. The poor convergence properties of the lambda iteration at low cooling load demand were obviated by a suitable Gradient method.

Nevertheless, it is well-known that chillers' power consumption could be a concave function, which makes OCL an NP-hard mixed-integer problem and poses a further challenge to the search of the optimal solution as global convergence of iterative methods cannot be guaranteed. This has motivated the development of several heuristic methods which do not guarantee optimality, but are effective to obtain a fair solution in a viable execution time.

Assuming quadratic models of power consumption, Geem [42] resorts to the Generalized Reduced Gradient (GRG) method. Salari et al. [43] show that the mixed integer problem can be solved using the general algebraic modeling system (GAMS). In parallel, many nature-inspired heuristic algorithms have been proposed, namely, genetic algorithm (GA) [44, 45], simulated annealing (SA) [46, 47], particle swarm optimization (PSO) [48, 49], evolution strategy (ES)[50], differential evolution (DE) [51], cuckoo search algorithm using differential operator (DCSA)

[52], differential search (DS) [53], improved firefly algorithm (IFA) [54], teaching-learning-based optimization (TLBO) [55] and improved invasive weed optimization (EIWO) [56].

Many of the these algorithms have been tested on the Hsinchu benchmark, a widely used case study consisting of a six-chiller system installed in a semiconductor factory located in the Hsinchu Scientic Garden (Taiwan). Hence, the performances achieved on the Hsinchu benchmark can be used to compare and rank alternative solution methods. A selection of the five best algorithms is reported in Table 5.1. In order of publication, they are IFA (2013), DCSA (2014), GAMS (2015), TLBO (2017) and EIWO (2018). It is seen that for the three highest cooling load demands (Q_{load} = 6858, 6477 and 6096 [kW]) all the algorithms achieve the same power consumption. In the remaining two cases (Q_{load} = 5717 and 5334 [kW]) the most promising algorithms would seem to be IFA and DCSA, which obtain the same results. Nevertheless, since an exact solver of the Hsinchu benchmark is not available, the final word on the optimality of solutions found by heuristic algorithms is still to be said.

In the practical management of a real chiller system, solving the OCL is not sufficient because there exist further dynamic constraints, namely minimum uptime and downtime requirements on chillers' operation. When these constraints are accounted for, the power consumption minimization problem goes under the name of Optimal Chiller Sequencing (OCS). Again, this is a problem that is hardly tractable without resorting to some heuristics. In particular, the knowledge of all future cooling loads is required, which raises the problem of forecasting it with reasonable accuracy, a task that can be successfully addressed only on a finite prediction horizon. In the literature, OCS solvers resulting from the combination of heuristic OCL algorithms with dynamic programming schemes have been proposed [14, 45, 57, 58].

The present chapter addresses both the OCL and OCS problems. Concerning the former one, two main issues are investigated. First of all, we derive an exact algorithm when the chillers' power consumption is a quadratic function of the Partial Load Ratio (PLR), as it happens for the Hsinchu benchmark. Our X-OCL algorithm allows to say the final word on the existing heuristic methods, highlighting also some erroneous results reported in the literature. The second issue has to do with the practical applicability of X-OCL to real-world plants. In particular, the execution time is compared with a state-of-art mixed integer solver and the adequacy of the quadratic power consumption model is discussed.

Concerning the OCS problem, we exploit the X-OCL algorithm to derive a lower bound on the minimum power consumption achievable by any OCS solver. Second, a greedy OCS algorithm leveraging on X-OCL is proposed and compared with [57]. Finally, the lower bound is used to quantitatively assess the degree of suboptimality ensuing from the lack of preview implicit in the greedy approach.

Literature benchmarks, i.e. the Hsinchu one and two OCS benchmarks are used for test and comparison. Moreover, an extensive dataset collected during two years in a semiconductor fab is used to build and demonstrate a comprehensive solution of both OCL and OCS, including the data-based estimation of the chillers' power consumption models. In particular, the potential energy saving with respect with the current HVAC energy management is assessed.

For sake of semplicity, in this chapter $Q_{e,i}$ and $P_{el,i}$ we will be denotes by as Q_i and P_i, respectively.

5.1 The Optimal Chiller Loading problem

Assuming n chillers operated in parallel, let Q_i, $i = 1, \ldots, n$ denote the cooling power delivered by the i-th chiller and $P_i = P_i(Q_i)$, $i = 1, \ldots, n$ the associated power consumption. For a prescribed overall cooling load demand Q_{load}, the goal of the OCL problem is finding the cooling powers Q_i, $i = 1, \ldots, n$ that the n chillers have to deliver in order to minimize the system total energy consumption:

$$\min_{Q_1,\ldots,Q_n} \sum_{i=1}^{n} P_i, \quad \text{s.t. } Q_{load} = \sum_{i=1}^{n} Q_i$$

For each chiller, PLR (Part Load Ratio) denotes the cooling load fraction, given by

$$PLR_i = \frac{Q_i}{Q_{100\%,i}}$$

where $Q_{100\%,i}$ is the maximum power supplied under full capacity operation. The vector of all PLR's is denoted by

$$\mathbf{PLR} = [PLR_1, \ldots, PLR_n]^T$$

When the i-th chiller is turned on, it should not operate under a minimum PLR, denoted by PLR_{min}. For the subsequent derivations it is convenient to introduce a binary variable δ_i that indicates the status of the $i-$th chiller and a real-valued variable x_i such that:

$$PLR_i = \delta_i \cdot x_i, \quad \delta_i \in \{0, 1\},\ x_i \in [PLR_{min,i}, 1]$$

The power consumption P_i of the i-th chiller is assumed to depend mainly on PLR_i and the condenser inlet water temperature T_i, that is $P_i = P_i(PLR_i, T_i)$. The consumption surface $P_i(PLR_i, T_i)$ is either obtained from laboratory experiments or field data collected during operation.

$Q_{load}[kW]$	i	IFA (2013)		DCSA (2014)		GAMS (2015)		TLBO (2017)		EIWO (2018)	
		PLR_i	$P_i[kW]$	PLR_i	$P_i[kW]$	PLR_i	$P_i[kW]$	PLR_i	$P_i[kW]$	PLR_i	$P_i[kW]$
6858 (90%)	1	0.812774	-	0.812726	-	0.8127	808.9736	0.8186	-	0.8127	-
	2	0.749527	-	0.749619	-	0.7496	740.7275	0.7523	-	0.7492	-
	3	1.000000	-	1.000000	-	1.0000	903.3450	1.0000	-	1.0000	-
	4	1.000000	-	1.000000	-	1.0000	781.4890	1.0000	-	1.0000	-
	5	1.000000	-	1.000000	-	1.0000	755.2010	1.0000	-	1.0000	-
	6	0.838603	-	0.838559	-	0.8386	748.8392	0.8297	-	0.8390	-
	$\sum$		**4738.576**		**4738.575**		**4738.5753**		**4738.54**		**4738.575**
6477 (85%)	1	0.727803	-	0.727731	-	0.7277	718.5040	0.727731	-	0.7275	-
	2	0.656174	-	0.656132	-	0.6561	641.1960	0.656132	-	0.6563	-
	3	1.000000	-	1.000000	-	1.0000	903.3450	1.000000	-	1.0000	-
	4	1.000000	-	1.000000	-	1.0000	781.4890	1.000000	-	1.0000	-
	5	1.000000	-	1.000000	-	1.0000	755.2010	1.000000	-	1.0000	-
	6	0.716408	-	0.716524	-	0.7165	621.9135	0.716524	-	0.7166	-
	$\sum$		**4421.649**		**4421.649**		**4421.6486**		**4421.65**		**4421.649**
6096 (80%)	1	0.642725	-	0.642735	-	0.6427	639.1411	0.6431	-	0.6427	-
	2	0.562642	-	0.562645	-	0.5626	553.8955	0.5621	-	0.5628	-
	3	1.000000	-	1.000000	-	1.0000	903.3450	1.000000	-	1.0000	-
	4	1.000000	-	1.000000	-	1.0000	781.4890	1.000000	-	1.0000	-
	5	1.000000	-	1.000000	-	1.0000	755.2010	1.000000	-	1.0000	-
	6	0.594504	-	0.594490	-	0.5945	510.6347	0.5946	-	0.5944	-
	$\sum$		**4143.706**		**4143.706**		**4143.7064**		**4143.64**		**4143.706**
5717 (75%)	1	0.842218	-	0.843697	-	0.0000	0.0000	0.55765	-	0.0000	-
	2	0.781365	-	0.783794	-	0.7150	702.4809	0.46918	-	0.7151	-
	3	0.000002	-	0.000001	-	1.0000	903.3450	0.99995	-	1.0000	-
	4	0.999995	-	1.000000	-	1.0000	781.4890	1.00000	-	1.0000	-
	5	1.000000	-	1.000000	-	1.0000	755.2010	1.00000	-	1.0000	-
	6	0.887053	-	0.883049	-	0.7934	700.0373	0.47250	-	0.7933	-
	$\sum$		**3840.063***		**3840.055***		**3842.5532**		**3904.70**		**3842.553**
5334 (70%)	1	0.759350	-	0.749969	-	0.0000	0.0000	0.64179	-	0.0000	-
	2	0.691121	-	0.682477	-	0.5835	572.3074	0.66219	-	0.5834	-
	3	0.000021	-	0.000012	-	1.0000	903.3450	0.33009	-	1.0000	-
	4	1.000000	-	1.000000	-	1.0000	781.4890	0.99059	-	1.0000	-
	5	1.000000	-	1.000000	-	1.0000	755.2010	0.99900	-	1.0000	-
	6	0.757897	-	0.776363	-	0.6217	534.0950	0.58047	-	0.6218	-
	$\sum$		**3507.286***		**3507.270***		**3546.4375**		**3642.51**		**3546.438**

TABLE 5.1: Comparison of the results of IFA, DCSA, GAMS, TLBO and EIWO on Hsinchu benchmark.

Note that the problem of finding the optimal part load ratios $\mathbf{PLR}^*$ can be formulated as a *mixed-integer nonlinear program* (MINLP):

$$\mathbf{PLR}^* := \arg\min_{\mathbf{PLR}} \quad \sum_{i=1}^{n} P_i(PLR_i) \tag{5.1a}$$

$$\text{subject to} \quad \sum_{i=1}^{n} PLR_i \cdot Q_{100\%,i} = Q_{load} \tag{5.1b}$$

$$PLR_i = \delta_i \cdot x_i \tag{5.1c}$$

$$PLR_{min} \leq x_i \leq 1 \tag{5.1d}$$

$$\delta_i \in \{0, 1\} \tag{5.1e}$$

In the above problem, two types of constraints are present: the *cooling demand constraint* (5.1b) and a set of *operational constraints* eqs. (5.1c) to (5.1e) regarding the admissible operating regions of the chillers.

In the following it is implicitly assumed that the cooling demand constraint is such that the admissible solution set is nonempty. Then, given that the constraints define a closed set of admissible solution, the cost function admits a minimum.

5.1.1 Quadratic power consumption model

With exception of particular structures, mixed-integer programming problems are classified as $\mathcal{NP}$-hard, which means that in the worst case, the solution time grows at least exponentially with the problem size. Although its combinatorial nature might suggest the use of heuristics, we will show that a significant subclass of industrial OCL problems, characterized by a quadratic power consumption model, may still be successfully attacked by a carefully designed exact method.

Assumption 5.1. *The power consumption P_i of the i-th chiller obeys the following model:*

$$P_i(PLR_i, T_i) = \begin{cases} \beta_{0,i} + \beta_{1,i} PLR_i + \beta_{2,i} PLR_i^2 + \beta_{3,i} T_i, & \text{if } PLR_{min} \leq PLR_i \leq 1 \\ 0, & \text{otherwise} \end{cases} \tag{5.2}$$

where $\beta_{p,i}$ are the model parameters. Moreover, it is assumed that $\beta_{2,i} \neq 0$

When the OCL problem is solved in a given time slot, the condenser inlet water temperature can be assumed to be known. Then, for a given condenser inlet water temperature T_i, in the interval $[PLR_{min}, 1]$ the consumption surface is a quadratic function of PLR_i alone:

$$P_i(PLR_i, T_i) = \begin{cases} a_i + c_i PLR_i + q_i PLR_i^2, & \text{if } PLR_{min} \leq PLR_i \leq 1 \\ 0, & \text{otherwise} \end{cases} \tag{5.3}$$

with $a_i = \beta_{0,i} + \beta_{3,i} T_i$, $c_i = \beta_{1,i}$, and $q_i = \beta_{2,i}$.

In view of the quadratic nature of the cost function, the system total energy consumption can be expressed in matrix form as follows:

$$P_{tot} = \sum_{i=1}^{n} P_i = \frac{1}{2}\mathbf{x}^T \mathbf{Q}\mathbf{x} + \mathbf{c}^T \mathbf{x} + \mathbf{a} \tag{5.4}$$

where

$$\begin{aligned} \mathbf{x} &= \{PLR_i\},\ \mathbf{x} \in \mathbb{R}^{n\times 1} \\ \mathbf{a} &= \{a_i\},\ \mathbf{a} \in \mathbb{R}^{n\times 1} \\ \mathbf{c} &= \{c_i\},\ \mathbf{c} \in \mathbb{R}^{n\times 1} \\ \mathbf{Q} &= diag\{2q_i\},\ \mathbf{Q} \in \mathbb{R}^{n\times n} \end{aligned} \tag{5.5}$$

Moreover, the equality constraint (5.1b) can be rewritten as

$$\mathbf{E}\mathbf{x} = d$$

where

$$\begin{aligned} \mathbf{E} &= \{Q_{100\%,i}\},\ \mathbf{E} \in \mathbb{R}^{1\times n} \\ d &= Q_{load},\ d \in \mathbb{R} \end{aligned} \tag{5.6}$$

5.1.2 Partition of the solution space

In view of (5.1), the admissible set $\mathcal{F}$ for PLR_i is

$$PLR_i \in \mathcal{F}, \quad \mathcal{F} = \bigcup_{s=0}^{3} \mathcal{F}_\sigma \tag{5.7}$$

$$\mathcal{F}_0 = \{0\}, \quad \mathcal{F}_1 = \{PLR_{min}\}, \tag{5.8}$$

$$\mathcal{F}_2 = \{1\}, \quad \mathcal{F}_3 = (PLR_{min}, 1) \tag{5.9}$$

where each subset $\mathcal{F}_\sigma$ is associated with one of the following operating conditions: switched off ($\sigma = 0$), minimum part load ($\sigma = 1$), maximum part load ($\sigma = 2$), intermediate part load ($\sigma = 3$). In the sequel, $\sigma_i \in \{0, 1, 2, 3\}$ will denote the state of the $i-$th chiller.

In order to satisfy the operational constraints of a multiple chiller system, the solution of the OCL problem must be searched within the cartesian product of the chillers' admissible sets, i.e.

$$\mathbf{PLR} \in \mathcal{S}, \quad \mathcal{S} = \mathcal{F}^n$$

Since the admissible set $\mathcal{F}$ of a single chiller can be partitioned in four subsets, the overall admissible set $\mathcal{S}$ can be partitioned in 4^n subsets $\mathcal{S}_j$, $j = 1, \ldots, 4^n$, each of which is in a one-to-one correspondence with the $n-$digit multichiller code

$$\mathbf{s}_j = \left[\begin{array}{ccc} s_{j1} & \ldots & s_{jn} \end{array} \right], \quad s_{ji} \in \{0, 1, 2, 3\}$$

formed by the state codes σ_i, $i = 1, \ldots, n$ of the n chillers.

To make an example, consider the case of $n = 3$ chillers. Then, the possible $4^n = 64$ subsets $\mathcal{S}_j$ are associated to the multi-chiller codes as follows:

$$\begin{aligned}
\mathcal{S}_1 &= \mathcal{F}_0 \times \mathcal{F}_0 \times \mathcal{F}_0 \leftrightarrow \mathbf{s}_1 = [0, 0, 0] \\
\mathcal{S}_2 &= \mathcal{F}_0 \times \mathcal{F}_0 \times \mathcal{F}_1 \leftrightarrow \mathbf{s}_2 = [0, 0, 1] \\
\mathcal{S}_3 &= \mathcal{F}_0 \times \mathcal{F}_0 \times \mathcal{F}_2 \leftrightarrow \mathbf{s}_3 = [0, 0, 2] \\
\mathcal{S}_4 &= \mathcal{F}_0 \times \mathcal{F}_0 \times \mathcal{F}_3 \leftrightarrow \mathbf{s}_4 = [0, 0, 3] \\
\mathcal{S}_5 &= \mathcal{F}_0 \times \mathcal{F}_1 \times \mathcal{F}_0 \leftrightarrow \mathbf{s}_5 = [0, 1, 0] \\
\mathcal{S}_6 &= \mathcal{F}_0 \times \mathcal{F}_1 \times \mathcal{F}_1 \leftrightarrow \mathbf{s}_6 = [0, 1, 1] \\
\mathcal{S}_7 &= \mathcal{F}_0 \times \mathcal{F}_1 \times \mathcal{F}_2 \leftrightarrow \mathbf{s}_7 = [0, 1, 2] \\
\mathcal{S}_8 &= \mathcal{F}_0 \times \mathcal{F}_1 \times \mathcal{F}_3 \leftrightarrow \mathbf{s}_8 = [0, 1, 3] \\
\mathcal{S}_9 &= \mathcal{F}_0 \times \mathcal{F}_2 \times \mathcal{F}_0 \leftrightarrow \mathbf{s}_9 = [0, 2, 0] \\
\ldots &= \ldots \\
\mathcal{S}_{64} &= \mathcal{F}_3 \times \mathcal{F}_3 \times \mathcal{F}_3 \leftrightarrow \mathbf{s}_{64} = [3, 3, 3]
\end{aligned}$$

The n elements of the set $\mathcal{S}_j$ are in a one-to-one correspondence with the chiller number $i, 1 \leq i \leq n$. Given the chiller number i, the notation $\mathcal{S}_j[i]$ will denote the operating condition, either a point or a range, of the i-th chiller. For instance, within $\mathcal{S}_8$, the three chillers $i = 1, 2, 3$ will operate at the following conditions:

$$\begin{aligned}
\mathcal{S}_8[1] &= 0 \\
\mathcal{S}_8[2] &= PLR_{min} \\
\mathcal{S}_8[3] &= (PLR_{min}, 1)
\end{aligned}$$

5.1.3 Divide and conquer strategy

For $j = 1, \ldots, 4^n$, let $\mathbf{QP}(j)$ indicate the OCL problem (5.1) restricted to the subset $\mathcal{S}_j$.

Problem 1. $\mathbf{QP}(j)$

$$\min_{\mathbf{PLR}\in\mathcal{S}_j} \quad \sum_{i=1}^{n} P_i(PLR_i) \tag{5.10a}$$

$$\text{subject to} \quad \sum_{i=1}^{n} PLR_i \cdot Q_{100\%,i} = Q_{load} \tag{5.10b}$$

The 4^n problems $\mathbf{QP}(j)$ can be partitioned in two subsets $\mathcal{C}$ and $\bar{\mathcal{C}}$:

- $\mathcal{C} = \{j \mid s_{ji} \neq 3, \forall i\}$: all the partial load ratios are fixed so that $\mathcal{S}_j$ has cardinality one;
- $\bar{\mathcal{C}}$, when there is at least one chiller operating at intermediate part load ($\exists i : s_{ji} = 3$).

Consider for example $\mathcal{S}_7$, whose multi-chiller code is $\mathbf{s_7} = [0, 1, 2]$: chiller #1 is switched off, chiller #2 is operating at minimum part load, and chiller #3 is operating at maximum part load. Within this subset, no optimization is actually needed because all the chillers' PLR's are fixed and

$$\mathcal{S}_7 = \begin{bmatrix} 0 & PLR_{min} & 1 \end{bmatrix}^T$$

has cardinality one. Therefore only a *feasibility check* is required: if constraint (5.10b) is satisfied, then $\mathbf{PLR}^*(7) = \mathcal{S}_7$. Otherwise, $\mathbf{QP}(7)$ does not admit a solution. It is easy to see that the number of elements of the subset $\mathcal{C}$ is 3^n.

As a second example, consider $\mathcal{S}_8$ whose multichiller code is $\mathbf{s_8} = [0, 1, 3]$: here the chiller #3 is operating at intermediate part load, i.e.

$$\mathcal{S}_8 = \left\{ \mathbf{PLR} \quad | \quad \mathbf{PLR} = \begin{bmatrix} 0 & PLR_{min} & x_3 \end{bmatrix}^T \right\}$$

with $x_3 \in \mathcal{F}_3$. In this case, $\mathbf{QP}(8)$ is a constrained quadratic programming problem in the unknown x_3.

As it will be shown later, for $j \in \bar{\mathcal{C}}$, the $\mathbf{QP}(j)$ problems enjoy a remarkable property: their optimal solution, if exists, is a critical point and no more than one critical point exists. In the following, $\mathbf{PLR}^*(j)$ will denote:

1. the optimal solution of $\mathbf{QP}(j)$, if $j \in \mathcal{C}$;
2. the unique feasible critical point of $\mathbf{QP}(j)$, if $j \in \bar{\mathcal{C}}$.

Let $\mathcal{A}$ denote the set of integers j s.t. $\mathbf{PLR}^*(j)$ exists. The associated value of the cost function will be denoted by

$$P^*_{tot}(j) := \sum_{i=1}^{n} P_i(PLR^*_i(j)) \tag{5.11}$$

Then, the key idea is to exploit the partition of the solution space by a two-step procedure:

1. Solve $\mathbf{QP}(j)$, obtaining $\mathbf{PLR}^*(j)$ and $P^*_{tot}(j)$ for $j = 1, \ldots, 4^n$;
2. Letting
$$j^* = \arg\min_{j \in \mathcal{A}} P^*_{tot}(j)$$
obtain the globally optimal part load vector as $\mathbf{PLR}^* = \mathbf{PLR}^*(j^*)$.

In the next subsection, it is shown how to reduce the inequality-constrained problems associated with $j \in \bar{\mathcal{C}}$ to Equality-constrained Quadratic Problems (EQP) for which a closed form solution is available.

5.1.4 Reduction to equality-constrained problems

For a given $j \in \bar{\mathcal{C}}$ we denote by $\mathcal{V}_j = \{i \in \{1, 2, \ldots, n\} \mid s_{ji} = 3\}$ the set of chillers operating at intermediate part loads, i.e. between PLR_{min} and 1. We also let $\kappa(j)$ denote the cardinality of $\mathcal{V}_j$. It is then possible to rewrite $\mathbf{QP}(j)$ as a reduced-order quadratic problem in the $\kappa(j) \leq n$ unknowns PLR_i, $i \in \mathcal{V}_j$.

$$\min_{PLR_i, i \in \mathcal{V}_j} \quad \sum_{i \in \mathcal{V}_j} P_i(PLR_i) \tag{5.12a}$$

$$\text{subject to} \quad \sum_{i \in \mathcal{V}_j} PLR_i \cdot Q_{100\%,i} = \tilde{Q}_{load} \tag{5.12b}$$

$$PLR_{min} < PLR_i < 1, \quad \forall i \in \mathcal{V}_j \tag{5.12c}$$

where

$$\tilde{Q}_{Load} = Q_{Load} - \sum_{i \in \overline{\mathcal{V}_j}} PLR_i \cdot Q_{100\%,i} \tag{5.13}$$

is the cooling load that must be supplied by the chillers operating at intermediate part loads.

Now, we associate to each $\mathbf{QP}(j)$ the corresponding Equality-constrained Quadratic Problem $\mathbf{EQP}(j)$, that is obtained by removing the inequality constraints (5.12c).

Problem 2. $\mathbf{EQP}(j)$

$$\min_{PLR_i, i \in \mathcal{V}_j} \quad \sum_{i \in \mathcal{V}_j} P_i(PLR_i) \tag{5.14a}$$

$$\text{subject to} \quad \sum_{i \in \mathcal{V}_j} PLR_i \cdot Q_{100\%,i} = \tilde{Q}_{load} \tag{5.14b}$$

It is convenient to rewrite $\mathbf{EQP}(j)$ in matrix form. For this purpose, we introduce a selection matrix $\mathbf{M}(\mathcal{V}_j) \in \mathbb{R}^{\kappa(j)\times n}$ that selects $\kappa(j)$ elements out of n.

Note that, being integers, the elements of $\mathcal{V}_j$ admit an obvious ordering. Then,

$$\mathbf{M}(\mathcal{V}_j) = [M]_{k,i} = \begin{cases} 1 & \text{if } i = k\text{-th element of } \mathcal{V}_j \\ 0 & \text{otherwise} \end{cases}$$

By short, $\mathbf{M}(\mathcal{V}_j)$ will be denoted by $\mathbf{M}_j$. The reduced matrices are thus given by:

$$\begin{aligned} \tilde{\mathbf{Q}} &= \mathbf{M}_j \mathbf{Q} \mathbf{M}_j^T,\ \tilde{\mathbf{Q}} \in \mathbb{R}^{\kappa\times\kappa} \\ \tilde{\mathbf{E}} &= \mathbf{E}\mathbf{M}_j^T,\ \tilde{\mathbf{E}} \in \mathbb{R}^{1\times\kappa} \\ \tilde{\mathbf{c}} &= \mathbf{M}_j \mathbf{c},\ \tilde{\mathbf{c}} \in \mathbb{R}^{\kappa\times 1} \\ \tilde{d} &= \tilde{Q}_{Load} \end{aligned} \tag{5.15}$$

Then, the $\mathbf{EQP}(j)$ can be restated as:

$$\tilde{\mathbf{x}}^*(j) := \arg\min_{\tilde{\mathbf{x}}} \ \frac{1}{2}\tilde{\mathbf{x}}^T \tilde{\mathbf{Q}} \tilde{\mathbf{x}} + \tilde{\mathbf{c}}^T \tilde{\mathbf{x}}, \quad \tilde{\mathbf{x}} \in \mathbb{R}^{\kappa\times 1} \tag{5.16a}$$

$$\text{s.t.} \quad \tilde{\mathbf{E}}\tilde{\mathbf{x}} = \tilde{d} \tag{5.16b}$$

By applying the first-order Karush-Kuhn-Tucker (KKT) necessary condition to the $\mathbf{EQP}(j)$ problem, the following linear system is obtained:

$$\begin{bmatrix} \tilde{\mathbf{Q}} & \tilde{\mathbf{E}}^T \\ \tilde{\mathbf{E}} & 0 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{x}}^* \\ \lambda^* \end{bmatrix} = \begin{bmatrix} -\tilde{\mathbf{c}} \\ \tilde{d} \end{bmatrix} \tag{5.17}$$

where $\tilde{\mathbf{x}}^*$ identifies a critical point, either maximum, minimum or saddle, and $\lambda^* \in \mathbb{R}$ is the associated Lagrange multiplier. In order to guarantee the existence of the solution, a technical assumption is introduced.

Assumption 5.2. $\tilde{\mathbf{E}}\tilde{\mathbf{Q}}^{-1}\tilde{\mathbf{E}}^T \neq 0$.

Note that Assumption 5.2 is immediately satisfied if $q_i > 0, \forall i$, that is when the quadratic power consumption curves (5.3) are all convex, although this is not necessary.

Theorem 5.3. *Under Assumption 5.1,* $\mathbf{EQP}(j)$ *admits a unique critical point* $\mathbf{x}^*$*, given by*

$$\lambda^* = -\frac{\tilde{d} + \tilde{\boldsymbol{E}}\tilde{\boldsymbol{Q}}^{-1}\tilde{\boldsymbol{c}}}{\tilde{\boldsymbol{E}}\tilde{\boldsymbol{Q}}^{-1}\tilde{\boldsymbol{E}}^T} \tag{5.18}$$

$$\tilde{\boldsymbol{x}}^* = -\tilde{\boldsymbol{Q}}^{-1}(\tilde{\boldsymbol{c}} + \tilde{\boldsymbol{E}}^T\lambda^*) \tag{5.19}$$

Proof. In view of Assumption 5.1, $\det(\mathbf{Q}) \neq 0$, because $q_i \neq 0, \forall i$. Moreover, Assumption 5.2 guarantees that $\tilde{\mathbf{E}}\tilde{\mathbf{Q}}^{-1}\tilde{\mathbf{E}}^T \neq 0$. Then, it is immediate to see that the KKT condition (5.17) admits (5.18)-(5.19) as unique solution.

Given the critical point $\tilde{\mathbf{x}}^*$, it is easy to obtain a critical point for $\mathbf{QP}(j)$, as well. For this purpose, it is convenient to introduce an auxiliary vector $\bar{\mathbf{x}} \in \mathbb{R}^{n\times 1}$, such that:

$$\bar{\mathbf{x}} = \begin{cases} \mathcal{S}_j[i], & \text{if } i \in \bar{\mathcal{V}}_j \\ 0, & \text{otherwise} \end{cases}$$

The candidate critical point for $\mathbf{QP}(j)$ is thus given by:

$$\widehat{\mathbf{PLR}}^*(j) = \bar{\mathbf{x}} + \mathbf{M}_j^T \tilde{\mathbf{x}}^* \tag{5.20}$$

The keystone of the solution procedure is the connection between the critical points of $\mathbf{QP}(j)$ and $\mathbf{EQP}(j)$, as stated in the following theorem.

Theorem 5.4. *The critical point* $\mathbf{PLR}^*(j)$ *exists iff* $\widehat{\mathbf{PLR}}^*(j) \in \mathcal{S}_j$. *In such a case,* $\mathbf{PLR}^*(j) = \widehat{\mathbf{PLR}}^*(j)$.

Proof. Sufficiency. Assume that the critical point $\widehat{\mathbf{PLR}}^*$ for EQP(j) belongs to $\mathcal{S}_j$. Observe that EQP(j) has less contraints that QP(j). Therefore, if a critical point for EQP(j) is feasible for QP(j) it is ipso facto a critical point for QP(j), as well.

Necessity. Assume that QP(j) admits a critical point, say $\mathbf{PLR}^*(j)$. Given that EQP(j) and QP(j) differ only for strict inequality constraints, any critical point for QP(j) is critical also for EQP(j). Since EQP(j) admits at most one critical point, necessity is proven.

5.1.5 Summary of the X-OCL Algorithm

We are now in a position to summarize the steps of the proposed algorithm, hereafter named X-OCL, for the exact solution of the OCL problem.

The partition of the solution set in the 4^n subsets $\mathcal{S}_j$, $j = 1, \ldots 4^n$ allows to divide the MINLP problem in 4^n sub-problems $\mathbf{QP}(j)$ (5.10). For each subset $\mathcal{S}_j$ two situations can occur:

1. all the chillers part load ratios are fixed;
2. there is at least one chiller working at intermediate part load.

In the former case, only a feasibility check is required to decide whether the PLR configuration is to be kept as a candidate solution. In the latter case, the solution of the $\mathbf{QP}(j)$ problem is reduced to the solution of an $\mathbf{EQP}(j)$ problem that admits a unique critical point, easily computable in closed form. As stated in Theorem 5.4, if the optimal solution of $\mathbf{EQP}(j)$ exists and belongs to $\mathcal{S}_j$, then it coincides with the solution of $\mathbf{QP}(j)$. Otherwise, $\mathbf{QP}(j)$ does not admit a solution.

Once all the $\mathbf{QP}(j)$ have been solved, the corresponding set of solutions $\{\mathbf{PLR}^*(j),\ j = 1,\dots,4^n\}$ include the optimal solution of the overall OCL problem, which can be found just by comparing the associated system power consumptions $P^*_{tot}(j)$. A pseudo-code summary is reported in Algorithm 1.

5.2 Test on Hsinchu benchmark model

5.2.1 Hsinchu cooling plant model

The Hsinchu chiller system, originally described in [44], has become a widely used benchmark for the testing and comparison of OCL algorithms [43, 52, 54–56]. The case study involves six chillers installed in a semiconductor factory located in Hsinchu Scientific Garden (Taiwan) with a 7,620 kW total cooling capacity. Quadratic models of the chillers' energy consumption were obtained and validated from data collected every 5 min over a 5 month period [44]. The benchmark problem assumes that the condenser inlet water temperature is 24.5°C. For convenience, the coefficients of the six chillers' energy consumption models (5.3) are reported in Table 5.2 and the corresponding P-PLR curves are displayed in Fig. 5.1. It is asked to solve the OCL problem for five different cooling loads, ranging from 70% to 90% of the system total cooling capacity. It is also required that the partial load ratio of each chiller never goes below 0.3. According to our notation, the following parameter settings are used:

- a_i, c_i, and q_i from Table 5.2
- $Q_{Load} = 90\%, 85\%, 80\%, 75\%$, and 70% of the chillers' maximum capacity ($\sum_{i=1}^{n} Q_{nom}$).
- $PLR_{min,i} = 0.3 \quad \forall i$

Algorithm 1 X-OCL

Input: $\mathbf{a}, \mathbf{c}, \mathbf{q} \in \mathbb{R}^{n\times 1}$, $\mathbf{s} \in \mathbb{R}^{4^n \times n}$, $Q_{load} \in \mathbb{R}$
Output: $\mathbf{PLR}^*, P^*_{tot}$

for $j = 1, \ldots, 4^n$ **do** $\mathcal{V}_j = \{i \in \{1, 2, \ldots, n\} : s_{ji} = 3\}$, $\kappa = \mid \mathcal{V}_j \mid$
  if $\mathcal{V}_j = \emptyset$ **then** ▷ **Feasibility check**
    all the chillers part load ratios are fixed
    $\mathbf{x} \leftarrow \mathcal{S}_j$
    if $\sum_{i=1}^{n} x_i \cdot Q_{100\%,i} = Q_{load}$ **then**
      $\mathbf{PLR}^*(j) = \mathbf{x}$
      $P^*_{tot}(j) = \sum_{i=1}^{n} P_i(PLR^*_i(j))$
    else
      $\mathbf{PLR}^*(j) = \emptyset$
      $P^*_{tot}(j) = \emptyset$
    end if
  else ▷ **solve the EQP(j) associated to QP(j)**
    $\mathcal{V}_j \Rightarrow \tilde{\mathbf{Q}}, \tilde{\mathbf{E}}, \tilde{\mathbf{c}}, \tilde{d}$
    $\lambda^* = -\frac{\tilde{d}+\tilde{\mathbf{E}}\tilde{\mathbf{Q}}^{-1}\tilde{\mathbf{c}}}{\tilde{\mathbf{E}}\tilde{\mathbf{Q}}^{-1}\tilde{\mathbf{E}}^T}$, $\lambda^* \in \mathbb{R}$
    $\tilde{\mathbf{x}}^* = -\tilde{\mathbf{Q}}^{-1}(\tilde{\mathbf{c}} + \tilde{\mathbf{E}}^T\lambda^*)$
    $\bar{\mathbf{x}} = \begin{cases} \mathcal{S}_j[i], & \text{if } i \in \bar{\mathcal{V}}_j \\ 0, & \text{otherwise} \end{cases}$
    $\widehat{\mathbf{PLR}}^*(j) = \bar{\mathbf{x}} + \mathbf{M}_j^T\tilde{\mathbf{x}}^*$
    if $\widehat{\mathbf{PLR}}^*(j) \in \mathcal{S}_j$ **then**
      $\mathbf{PLR}^*(j) = \widehat{\mathbf{PLR}}(j)$
      $P^*_{tot}(j) = \sum_{i=1}^{n} P_i(PLR^*_i(j))$,
    else
      $\mathbf{PLR}^*(j) = \emptyset$
      $P^*_{tot}(j) = \emptyset$
    end if
  end if
end for

$j^* = \min_j P_{tot}(j)$
$\mathbf{PLR}^* = \mathbf{PLR}^*(j^*)$
$P^*_{tot} = P_{tot}(j^*)$

5.2.2 OCL benchmark: results

The OCL solutions provided by the five literature methods and X-OCL are reported in Table 5.3. For the three highest cooling load demands ($Q_{Load} = 6858, 6477$ and 6096 [kW]) the X-OCL optimal solution coincided with the common solution provided by the five algorithms, thus confirming that they had reached the optimum. In the remaining two cases ($Q_{Load} = 5717$ and 5334 [kW]), the solution computed by X-OCL coincided with those of GAMS adn EIWO, which, however, are apparently outperformed by IFA and DCSA. This outcome seems to contradict the exact nature of X-OCL. As a matter of fact, a closer look at the solution provided by IFA

TABLE 5.2: Hsinchu benchmark: P-PLR curves coefficients

Chiller	a_i	c_i	q_i	Q_{nom} [RT]
1	399.345	-122.12	770.46	1280
2	287.116	80.04	700.48	1280
3	-120.505	1525.99	-502.14	1280
4	-19.121	898.76	-98.15	1280
5	-95.029	1202.39	-352.16	1250
6	191.750	224.86	524.04	1250

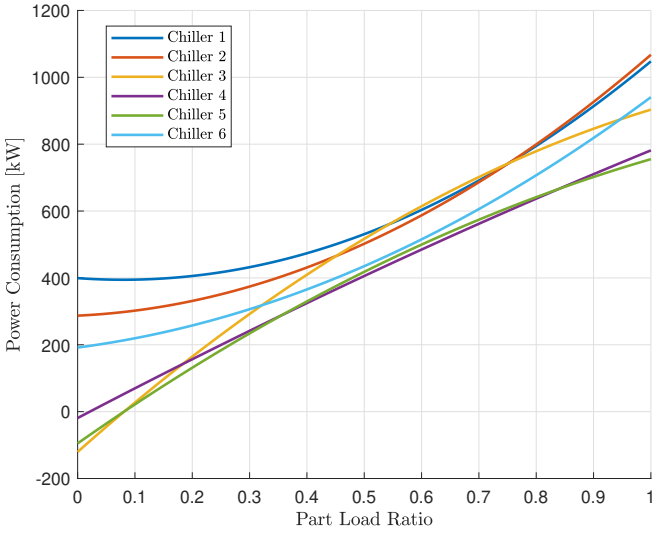

FIGURE 5.1: Hsinchu benchmark: P-PLR curves at $T = 24.5°C$

and DCSA reveals that four wrong values of power consumption (identified by asterisks in Table 5.1) were published in [52] and [54].

- $P_{tot,IFA}$ at load 75% and 70%;
- $P_{tot,DCSA}$ at load 75% and 70%.

For example, Table 5.1 reports $P_{tot,IFA} = P_{tot,DCSA} \approx 3507.3$ [kW] at 70% cooling load, which is inconsistent with the PLR_i reported in the same table. Such inconsistency is easily verified by plugging the PLR_i's into the chillers' power consumption equations (5.3) in order to obtain the individual chiller power consumptions reported in the $P_i[kW]$ column (incidentally, the papers [52, 54] describing algorithms IFA and DCSA were the only ones not reporting these individual consumptions). The source of the error is the mechanical use of the chiller's quadratic power consumption model out of its operational range, that is in correspondence of null partial load ratio. Obviously, the associated power consumption is null as well, but if the quadratic model has a negative constant term, as is the case for chillers 3-5, the formula will return a negative power consumption as if a turned-off chiller could generate free power.

$Q_{Load}[kW]$	i	IFA		DCSA		GAMS		TLBO		EIWO		X-OCL	
		$PLR_i[-]$	$P_i[kW]$	$PLR_i[-]$	$P_i[kW]$	$PLR_i[-]$	$P_i[kW]$	$PLR_i[-]$	$P_i[kW]$	$PLR_i[-]$	$P_i[kW]$	$PLR_i[-]$	$P_i[kW]$
6858 (90%)	1	0.812774	809.0541	0.812726	808.9999	0.8127	808.9736	0.8186	815.6654	0.8127	808.9705	0.8127	809.0002
	2	0.749527	740.6313	0.749619	740.7353	0.7496	740.7275	0.7523	743.7705	0.7492	740.2618	0.7496	740.7352
	3	1.000000	903.3450	1.000000	903.3450	1.0000	903.3450	1.0000	903.3450	1.0000	903.3450	1.0000	903.3450
	4	1.000000	781.4890	1.000000	781.4890	1.0000	781.4890	1.0000	781.4890	1.0000	781.4890	1.0000	781.4890
	5	1.000000	755.2010	1.000000	755.2010	1.0000	755.2010	1.0000	755.2010	1.0000	755.2010	1.0000	755.2010
	6	0.838603	748.8520	0.838559	748.8035	0.8386	748.8392	0.8297	739.0666	0.8390	749.2903	0.8386	748.8030
	$\sum$		**4738.576**		**4738.575**		**4738.5753**		**4738.54**		**4738.575**		**4738.5733**
6477 (85%)	1	0.727803	718.5742	0.727731	718.5023	0.7277	718.5040	0.727731	718.5023	0.7275	718.2715	0.7277	718.5010
	2	0.656174	641.2379	0.656132	641.1959	0.6561	641.1960	0.656132	641.1959	0.6563	641.3638	0.6561	641.1960
	3	1.000000	903.3450	1.000000	903.3450	1.0000	903.3450	1.000000	903.3450	1.0000	903.3450	1.0000	903.3450
	4	1.000000	781.4890	1.000000	781.4890	1.0000	781.4890	1.000000	781.4890	1.0000	781.4890	1.0000	781.4890
	5	1.000000	755.2010	1.000000	755.2010	1.0000	755.2010	1.000000	755.2010	1.0000	755.2010	1.0000	755.2010
	6	0.716408	621.8000	0.716524	621.9132	0.7165	621.9135	0.716524	621.9132	0.7166	621.9874	0.7165	621.9136
	$\sum$		**4421.649**		**4421.649**		**4421.6486**		**4421.65**		**4421.649**		**4421.6466**
6096 (80%)	1	0.642725	639.1269	0.642735	639.1356	0.6427	639.1411	0.6431	639.4526	0.6427	639.1052	0.6427	639.1359
	2	0.562642	553.8980	0.562645	553.9006	0.5626	553.8955	0.5621	553.4276	0.5628	554.0352	0.5626	553.9010
	3	1.000000	903.3450	1.000000	903.3450	1.0000	903.3450	1.000000	903.3450	1.0000	903.3450	1.0000	903.3450
	4	1.000000	781.4890	1.000000	781.4890	1.0000	781.4890	1.000000	781.4890	1.0000	781.4890	1.0000	781.4890
	5	1.000000	755.2010	1.000000	755.2010	1.0000	755.2010	1.000000	755.2010	1.0000	755.2010	1.0000	755.2010
	6	0.594504	510.6442	0.594490	510.6324	0.5945	510.6347	0.5946	510.7257	0.5944	510.5561	0.5945	510.6325
	$\sum$		**4143.706**		**4143.706**		**4143.7064**		**4143.64**		**4143.706**		**4143.3704**
5717 (75%)	1	0.842218	843.0026	0.843697	844.7431	0.0000	0.0000	0.55765	570.8354	0.0000	0.0000	0.0000	0.0000
	2	0.781365	777.3214	0.783794	780.1780	0.7150	702.4809	0.46918	478.8657	0.7151	702.5557	0.7150	702.4808
	3	0.000002	0.0000	0.000001	0.0000	1.0000	903.3450	0.99995	903.3189	1.0000	903.3450	1.0000	903.3450
	4	0.999995	781.4855	1.000000	781.4890	1.0000	781.4890	1.00000	781.4890	1.0000	781.4890	1.0000	781.4890
	5	1.000000	755.2010	1.000000	755.2010	1.0000	755.2010	1.00000	755.2010	1.0000	755.2010	1.0000	755.2010
	6	0.887053	803.5604	0.883049	798.9460	0.7934	700.0373	0.47250	414.9916	0.7933	699.9229	0.7934	700.0374
	$\sum$		**3960.579**		**3960.558**		**3842.5532**		**3904.70**		**3842.553**		**3842.5532**
5334 (70%)	1	0.759350	750.8680	0.749969	741.1047	0.0000	0.0000	0.64179	638.3158	0.0000	0.0000	0.0000	0.0000
	2	0.691121	677.0164	0.682477	668.0074	0.5835	572.3074	0.66219	647.2751	0.5834	572.2236	0.5835	572.3068
	3	0.000021	0.0000	0.000012	0.0000	1.0000	903.3450	0.33009	328.4962	1.0000	903.3450	1.0000	903.3450
	4	1.000000	781.4890	1.000000	781.4890	1.0000	781.4890	0.99059	774.8702	1.0000	781.4890	1.0000	781.4890
	5	1.000000	781.2010	1.000000	755.2010	1.0000	755.2010	0.99900	754.7026	1.0000	755.2010	1.0000	755.2010
	6	0.757897	663.1834	0.776363	682.1826	0.6217	534.0950	0.58047	489.8474	0.6218	534.1803	0.6217	534.0957
	$\sum$		**3627.758**		**3627.98**		**3546.4375**		**3642.51**		**3546.438**		**3546.6437**

TABLE 5.3: Comparison of the results of IFA, DCSA, GAMS, TLBO, EIWO and X-OCL on Hsinchu benchmark.

TABLE 5.4: Published chillers' power consumptions and the xorresponding corrected values.

$Q_{load}[kW]$	i	IFA			DCSA		
		$PLR_i[-]$	$P_{p,i}[kW]$	$P_{c,i}[kW]$	$PLR_i[-]$	$P_{p,i}[kW]$	$P_{c,i}[kW]$
5717 (75%)	1	0.842218	843.0026	843.0026	0.843697	844.7431	844.7431
	2	0.781365	777.3214	777.3214	0.783794	780.1789	780.1789
	3	0.000002	-120.5019	**0.0000**	0.000001	-120.5035	**0.0000**
	4	0.999995	781.4855	781.4855	1.000000	781.4890	781.4890
	5	1.000000	755.2010	755.2010	1.000000	755.2010	755.2010
	6	0.887053	803.5604	803.5604	0.883049	798.9460	798.9460
	$\sum$		3840.0690	**3960.5709**		3840.0545	**3960.5580**
5334 (70%)	1	0.759350	750.8680	750.8680	0.749969	741.1047	741.1047
	2	0.691121	677.0164	677.0164	0.682477	668.0074	668.0074
	3	0.000021	-120.4730	**0.0000**	0.000012	-120.4867	**0.0000**
	4	1.000000	781.4890	781.4890	1.000000	781.4890	781.4890
	5	1.000000	755.2010	755.2010	1.000000	755.2010	755.2010
	6	0.757897	663.1834	663.1834	0.776363	682.1826	682.1826
	$\sum$		3507.2848	**3627.7578**		3507.4980	**3627.9847**

In Table 5.4 the published chillers' power consumption $P_{p,i}$ and the corresponding corrected values $P_{c,i}$ are reported for both IFA and DCSA. Once the errors have been corrected, IFA and DCSA are no more optimal at 75% and 70% loads. At load 75% $P_{tot,IFA}$ rises from 3840.0690 to 3960.5709 [kW] (+3.14%), while at 70% it rises from 3507.2848 to 3627.7578 [kW] (+3.43%). Analogously, for DCSA $P_{tot,DCSA}$ goes from 3840.0545 to 3960.5580 [kW] (+3.14%) at 75% and from 3507.2848 to 3627.7578 [kW] (+3.43%) at 70%.

In conclusion, on the Hsinchu benchmark, GAMS and EIWO prove to be the best heuristic algorithms as their solutions coincide with the optimal ones computed by X-OCL, for all considered loads.

5.3 The Optimal Chiller Sequencing problem

The cooling load demand of a building can be subject to significant variations during the day. Consequently, solving the OCL problem in each time step t (for example, 20 min), just ignoring minimum up/down time constraints on the chillers, could lead to frequent switchings (chiller startups and shutdowns). In order to preserve chillers from excessive mechanical stress and increase their operating life, each machine should not be switched off before a minimal up-time is reached. Analogously, it should not be switched on too quickly. To comply with these requirements, minimum up/down-time constraints must be enforced in the formulation of the so-called Optimal Chiller Sequencing (OCS) problem. In its full formulation, OCS is a dynamical problem, because, in order to minimize the cumulative power consumption, the current decision should take into account also future constraints. As a consequence, the solution approaches proposed in the literature range from dynamic programming to heuristic methods designed to alleviate the complexity of the problem.

5.4 A lower bound to the OCS problem

Any solution to the OCS problem must face some level of approximation. Even when dynamic programming is used, there is the necessity of forecasting future loads, which introduces a suboptimality margin with respect to the ideal solution based on perfect knowledge of the future load profile. It is therefore of interest the availability of an easy-to-compute limit of performance against which the results of heuristic methods can be benchmarked. In order to derive such a limit, one can consider the *relaxed OCS problem* (R-OCS), that is an OCS problem without up- and down-time constraints. The relaxed OCS problem boils down to a sequence of independent OCL problems to be solved at each step in correspondence of the associated load. The availability of an exact OCL solver, such as X-OCS, makes it possible to compute the exact solution of the relaxed OCS as well. Notably, in view of the independence of the OCL problems, what matters is not the load sequence but the load distribution, so that the R-OCS bound could be easily derived based on statistical distributions reflecting different production and weather scenarios.

Such a bound can be used to quantitatively assess the existing margin of improvement for a given heuristic OCS solver. In fact, if the achieved power consumption is close enough to the bound, there is no scope for the search of further improvements. Along this direction, in the following section, the X-OCL solver is used to derive a greedy OCS algorithm, whose performance is then assessed against the R-OCS bound.

5.5 X-OCS, a greedy approach to OCS

X-OCS is a greedy algorithm that reduces OCS to a sequence of OCL problems, solvable through X-OCL. The approach is greedy because at each time step, future constraints are ignored, and the optimal OCL solution compatible with the current minimum up/down time constraint is searched for. In mathematical form, the greedy OCS problem can be written as a mixed-integer quadratic problem with linear constraints:

Problem 3. greedy OCS

$$\mathbf{PLR}^*(t) := \arg\min_{\mathbf{PLR}} \quad \sum_{i=1}^{n} P_i(PLR_i(t), T_i(t)) \tag{5.21a}$$

$$\text{subject to} \quad \sum_{i=1}^{n} PLR_i(t) \cdot Q_{100\%,i} = Q_{load}(t) \tag{5.21b}$$

$$PLR_i(t) = \delta_i(t) \cdot x_i(t) \tag{5.21c}$$

$$PLR_{min} \leq x_i(t) \leq 1 \tag{5.21d}$$

$$(T_i^{ON}(t-1) - MUT_i)(\delta_i(t) - \delta_i(t-1)) \leq 0 \tag{5.21e}$$

$$(T_i^{OFF}(t-1) - MDT_i)(\delta_i(t) - \delta_i(t-1)) \geq 0 \tag{5.21f}$$

$$\delta_i \in \{0, 1\} \tag{5.21g}$$

where MUT_i and MDT_i are the $i-$th chiller's minimum up-time and minimum down-time limits, expressed in number of time steps. The time counters $T_i^{ON}(t)$ and $T_i^{OFF}(t)$, for which a chiller has been continuously on/off until the time t, are expressed as:

$$T_i^{ON}(t) = (1 + T_i^{ON}(t-1))\delta_i(t) \tag{5.22}$$

$$T_i^{OFF}(t) = (1 + T_i^{OFF}(t-1))(1 - \delta_i(t)) \tag{5.23}$$

It is easy to observe that the greedy OCS problem is an OCL problem with the two additional constraints (5.21e - 5.21f) which force some chillers to be online/offline depending on their previous states $\delta(\tau), \tau = t-1, t-2, \ldots$. At each time step t, two situations can occur:

- all the chillers' states $\delta_i(t)$ are free, i.e. all the chillers have been online/offline for more time steps than those prescribed by MUT/MDT
- there is at least one chiller, say the i-th one, whose state $\delta_i(t)$ is constrained to be online or offline ($\delta_i(t) = 1$ or $\delta_i(t) = 0$) by the the MUT or MDT.

Concerning the first case, the minimum up/down-time constraints are not active, therefore the step of the greedy OCS boils down to an OCL problem and its optimal solution can be found

by the X-OCL algorithm. In the second case, instead, at each time t, the optimal solution is found by applying the X-OCL algorithm over a suitable subset of the OCL solution space $\mathcal{S}$. At each time step t, the solution space of the greedy OCS is therefore given by:

$$\mathcal{P}(t) = \{j \in \{1, \dots, 4^n\} \mid s_{j,i} \neq 0,\ \forall i \in \mathcal{B}^{ON}(t), s_{j,i} = 0,\ \forall i \in \mathcal{B}^{OFF}(t)\}$$

where

$$\mathcal{B}^{ON}(t) = \{i \in \{1, \dots, n\} \ \mid (T_i^{ON}(t-1) - MUT_i) < 0\}$$
$$\mathcal{B}^{OFF}(t) = \{i \in \{1, \dots, n\} \ \mid (T_i^{OFF}(t-1) - MDT_i) < 0\}$$

are the set of chillers which must be turned on and off, respectively.

The idea is to exploit the partitions of the solution space $\mathcal{P}_j$ by the typical two-step procedure of X-OCL:

1. Solve $\mathbf{QP}(j)$, obtaining $\mathbf{PLR}^*(j)$ and $P^*_{tot}(j)$ for $j \in \mathcal{A}(t) \cap \mathcal{P}(t)$;
2. Letting
$$j^* = \arg \min_{j \in \mathcal{A}(t) \cap \mathcal{P}(t)} P^*_{tot}(j)$$
obtain the globally optimal part load vector as $\mathbf{PLR}^* = \mathbf{PLR}^*(j^*)$.

Herein, $\mathcal{A}(t)$ which denotes the set of integers j s.t. $\mathbf{PLR}^*(j)$ exists, is a function of t because the load changes with time.

5.5.1 OCS Simulated example

In this section, the performance of the X-OCS method is assessed on two case studies taken from the literature [57]. Case study 1 involves a hotel in Taipei with two 450 refrigeration tons (RT) chillers and two 1000 RT chillers, while Case study 2, again from the Hsinchu Science Industrial District, features nine 1250 RT chillers. In both cases, chillers are described by their COP-PLR curves, expressed by a second-order polynomial model:

$$COP_i = \alpha_i + \beta_i PLR_i + \gamma_i PLR_i^2 \tag{5.24}$$

where α_i, β_i and γ_i are the chiller's coefficients, reported in Table A.1.

The first step was the identification of quadratic power consumption models (5.3) using data sampled from the COP-PLR curves. Further details regarding the identification procedure are reported in Appendix A. The coefficients of the quadratic P-PLR curves in Fig. 5.2 are reported in Table A.2.

Systems	Chiller	a_i	c_i	q_i	Q_{nom}
Case 1	1	243.58	-398.01	504.00	450
	2	130.81	-103.53	309.65	450
	3	417.51	-444.57	771.99	1000
	4	383.79	-347.84	611.64	1000
Case 2	1	95.54	321.92	103.60	1250
	2	170.68	41.43	235.46	1250
	3	371.09	-307.99	693.76	1250
	4	477.85	-217.09	733.35	1250
	5	433.17	-186-20	810.89	1250
	6	104.21	358.47	205.50	1250
	7	272.33	116.22	457.96	1250
	8	218.68	-20-94	333.72	1250
	9	191.69	276.56	429.51	1250

TABLE 5.5: Estimated parameters of chillers' P-PLR curves

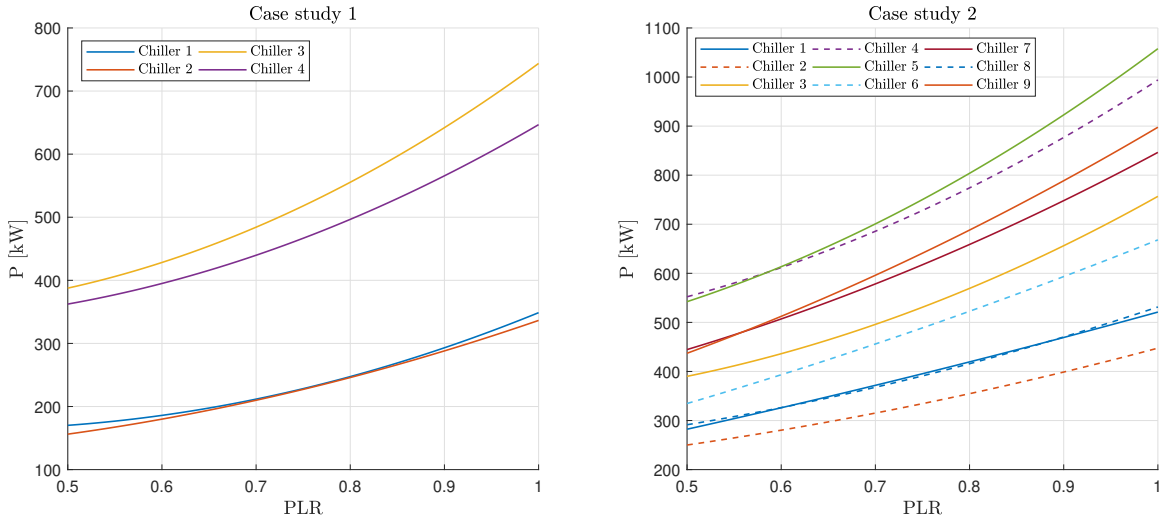

FIGURE 5.2: Estimated chillers' P-PLR curves on case study 1 (left) and case study 2 (right).

The aim of both benchmarks is to compute the sequence of chillers' partializations over one day, assuming 20-min stages, so as to minimize the cumulative power consumption, while satisfying the cooling demand constraint at each stage. The load demand profiles are represented in Figure 5.2.

The parameters settings were as follows:

- $PLR_{min} = 0.5$
- $MUT_i = 3 \quad \forall i$

- $MDT_i = 1 \quad \forall i$

Recall that $MUT_i = 3$ means that the $i-$th chiller must be on at least 3 steps before being turned off. Likewise, $MDT_i = 1$ indicates that the $i-$th chiller, once turned off, must remain off at least 1 step.

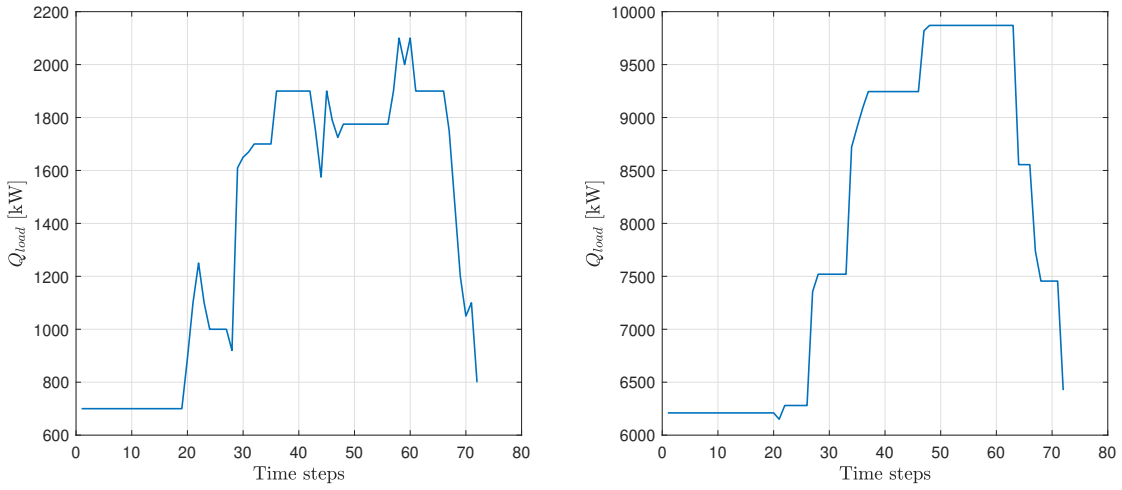

FIGURE 5.3: Cooling load demand profile for case study 1 (left) and case study 2 (right).

The results obtained via X-OCS were compared with those obtained by Dynamic Programming, as reported in [57]. Dynamic Programming was used under the ideal condition that all future loads are known in advance.

For sake of comparability with DP, the chillers power consumption associated with X-OCS was evaluated by plugging the PLRs computed by X-OCS into the original benchmark's COP model [57] and not in the approximate quadratic power consumption curves used by X-OCS.

5.5.2 OCS benchmark: results

The results are shown in Tables 5.6 and 5.7. In the case study 1 the power consumption obtained by the X-OCS method was, at each stage, lower or equal than that obtained by the DP method. From stage 1 to stage 29 the DP and X-OCS methods selected the same chillers. The MUT and MDT constraints were not active, so the solutions coincided with the OCL ones. At stage 30, the OCL solution had chillers 4 and 3 switched on. However, the MUT constraint forced the DP and X-OCS to leave chillers 1 and 2 switched on. The same goes for stage 31. At stage 32, 44 and 59 the X-OCS method performed better than the DP one. Although the chillers were not constrained by minimum up/down time limits, apparently the DP method, as implemented in [57], could not find the global minimum.

Stage	*Load*	X-OCL		DP		X-OCS		
		kW	*chiller*	*kW*	*chiller*	*kW*	*chiller*	$[PLR_1, PLR_2, PLR_3, PLR_4]$
1	700	441.13	4	441.15	4	441.13	4	[0.00, 0.00, 0.00, 0.70]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
20	890	556.53	4	556.55	4	556.53	4	[0.00, 0.00, 0.00, 0.89]
21	1100	695.32	4 2	695.56	4 2	695.32	4 2	[0.00, 0.63, 0.00, 0.82]
22	1250	801.66	4 2	801.78	4 2	801.66	4 2	[0.00, 0.73, 0.00, 0.92]
23	1100	695.32	4 2	695.56	4 2	695.32	4 2	[0.00, 0.63, 0,00, 0.82]
24	1000	634.42	4 2	636.73	4 2	634.42	4 2	[0.00, 0.50, 0,00, 0.77]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
28	920	579.96	4	579.98	4	579.96	4	[0.00, 0.00, 0.00, 0.92]
29	1610	1048.59	4 2 1	1048.43	4 2 1	1048.59	4 2 1	[0.75, 0.75, 0.00, 0.94]
30	1650	1081.63	4 3	1082.07	4 2 1	1082.07	4 2 1	[0.76, 0.77, 0.00, 0.96]
31	1670	1096.43	4 3	1099.69	4 2 1	1099.43	4 2 1	[0.77, 0.78, 0.00, 0.97]
32	1700	1118.96	4 3	1122.12	4 2 3	1118.96	4 3	[0.00, 0.00, 0.78, 0.91]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
36	1900	1246.74	4 3 2	1247.61	4 2 3	1246.74	4 3 2	[0.00, 0.67, 0.74, 0.86]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
43	1750	1150.68	4 3 2	1151.60	4 2 3	1150.68	4 2 3	[0.00, 0.61, 0.69, 0.79]
44	1575	1020.82	4 2 1	1030.56	4 3	1020.82	4 2 1	[0.74, 0.73, 0.00, 0.92]
45	1900	1246.81	4 3 2	1250.32	4 3 1	1249.42	4 3 1	[0.70, 0.00, 0.73, 0.85]
46	1790	1175.08	4 3 2	1180.75	4 3 1	1179.56	4 3 1	[0.67, 0.00, 0.69, 0.79]
47	1725	1135.32	4 3 2	1142.76	4 3 1	1135.32	4 3 2	[0.00, 0.50, 0.70, 0.80]
48	1775	1165.89	4 3 2	1166.78	4 3 2	1165.89	4 3 2	[0.00, 0.62, 0.69, 0.80]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
57	1900	1246.75	4 3 2	1247.61	4 3 2	1246.75	4 3 2	[0.00, 0.67, 0.74, 0.86]
58	2100	1385.03	4 3 2 1	1386.08	4 3 2 1	1385.03	4 3 2 1	[0.68, 0.63, 0.70, 0.81]
59	2000	1319.66	4 3 1	1326.20	4 3 2 1	1319.66	4 3 1	[0.73, 0.00, 0.77, 0.90]
60	2100	1385.03	4 3 2 1	1386.08	4 3 2 1	1385.03	4 3 2 1	[0.68, 0.63, 0.70, 0.81]
61	1900	1246.75	4 3 2	1247.61	4 3 2	1246.75	4 3 2	[0.00, 0.67, 0.74, 0.86]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
67	1750	1150.68	4 3 2	1151.60	4 3 2	1150.68	4 3 2	[0.00, 0.61, 0.69, 0.79]
68	1475	948.04	4 2 1	948.08	4 2 1	948.04	4 2 1	[0.71, 0.67, 0.00, 0.85]
69	1200	763.59	4 2	765.20	4 1	765.45	4 1	[0.72, 0.00, 0.00, 0.88]
70	1050	664.47	4 2	671.50	4 1	670.46	4 1	[0.65, 0.00, 0.00, 0.76]
71	1100	695.32	4 2	700.10	4 1	695.32	4 2	[0.00, 0.63, 0.00, 0.82]
72	800	495.80	4	495.82	4	539.23	4 2	[0.00, 0.50, 0.00, 0.57]
P_{day}		**64432.56** kW		**64883.84**$^*kW$		**64495.91** kW		

TABLE 5.6: Case study 1

In the second case study, the cooling load demand varies slowly over time, so that the MUT and MDT constraints were never active. Therefore, the X-OCS, notwithstanding its greedy nature, attains the best achievable performance bound R-OCS. For the majority of the cooling loads, DP and X-OCS gave the same results, the only exceptions being at stages 22, 48-62, 68-72 where X-OCS performed slightly better. For both the case studies, the cumulative daily power consumptions obtained by DP (marked by asterisks) had been reported incorrectly in [57]:

- Case study 1: $P_{day} = 645220.08$ $[kW]$ instead of 64883.84 $[kW]$
- Case study 2: $P_{day} = 298425.69$ $[kW]$ instead of 289525.25 $[kW]$

Stage	*Load*	X-OCL		DP		X-OCS		
		kW	*chiller*	*kW*	*chiller*	*kW*	*chiller*	$[PLR_1, PLR_2, PLR_3, PLR_4, PLR_5, PLR_6, PLR_7, PLR_8, PLR_9]$
1	6210	2899.68	8 6 3 2 1	2899.69	8 6 3 2 1	2899.68	8 6 3 2 1	[1.00 1.00 0.97 0.00 0.00 1.00 0.00 1.00 0.00]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
21	6150	2848.50	8 6 3 2 1	2852.34	8 6 3 2 1	2848.50	8 6 3 2 1	[1.00 1.00 0.92 0.00 0.00 1.00 0.00 1.00 0.00]
22	6280	2974.84	8 6 3 2 1 9	3039.86	8 6 3 2 1 7	2974.84	8 6 3 2 1 9	[1.00 1.00 0.72 0.00 0.00 0.80 0.00 1.00 0.50]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
27	7355	3658.93	8 6 3 2 1 7	3672.69	8 6 3 2 1 7	3658.93	8 6 3 2 1 7	[1.00 1.00 0.93 0.00 0.00 1.00 0.95 1.00 0.00]
28	7520	3768.28	8 6 3 2 1 7 9	3847.81	8 6 3 2 1 7 9	3768.28	8 6 3 2 1 7 9	[1.00 1.00 0.79 0.00 0.00 1.00 0.73 1.00 0.50]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
34	8720	4659.67	8 6 3 2 1 7 9	4659.65	8 6 3 2 1 7 9	4659.67	8 6 3 2 1 7 9	[1.00 1.00 0.98 0.00 0.00 1.00 1.00 1.00 1.00]
35	8910	4791.00	8 6 3 2 1 7 9 4	4891.13	8 6 3 2 1 7 9 4	4791.00	8 6 3 2 1 7 9 4	[1.00 1.00 0.85 0.74 0.00 1.00 0.83 1.00 0.70]
36	9090	4917.89	8 6 3 2 1 7 9 4	5015.07	8 6 3 2 1 7 9 4	4917.89	8 6 3 2 1 7 9 4	[1.00 1.00 0.88 0.77 0.00 1.00 0.87 1.00 0.74]
37	9245	5032.60	8 6 3 2 1 7 9 4	5125.26	8 6 3 2 1 7 9 4	5032.60	8 6 3 2 1 7 9 4	[1.00 1.00 0.91 0.80 0.00 1.00 0.91 1.00 0.78]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
47	9820	5511.19	8 6 3 2 1 7 9 4	5552.25	8 6 3 2 1 7 9 4	5511.19	8 6 3 2 1 7 9 4	[1.00 1.00 1.00 0.90 0.00 1.00 1.00 1.00 0.96]
48	9870	5552.25	8 6 3 2 1 7 9 4 5	5582.17	8 6 3 2 1 7 9 4	5552.25	8 6 3 2 1 7 9 4 5	[1.00 1.00 0.87 0.76 0.67 1.00 0.86 1.00 0.73]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
63	8780	4703.10	8 6 3 2 1 7 9 4	4804.11	8 6 3 2 1 7 9 4	4703.10	8 6 3 2 1 7 9 4	[1.00 1.00 0.83 0.73 0.00 1.00 0.80 1.00 0.66]
64	8555	4511.80	8 6 3 2 1 7 9	4543.81	8 6 3 2 1 7 9	4511.80	8 6 3 2 1 7 9	[1.00 1.00 0.96 0.00 0.00 1.00 1.00 1.00 0.88]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
67	7740	3907.96	8 6 3 2 1 7 9	3985.58	8 6 3 2 1 7 9	3907.96	8 6 3 2 1 7 9	[1.00 1.00 0.81 0.00 0.00 1.00 0.76 1.00 0.62]
68	7455	3728.52	8 6 3 2 1 7 9	3745.20	8 6 3 2 1 7	3728.52	8 6 3 2 1 7 9	[1.00 1.00 0.77 0.00 0.00 0.98 0.71 1.00 0.50]
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
72	6425	3055.91	8 6 3 2 1 9	3115.53	8 6 3 2 1 7	3055.91	8 6 3 2 1 9	[1.00 1.00 0.74 0.00 0.00 0.89 0.00 1.00 0.50]
P_{day}		**296440.08** kW		**289525.25*** kW		**295605.88** kW		

TABLE 5.7: Case study 3

As a matter of fact, the daily power consumption values reported in the paper did not match with the sums of the power consumptions at each step, which we used for the correction.

5.6 Validation on field data

So far, the performances of alternative methods have been compared on OCL and OCS benchmark problems whose quadratic consumption models were taken from the literature. Moreover, a limited number of loads were considered.

In this section, the feasibility of HVAC efficient management based on the exact solution of the OCL problem is validated against a 2-year long real-world scenario that includes the estimation of the chillers' consumption models from the field data introduced in Chapter 2. Quarter-hourly data were available at different working conditions over a period of almost 2 years, from February 2017 to January 2019. Collected data included: temperatures PLR cooling capacities power consumption

The time series of the cooling load demand that the five-chillers system had to satisfy is shown in Fig. 5.4

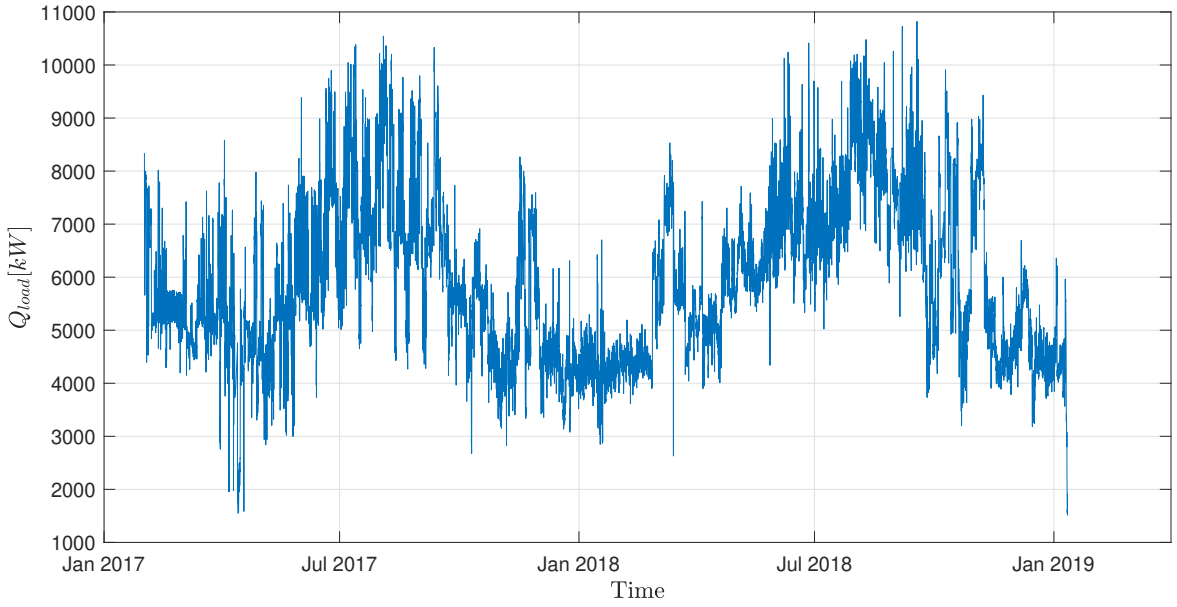

FIGURE 5.4: Cooling load demand time series

The chillers were subject to the following operating constraints:

- $PLR_{min} = 0.2$
- $MUT_i = 4 \quad \forall i$

- $MDT_i = 2 \quad \forall i$

5.6.1 Chiller energy consumption models

The chillers power consumption models were estimated using the evaporator cooling capacity Q_{evap} [kW] and the condenser inlet water temperature T [$^\circ C$], as covariates, and the compressor power consumption data P [kW], as target.

For each chiller, the dataset was randomly partitioned in two datasets: 70% for training and 30% for testing, respectively. The parameters $\beta_{p,i}$ of the model (5.2) were estimated via least squares fitting of the training data, discarding data with $PLR < PLR_{min} = 0.2$. The values of the estimated parameters $\beta_{p,i}$ are reported in Table 5.8 together with their percent coefficient of variation, defined as $\text{CV}\% = 100 \times \text{SE}(\beta_{p,i})/|\beta_{p,i}|$.

TABLE 5.8: Estimated parameters of the quadratic power consumption model (5.2) and corresponding $CV\%$

Chiller	$\beta_{0,i}$ ($CV\%$)	$\beta_{1,i}$ ($CV\%$)	$\beta_{2,i}$ ($CV\%$)	$\beta_{3,i}$ ($CV\%$)	Q_{nom} [kW]
1	49.6367 (5.24)	124.7681 (2.50)	269.5769 (0.82)	3.1950 (3.21)	2700
2	56.2047 (11.97)	154.9900 (7.98)	278.0211 (2.95)	2.2997 (8.32)	2700
3	-10.9883 (22.49)	520.3125 (0.85)	-45.1983 (6.98)	-0.3309 (24.57)	2700
4	-159.0637 (6.15)	112.2988 (19.80)	15.7524 (8.21)	14.1461 (0.74)	2700
5	46.7748 (27.74)	461.7221 (1.84)	-8.3474 (79.87)	-1.1835 (47.23)	2700

In Figure 5.5 the surfaces $P(PLR, T)$ are displayed against the validation data for each of the five chillers. It is seen that the the quadratic model, in spite of its simplicity, predicts well the consumptions at different operating conditions, as also confirmed by the Goodness-of-Fit (GOF) plots in Figure (5.6).

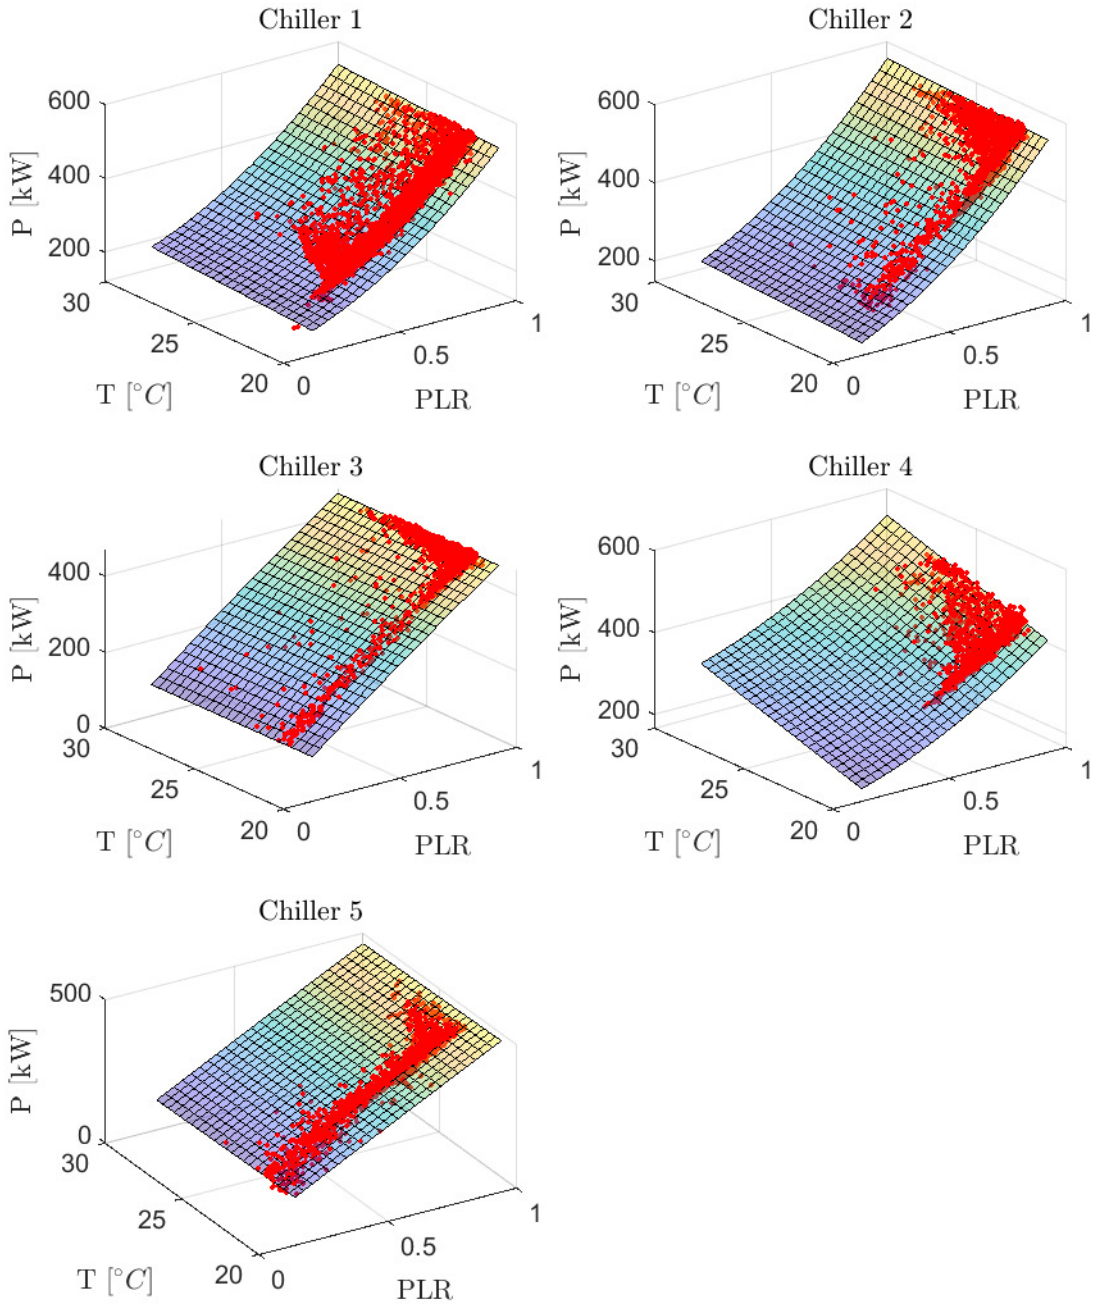

FIGURE 5.5: Three-dimensional representation of power consumption models. Red dots represent the experimental data.

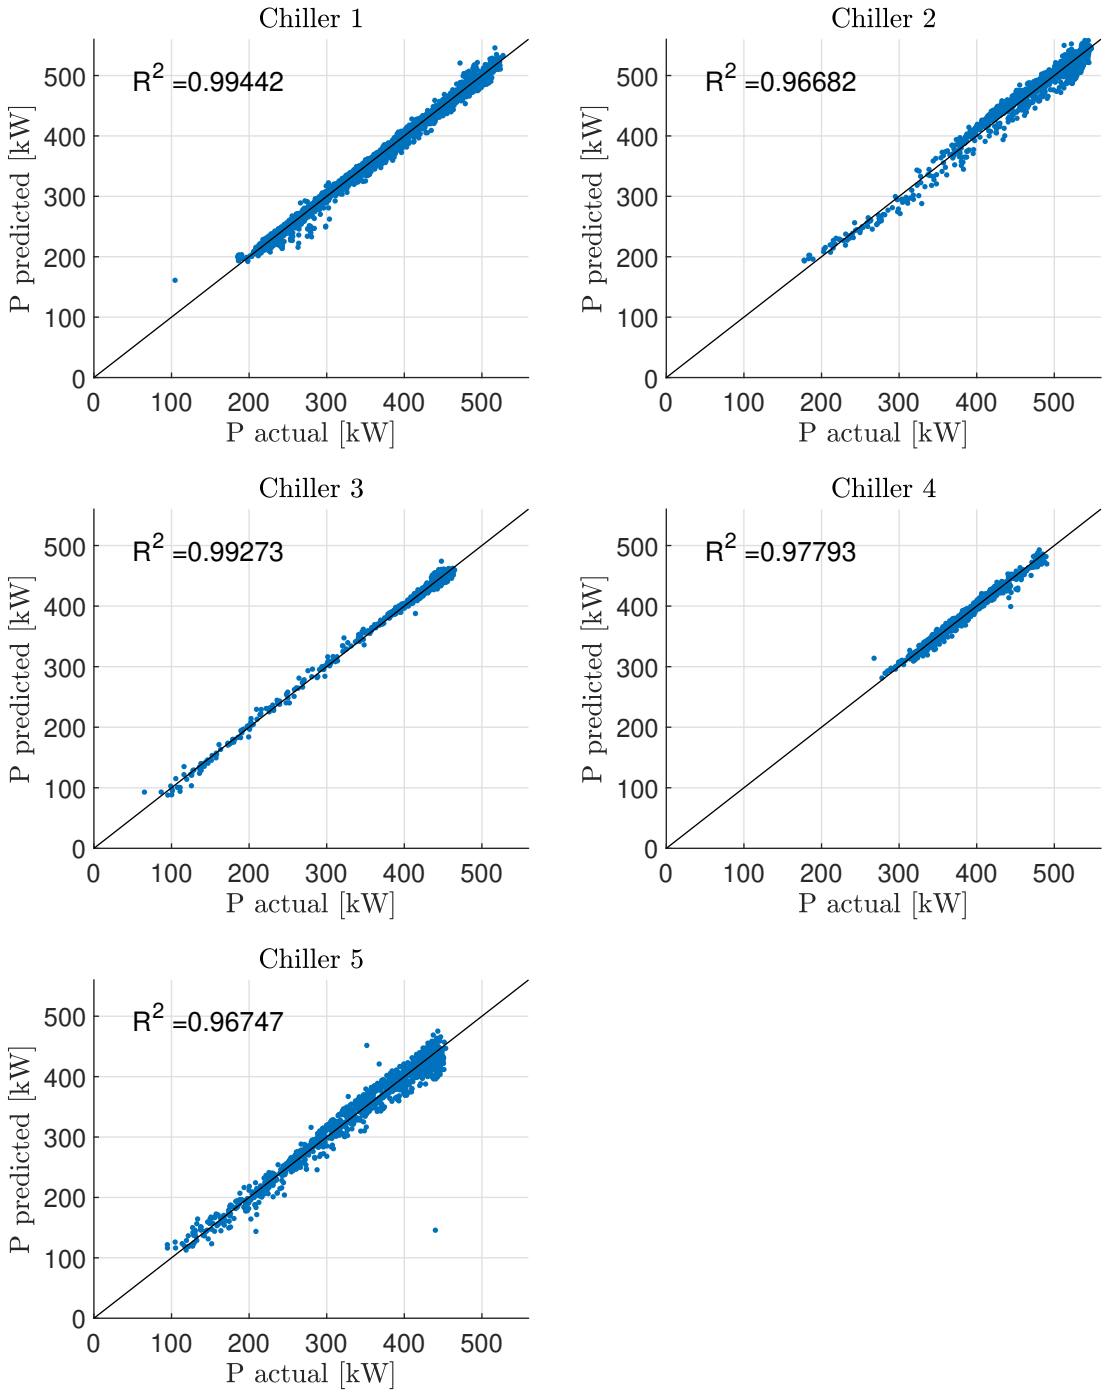

FIGURE 5.6: Goodness-of-Fit (GOF) plots of the quadratic power consumption models. Blue dots are the validation data.

Concerning the role of the condenser inlet water temperature T, the 3-D histogram of the covariates (Q_{evap}, T) in Fig. 5.7 shows that T is mainly concentrated in a narrow range centered

around the setpoint, which is 21.5 $°C$. Following what usually done in the literature benchmarks, one could neglect temperature variations around the set point and solve the OCL and OCS problems using the chillers consumption models at 21.5 $°C$, reported in Figure 5.9. However, the inspection of Fig. 5.5 shows that for some chillers the power consumption is significantly affected by the temperature, especially in summer. Therefore, differently from other literature studies, OCL and OCS solutions were computed based on the complete model $P(PLR, T)$.

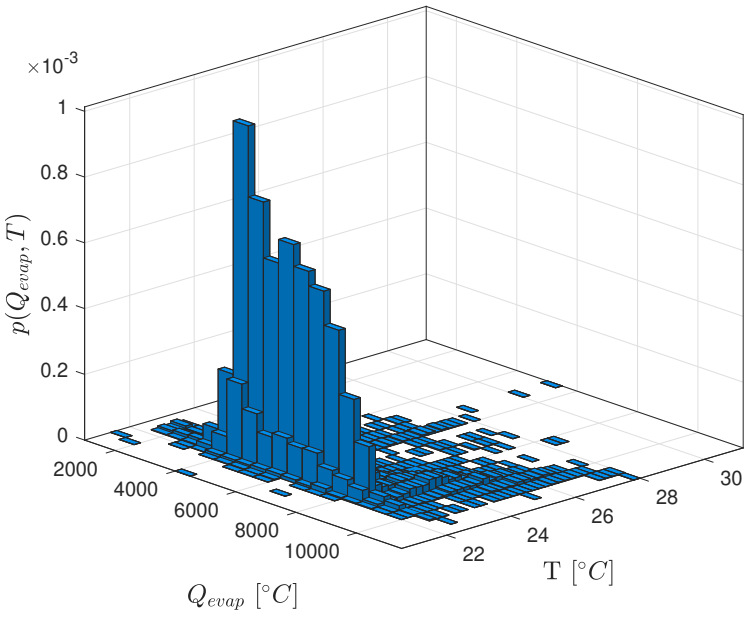

FIGURE 5.7: 3-D histogram of the pairs (Q_{load}, T) recorded from February 2017 to January 2019.

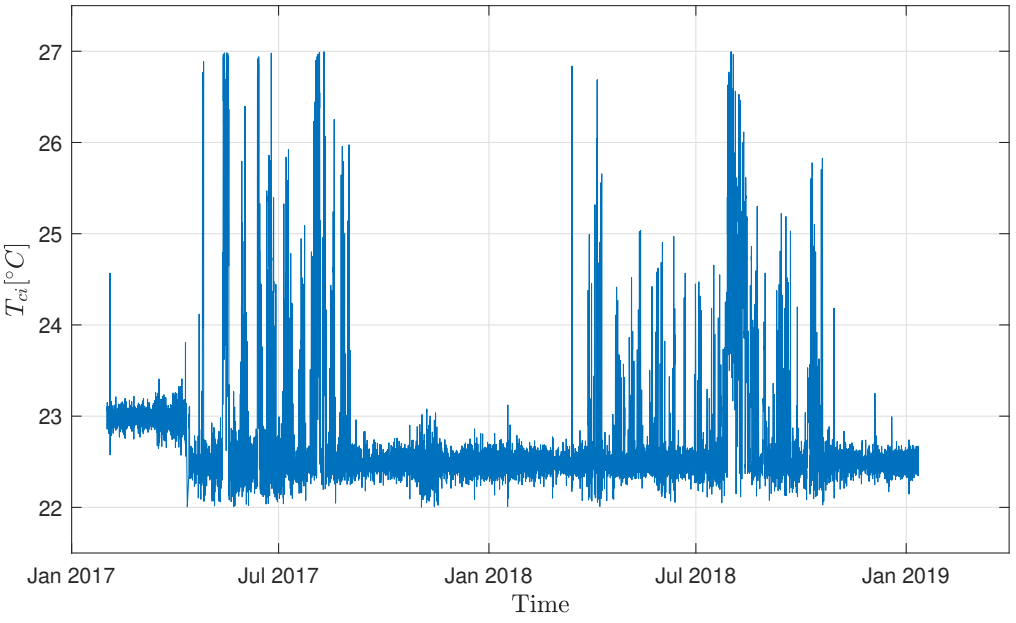

FIGURE 5.8: Condenser inlet water temperature time series

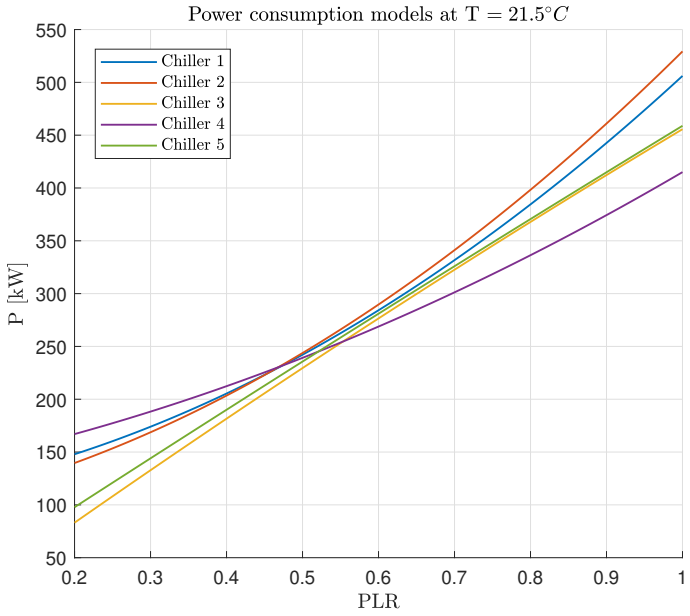

FIGURE 5.9: P-PLR curves of Field data benchmark at $T = 21.5\ ^{\circ}C$.

5.6.2 Real HVAC system: assessment of potential savings

The field data were used to perform a retrospective analysis of the efficiency of the HVAC system management. More precisely, starting from the historical decisions and the associated cumulative power consumption, two comparisons were performed for the 2-year OCS problem. First of all, the lower bound R-OCS on the best achievable consumption was computed in order to quantitatively assess the potential improvement margin. Being a lower bound, R-OCS may be overly optimistic, so that it is important to evaluate the performance that can be obtained in practice. This was done by running the X-OCS solver, whose energy consumption could then be compared with the (ideal) R-OCS bound and the historically recorded power consumption.

The cumulative consumption recorded during the 2-year monitoring was $1.758 \times 10^6[kW]$. This figure can be compared with the R-OCS lower bound, equal to $1.600 \times 10^6[kW]$. This means that the potential margin of improvement is not larger than 8.97%.

When the X-OCS algorithm was applied, the cumulative power consumption was $1.601 \times 10^6[kW]$. As a matter of fact, for this HVAC system, the loss of performance due to the suboptimality of the greedy algorithm is definitely negligible (it is less than 0.1%). In Figures 5.10 and 5.11, it is seen that, on a weekly basis, the difference between the consumption achieved by X-OCS and the lower bound is always less than $15[kWh]$.

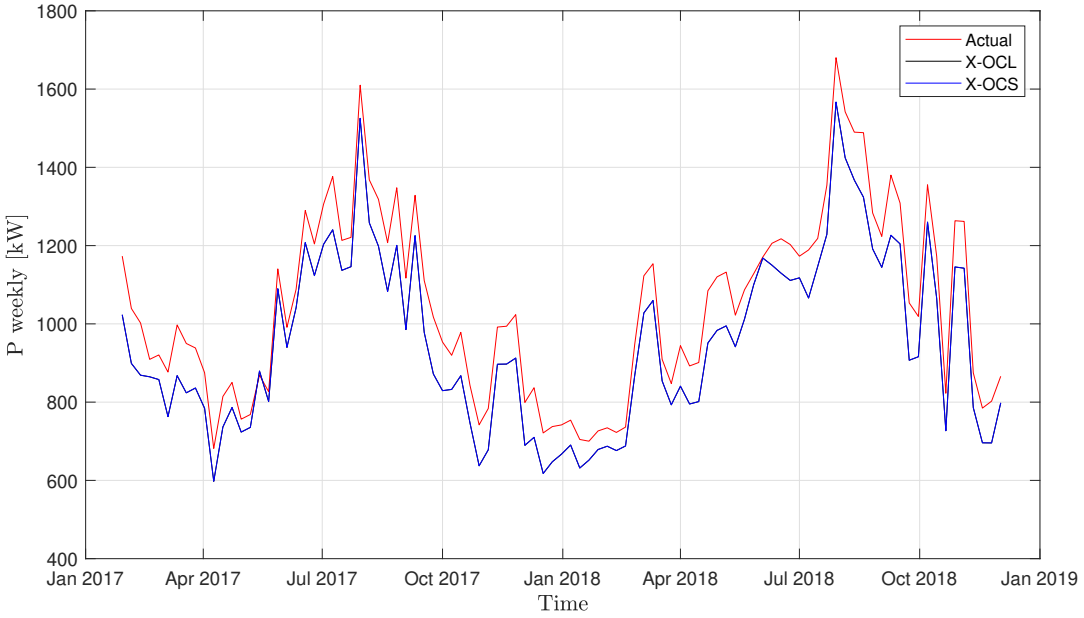

FIGURE 5.10: Weekly power consumption time serie: comparison bewteen actual recorded data (red), the lower bound R-OCL (black) and the consumption associated with X-OCS (blue).

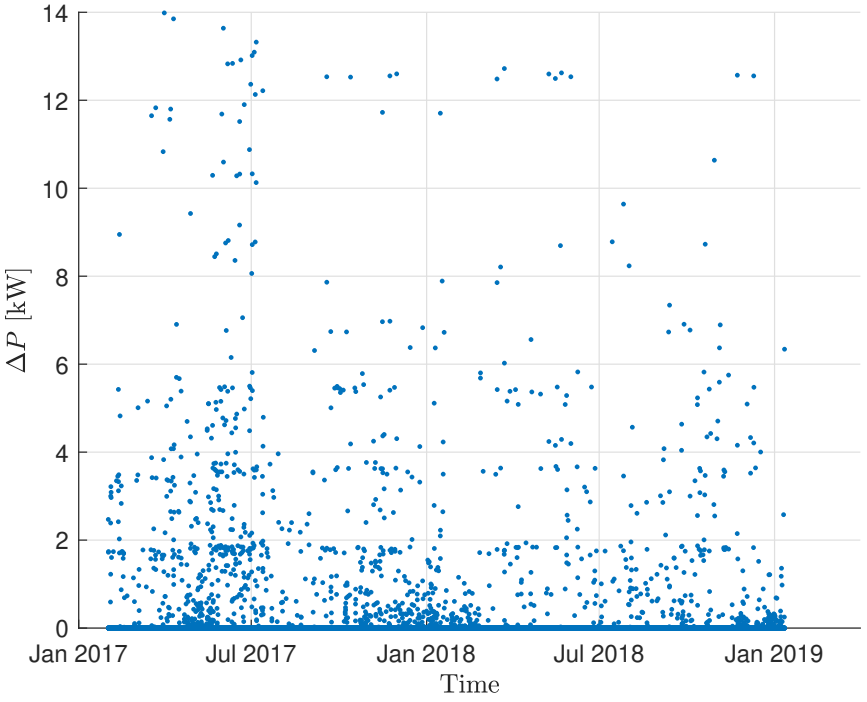

FIGURE 5.11: Difference between the actual recorded power consumption and the one associated with X-OCS (hourly data)

5.7 Execution time

The X-OCL and X-OCS algorithms, coded in Matlab®, were executed on a standard laptop (Intel(R) i7-7500U dual-core with hyperthreading, RAM 16GB, 2.7 GHz). While no explicit parallelization of the algorithm was implemented, the solution of the 4^n QP problems was formulated as a unique algebraic computation using sparse matrices. This means that the algorithm may have benefited from some optimization automatically enforced by the Matlab® compiler. For the sake of comparison, the Hsinchu benchmark was also solved using the CPLEX solver for costrained mixed integer problems as implemented in the GAMS environment, used by [43].

Concerning the Hsinchu OCL benchmark, the computational cost of X-OCL was negligible ($0.11s$) and, more importantly definitely lower than the time ($1.99s$) spent by GAMS.

Coming now to the OCS benchmarks, the larger number of loads makes it possible to asses the average computation time per load. Not surprisingly, the computation time for X-OCS is larger than that for R-OCS, because the greedy algorithm performs the additional task of looking for solutions that satisfy the up/down-time constraints.

For a small number of chillers, the average computation time per load of X-OCS is remarkably small: $0.29/72 = 4 \times 10^{-3}$s for OCS benchmark 1 (4 chillers) and $497.13/68110 = 7 \times 10^{-3}$s for the OCS on field data (5 chillers). In the latter case, X-OCS took less than 6 min to solve the OCS problem over 2-year data with quarter-hour sampling.

As expected, in view of the exponential growth of the number of QP problems, the maximal average computation time per load is found in correspondence of OCS benchmark 2, where 9 chillers are present. Nevertheless, the average time per load amounts to $1.96s$, which is totally affordable. In Fig. 5.12 the average computation time per load is displayed as a function of the number of chillers. The exponential growth is apparent but even for a large HVAC system made of 9 chillers the computational cost is all but prohibitive.

Table 5.9: Execution times for OCL and OSC benchmarks

	N° chillers	N° loads	X-OCL [s]	R-OCS [s]	X-OCS [s]	GAMS [s]
Hsinchu benchmark	6	5	0.722	-	-	1.991
OCS benchmark 1	4	72	-	0.302	0.368	-
OCS benchmark 2	9	72	-	226.77	228.29	-
Field data	5	68110	-	354.03	573.96	-

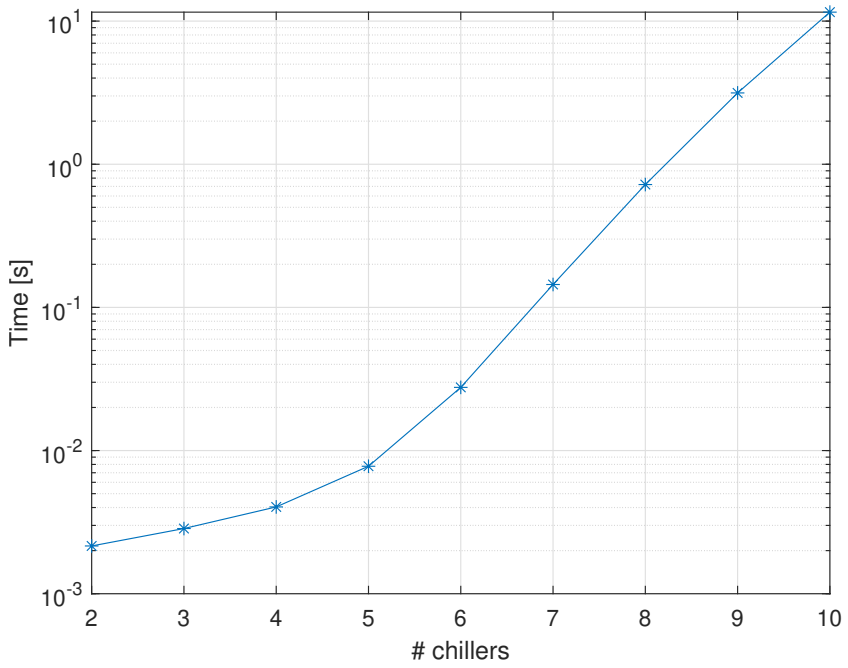

FIGURE 5.12: Average computation time per load

5.8 Discussion

The main purpose of the chapter was the derivation of an exact algorithm for the solution of the OCL problem. This goal was successfully completed by a decomposition approach that exploits a suitable partition of the solution space.

The availability of an exact method has been immediately exploited along two directions. First, it became possible to give a definitive assessment on the performances of different literature methods that had been applied to some consolidated benchmark problems. By the way, the comparison with the exact solution revealed that some unrealistic performances had been declared in the literature, due to erroneous extrapolations of the power consumption curve that became negative at low part loads ratios. The exact method has been exploited also in relation with the optimal chiller sequencing problem. For a given cooling demand profile, if the dynamic up/down-time constraints are neglected, a sequence of OCL problems can be exactly solved to yield a lower bound limit, called R-OCS, to the optimal performance achievable by any method complying with the dynamic constraints. For a given OCS algorithm, the comparison of its performance with the bound provides an assessment of the potential margins for further improvement.

An outcome of this study is the remarkable numerical efficiency of the proposed exact algorithm. This raises the question of its practical applicability in place of heuristic approaches discussed

in the literature. Two major objections may be posed to a generalized adoption of the exact solution: (i) the explosion of the computational cost with the number of chillers, (ii) the need to make overly restrictive assumptions on the shape of the power consumption curves.

On the first side, it is true that the OCL problem as formulated in 5.1 is NP-hard, which implies an exponential growth of the computations needed for its solution. This is confirmed also by Fig. 5.12, where the experimental growth of computation time as a function of the number n of chillers is exponential (it is linear in the semilog scale). At the same time, the figure shows that even for a medium/large-sized chiller system, e.g. 6-9 chillers, the computational cost for a single load is less than 2 s. This suggests that even for chiller systems used in large semiconductor factories, the cost of the exact solution is not prohibitive. The numerical efficiency is a direct consequence of the tiny number of computations required to solve the elementary EQP problems, see 5.14. Moreover, the partition strategy underlying X-OCL, implies that it is totally parallelizable into 4^n threads, a feature that has not been explicitly exploited and that could further speed up the solution.

The second objection to the general applicability of the exact solution has to do with the quadratic assumption made on the power consumption curve. Even if the majority of benchmarks share this assumption, there is no stringent reason to rule out other functional descriptions. Nevertheless, when confronted with real data, see the Field data benchmark, we found that a quadratic power consumption fitted well the recorded data, see Figure 5.5. Even when a single quadratic function were not adequate, it would still be possible to switch to a piecewise quadratic description. In that case, it would be rather immediate to generalize the exact algorithm by just increasing the number of partitions $\mathcal{S}_j$, in such a way the problems still boils out to a set of easy-to-solve EQP problems.

In view of its numerical efficiency, also the application of X-OCS to the solution of the OCS problem appears very promising. Indeed, in the OCS benchmarks taken from the literature and in the Field data OCS benchmark, the performance of X-OCS is very close to the lower bound, implying that there is no scope for the use of more sophisticated algorithms.

6

Thermal Comfort in Air-Conditioned Buildings

The Heating, Ventilation and Air-Conditioning (HVAC) systems are responsabile for about 30% of Greenhouse Gas Emissions (GHG) [59]. Progress towards sustainable buildings is advancing, see e.g. the concept of Nearly Zero Energy Buildings NZEB [60], but improvements are still not keeping up with a rising demand for energy services. The energy intensity per square meter of the global buildings sector needs to improve on average by 30% by 2030 (compared to 2015) to be on track to meet global climate goals set forth in the Paris Agreement [61]. In this context, the human factor also plays a key role in terms of occupants' behavior and needs. As a matter of fact, the main challenge for the HVAC control is the trade-off between building energy consumption and occupants' thermal comfort.

Two major types of control strategies have been proposed for heating and cooling systems. The first one is based on different thermostat adjustments such as simply setting a higher summer set point temperature (SST) or implementing a wide range of indoor design temperatures in correspondence of different day hours and outdoor conditions. The second strategy is to dynamically adjust the set point temperature based on adaptive comfort models.

In the 70s, the so-called adaptive theory introduced the concept of physiological adaptation (in terms of acclimatization). This means that people's thermal comfort sensation changes with a number of context variables, such as the indoor and outdoor air temperatures. This dependence can be exploited to save energy and costs and nevertheless maintain the building thermal comfort. For example, raising indoor temperatures in summertime not only would reduce the cooling requirements, but would also lead to widespread energy savings, in that such intervention can be applied to both new and existing buildings [62]. For these reasons, the availability of effective comfort models able to predict the thermal sensations of the occupants may give a valuable contribution to the efficient management of HVAC systems.

6.1 Neutral Temperature definition and estimation

Thermal comfort is defined as *"the state of mind, which expresses satisfaction with the thermal environment'*; a definition easily understandable, but hard to capture in terms of physical parameters. An essential requirement for its fulfillment is providing an indoor air temperature regarded as satisfactory by the majority of occupants. In the indoor thermal comfort literature, such a temperature goes under the name of *neutral temperature*. Notably, its semantics has undergone several changes over the years due to the debate between conventional and "adaptive" schools of thought. The conventional or "static" model views occupants as passive recipients of thermal stimuli driven by the physics of the thermal balance between the body and the environment.

Experimental measures of thermal comfort are usually expressed according to the ASHRAE thermal sensation scale: -3 cold, -2 cool, -1 slightly cool, 0 neutral, +1 slightly warm, +2 warm, +3 hot. For the development of thermal control strategies it is essential to have a thermal comfort model capable of predicting the occupants' comfort as a function of context variables such as air temperature, mean radiant temperature, air speed, humidity, metabolic rate and clothing insulation. Given the model, the neutral temperature can be searched for in correspondence of the other context variables, providing the set point for the HVAC system.

The most popular model of thermal comfort is Fanger's *Predicted Mean Vote (PMV)* [63], a nonlinear model that predicts the expected comfort vote on the ASHRAE thermal sensation scale in dependence of a number of context variables. The model parameters were calibrated by Fanger based on the thermal sensation votes of subjects who were subject to standardized clothing and activities, not necessarily reflecting the variable conditions of everyday life.

In fact, as observed by Nicol and Humphreys [7], the predictions of the PMV model proved no better than those obtained with a much simpler model that used the indoor air temperature as the only independent variable. A possible explanation was formulated in terms of the so-called adaptive theory, according to which people are able to adapt to the environment climate where they stay [64, 65].

Subsequent work by Humphreys [66] found that the thermal comfort depended not only on the indoor temperature, but, to some extent, also on the outdoor climate. In order to validate these findings a number of meta-analysis were conducted using field surveys coming from a wide range of environments [66–69]. In this framework, de Dear [69] proposed the following method for predicting the neutral indoor air temperature which has been thereafter widely adopted. Starting from ASHRAE scores collected in different buildings the procedure derives a neutral temperature characteristic of each specific building. The following steps refer to a generic building:

1. The indoor air temperature is divided into bins of half-degree and the average value $ASH(k)$ of individual subjects' thermal votes is computed for the k-th bin.

2. A weighted linear regression model is fitted linking the mean sensation vote to the indoor air temperature:
$$ASH(k) = f(T_{in}(k)) = \alpha + \beta T_{in}(k)$$

3. The neutral temperature T_N is derived by solving the regression line for a mean vote of zero.
$$T_N = -\alpha/\beta$$

4. The outdoor temperature T_{out} is computed as the average of recorded outdoor temperatures.

The procedure yields pairs of outdoor and neutral temperatures $(T_{out}(i), T_N(i)), i = 1, \ldots, N_b$, where N_b is the number of buildings. These pairs can then be used to identify a model that, according to the adaptive hypothesis, links the neutral temperature to the outdoor one, i.e $T_N = g(T_{out})$. The simplest model is linear with positive slope which is consistent with the idea that if the outdoor climate is hot, the indoor thermal comfort is achieved in correspondence of a higher indoor temperature.

From the statistical point of view, the above procedure presents several shortcomings. In particular, is it a good practice to average the data over the bins? If the final goal is prediction of the neutral temperature as a function of the indoor and outdoor air temperature, is it really necessary to go through two distinct steps (regression of ASH on T_{in} and then regression of T_N on T_{out})? Does the adaptive effect, which explains the occupants' neutral sensation as a function of the outdoor air temperature, actually exist or is it an artifact of the indirect procedure?

In order to answer these questions, in the present chapter a direct modeling approach is proposed to establish the possible dependence of the neutral temperature on the outdoor one. In particular, binning is avoided and the building by building approach is replaced by the direct use of individual ASH scores and outdoor temperatures avoiding averaging at building level.

6.2 Thermal comfort survey data

The present study uses experimental data drawn from the ASHRAE RP-884 project database, which consists of about 21,000 sets of thermal comfort data pooled from a number of surveys conducted by several research groups around the world [70]. The surveys included *basic identifiers* (i.e. building code, subject information and date), *thermal questionnaire responses* (i.e. thermal sensation, acceptability, metabolic rate and clothing insulation), *calculated indices* (i.e.

operative temperature, new effective temperature and predicted mean vote) and *indoor and outdoor climate observations* (i.e. air temperature, speed and humidity). The surveys involve Naturally Ventilated (NV), Air-Conditioned (HVAC), and Mixed Mode (MM) buildings, the majority of which contain offices (136 building out of 160).

In the present chapter, only surveys carried out in summer for HVAC buildings are considered. The detection of contradictory responses was used to identify and remove misleading questionnaires. An example of a contradictory response is when a participant declares that the thermal environment is not acceptable but, in the preference question, does not desire any change. The selected datasets (Jakarta (06), Montreal (09), Brisbane (11), Melbourne (15), San Francisco (32), Singapore (41) and San Ramon(44)), were put together.

A preliminary analysis showed that the occupants' thermal sensation vote ASH is correlated with the indoor and the outdoor air temperature, while it is only marginally correlated with the other independent variables (i.e. air speed, air humidity, clothing insulation, etc.), which confirms the adaptive hypothesis recalled in Section II. Accordingly, our subsequent analysis is focused on the triplets (T_{in}, T_{out}, ASH). The ASH scores exhibit high variability, see Figure 6.1. For instance, for almost any indoor temperature a wide range of votes is found as a consequence of the subjectivity of thermal sensations.

TABLE 6.1: Sources of raw data for the RP-884 databaset

Researcher	**Experiment location**	**Sample size**	N_b
Tri Karyono (Sheffield, UK)	Jakarta, Indonesia	458	5
de Dear (PhD data)	Brisbane, Australia	564	5
Donnini ASHRAE RP-821	Montreal, Canada	443	12
de Dear (PhD data)	Melbourne, Australia	512	4
Brager ASHRAE RP462	Bay Area, California	673	7
de Dear, Foo and Leow	Singapore	333	1
Benton + Brager (ACT^2)	San Ramon, CA	96	1

6.3 Exploratory analysis of thermal comfort data

In this section a simple approach for exploring the possible link between the outdoor temperature and the indoor comfort temperature is illustrated. In order to verify whether the adaptive hypothesis holds true, we consider a subset of the original data, selecting only the occupants whose thermal sensation vote is equal to zero ($ASH = 0$). For this subset, the scatter plot of the indoor temperature against the outdoor one is displayed in Fig. 6.3. Note that these

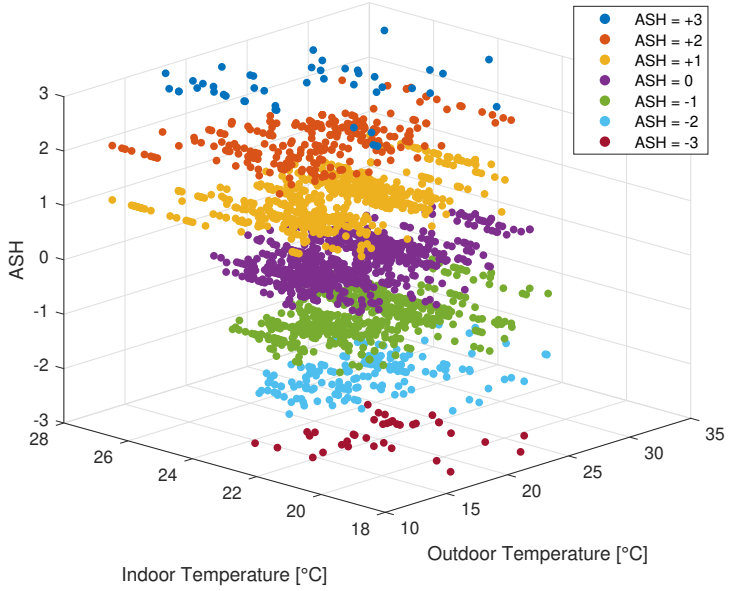

FIGURE 6.1: Overall summer dataset. Coloured dots: occupants' thermal sensation votes on the ASHRAE scale [-3; +3] as a function of T_{in} and T_{out}.

indoor temperatures are regarded as comfortable by the occupants. Therefore, if the outdoor temperature affects the occupants' sensibility, some dependence between the indoor and outdoor temperatures should be observed.

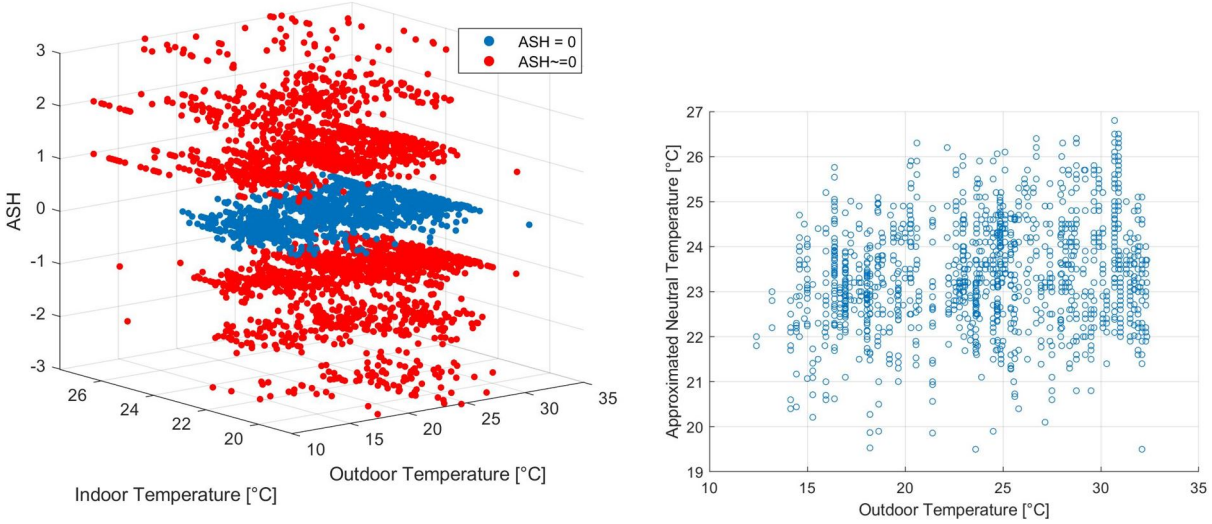

FIGURE 6.2: Left: blue dots represent the selected data corresponding to ASH = 0 while the red ones the remaining. Right: the blue dots, displayed in the three dimensional space in the left panel, are here shown in the plane Outdoor-Indoor air temperature.

In spite of the fact that only indoor temperatures associated with null ASH score are being considered, the data is still remarkably noisy. For a given outdoor temperature, there is a wide range of indoor temperature that are regarded as neutral. It seems therefore reasonable to use regression models of low complexity to predict the neutral temperature as as a function

of the indoor one. Herein, three methods are used: linear regression, Lowess regression, and MLP Artificial Neural Network (ANN) with McKay regularization. The first two methods have already been used in the context of neutral temperature modeling [71].

A Lowess regression of linear type with 0.25 span (i.e proportion of data points used) was employed. As for the MLP ANN, a one-hidden-layer network of fully connected neurons with activation function *tanh* and linear output was considered. When a fully connected structure is assumed, overparametrization may occur even for a small number of neurons. Following [28], in order to prevent overfitting, a quadratic regularization penalty on the weights is added to the conventional squared loss. The regularization parameter was tuned according to McKay's criterion and the number of neurons decided through crossvalidation with $70-30$ split ratio.

The three regression curves, displayed in Figure 6.3, exhibit a good agreement. Some differences are observed for low (12) and high (32) temperatures, where linear regression predicts a higher comfort temperature. The main finding is the presence of a positive slope which is in keeping with the adaptive theory according to which people adapt, to some extent, their thermal comfort to the outdoor air temperature. Indeed, when outdoor temperature increases, the indoor comfort temperature increases as well: while for $T_{out} = 20°C$ the predicted neutral temperature is about $23°C$, when $T_{out} = 30°C$, the predicted neutral temperature becomes $23.5°C$. This could be explained either with adaptation to the climate or with adaptation of clothing to the seasons. In terms of HVAC energy consumption, this result could be exploited to set a higher value of the indoor temperature set point when, for instance, $T_{out} = 30°C$, thus achieving some energy saving.

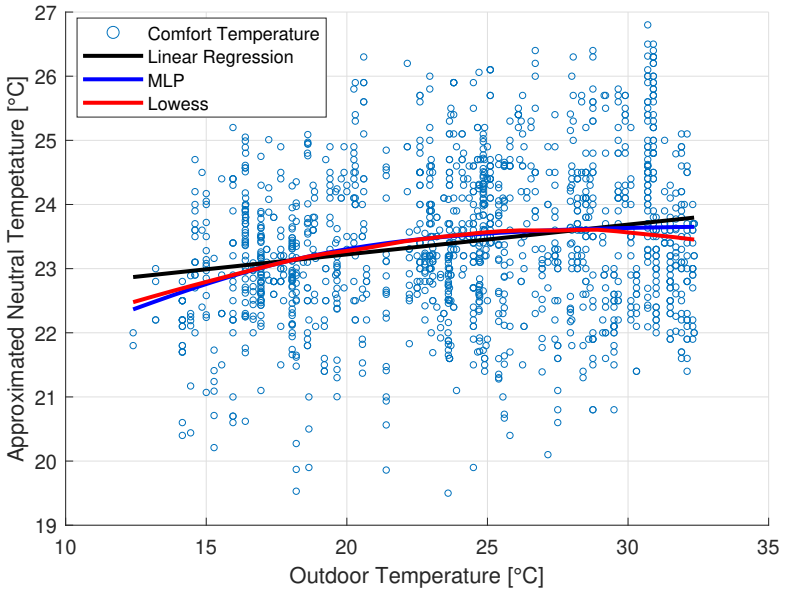

FIGURE 6.3: Blue dots: data points (T_{out}, T_{in}) corresponding to $ASH = 0$; Black line: simple linear regression model; Blue line:Lowess regression; Red line: MLP Neural Network model.

In view of this exploratory analysis, there is scope for a more systematic identification of the dependence of the neutral temperature on the outdoor one. Differently from what done in this section, hereafter we will take into account the complete dataset, instead of just the occupants whose vote is neutral.

6.4 Direct Modeling of Thermal Comfort

A main shortcoming of the standard approach to neutral temperature modeling reviewed in Section 6.1 is the need of two steps: for each building, neutral and outdoor temperatures are obtained by gathering and processing data at building level which are then used to establish the link between outdoor ad neutral temperature. In this section, a new approach is proposed that is rooted in the very definition of neutral temperature and uses directly all the occupants data, avoiding preprocessing at building level. Recall that, according to the ANSI/ASHRAE Standard 55-2010, the neutral temperature is the indoor air temperature corresponding to a zero mean vote ($\mathcal{E}(ASH) = 0$) on the thermal sensation scale by a sample of building occupants.

Let $y = \mathcal{E}(ASH|T_{in}, T_{out}) = f(T_{in}, T_{out})$ denote the expected vote conditional on T_{in} and T_{out}. Then, according to its definition, the neutral temperature $T_N = T_N(T_{out})$ is just the indoor temperature that, for a given T_{out} yields $y = 0$, that is the solution of the equation

$$0 = f(T_N, T_{out}) \tag{6.1}$$

This observation is the basis of the direct approach that consists in estimating the function $f(\cdot, \cdot)$ from the triplets (T_{in}, T_{out}, ASH). Below, two approaches are considered.

6.4.1 Polynomial regression model

A first approach to estimating $f(\cdot, \cdot)$ amounts to finding a surface that best fits the data according to the least squares criterion.

In particular, a polynomial model was used to predict the occupants' thermal sensation votes (ASH) as a function of T_{in} and T_{out}. The model's regressors were automatically selected by a stepwise regression algorithm which chose the following quadratic model without interaction terms:

$$y = \hat{f}(T_{out}, T_{in}) = \beta_0 + \beta_1 T_{out} + \beta_2 T_{in} + \beta_3 T_{out}^2 + \beta_4 T_{in}^2$$

The estimated surface $f(T_{out}, T_{in})$ is displayed in Fig. 6.4. As expected y depends strongly on T_{in} but, to a less extent, also on T_{out}, thus confirming the adaptive theory. The points associated

with zero expected vote, which are highlighted in red, identify the solution of equation (6.1), which, in turn, provides the neutral temperature model $\hat{T}_N = g(T_{out})$.

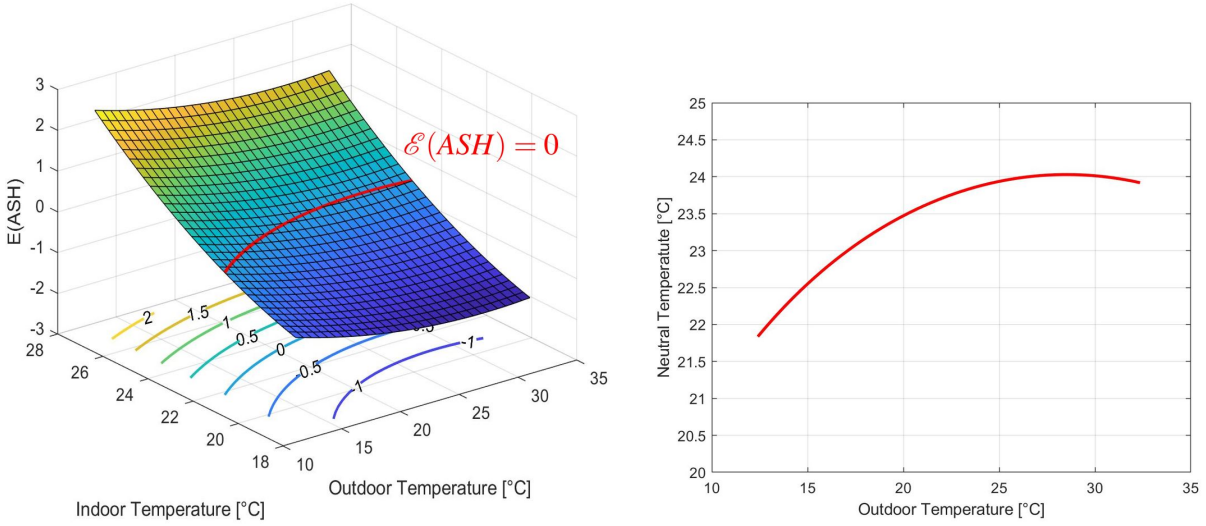

FIGURE 6.4: Polynomial model: Left, occupants' expected mean vote $\mathcal{E}(ASH \mid T_{out}, T_{in})$. Red curve: intersection of the surface with $\mathcal{E}(ASH) = 0$. When projected onto the plane (T_{out}, T_{in}), the red curve provides the neutral temperature model $\hat{T}_N = g(T_{out})$. Right, the estimated neutral temperature curve in the plane outdoor air temperature - indoor air temperature

6.4.2 Multinomial logistic model

A model for the expected vote can also be obtained as a byproduct of a complete model of the occupants' votes. Given the discrete 7-levels scale of thermal sensation votes, this is a multinomial model that, for any pair (T_{in}, T_{out}), returns the seven probabilities

$$P_v(T_{in}, T_{out}) = \Pr(ASH = v), v \in \mathcal{A} \tag{6.2}$$

where $\mathcal{A} = \{-3, -2, -1, 0, 1, 2, 3\}$. From the multinomial model, the function $f(T_{out}, T_{in})$ is immediately obtained as

$$y = f(T_{out}, T_{in}) = \sum_{v \in \mathcal{A}} v P_v(T_{in}, T_{out})$$

As far as the estimation procedure is concerned, in the literature this kind of model has been calibrated through the repeated estimation of binary probit models [72]. However, the ordinal nature of the ASHRAE' votes scale suggests the use of a generalized ordered logit (*gologit*) model [73] that processes all the data simultaneously. Starting from a second order polynomial

model, a backward stepwise regression yielded the following model:

$$\begin{aligned} \Pr(ASH > v \mid T_{out}, T_{in}) &= \frac{\exp\left(\alpha_v + X^T\beta_v\right)}{1 + \left[\exp\left(\alpha_v + X^T\beta_v\right)\right]}, v \in \mathcal{A} \\ X_i &= \begin{bmatrix} T_{out} & T_{in} & T_{in}T_{out} & T_{out}^2 \end{bmatrix}^T \end{aligned}$$

The maximum likelihood estimate was obtained by the *mnrfit* function of the Matlab Statistics and Machine Learning Toolbox.

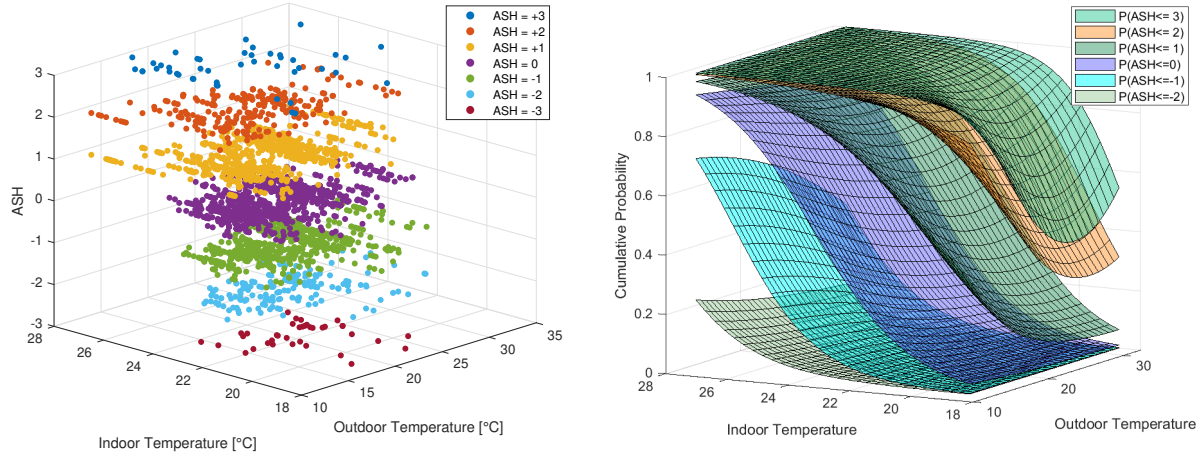

FIGURE 6.5: Left: Overall summer dataset. Right: Logistic cumulative distribution functions

The estimated multinomial model is visualized in Fig 6.5, where the cumulative probabilities $\Pr(ASH \leq v \mid T_{out}, T_{in})$ are plotted as functions of T_{out} and T_{in}. This multinomial model provides the most complete probabilistic description of thermal comfort sensations as it can predict, for any pair (T_{in}, T_{out}), the percent of occupants expressing any of the seven votes. In particular, it can be used to obtain important indicators such as the predicted percentage of dissatisfaction (PPD).

Eq. (6.2) can be used to obtain $y = \mathcal{E}(ASH|T_{in}, T_{out})$ from the multinomial model. The resulting surface is plotted in Fig. (6.6). The overall surface is in fairly good agreement with that in Fig. (6.4). Again, the points solving equation (6.1) are highlighted in red and show an even better agreement with the red curve produced by the polynomial regression model, see Fig. (6.4).

6.5 Alternative definition of Neutral Temperature

The standard definition of neutral temperature relies on the mean vote of occupants. An advantage of such a definition is that it is rather straightforward to obtain a model for the mean

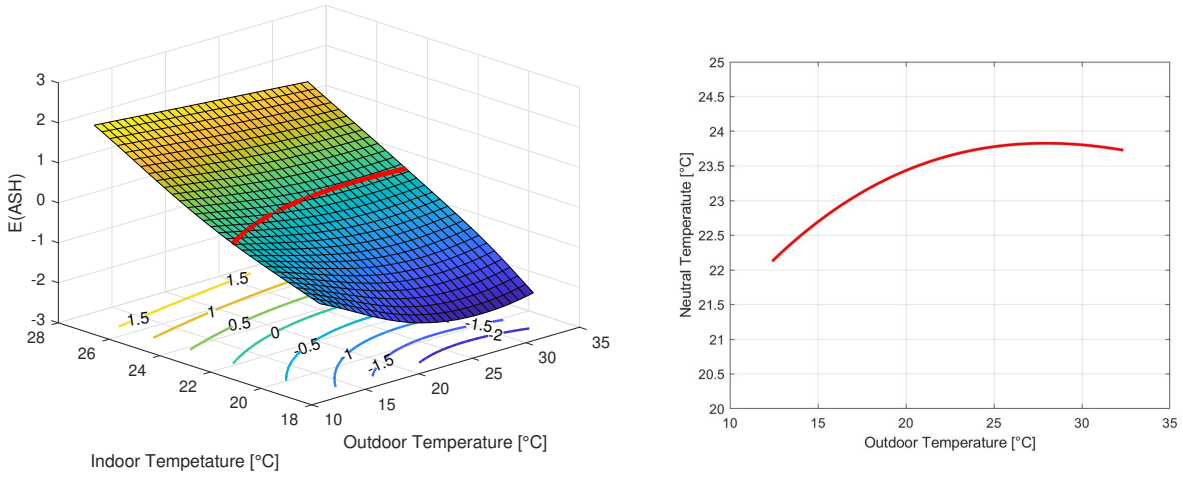

FIGURE 6.6: Ordered logit model. Left panel: occupants' expected mean vote $\mathcal{E}(ASH \mid T_{out}, T_{in})$. Red curve: intersection of the surface with $\mathcal{E}(ASH) = 0$. When projected onto the plane (T_{out}, T_{in}), the red curve provides the neutral temperature model $\hat{T}_N = g(T_{out})$. Right panel: the estimated neutral temperature curve in the plane outdoor air temperature - indoor air temperature

vote just by fitting a surface to the experimental votes, see Section 6.4.1. However, one may be interested in achieving the temperature that maximizes the percentage of satisfied occupants. This calls for the estimation of a probabilistic model of the occupants' thermal sensation votes, a task that has already been successfully performed in Section 6.4.2. Accordingly, we introduce an alternative definition of neutral temperature aimed at maximizing the occupants' consensus.

Definition (probabilistic neutral temperature): For a given T_{out} the probabilistic neutral temperature $T_N^* = T_N^*(T_{out})$ is defined as the indoor air temperature which maximizes the probability of a zero thermal sensation vote:

$$T_N^*(T_{out}) = \arg\max_{T_{in}} \Pr(ASH = 0 \mid T_{out}, T_{in})$$

It goes without saying that the function $T_N^*(T_{out})$ can be easily derived from the *ordered logit model* presented in Section 6.4.2. In Fig. 6.7, where $\Pr(ASH = 0 \mid T_{out}, T_{in})$ is plotted against T_{out} and T_{in}, the red curve highlights the maximum of $\Pr(ASH = 0)$ for different values of T_{out}. Then, $\hat{T}_N^* = g(T_{out})$ is just the projection of the red line on the plane (T_{out}, T_{in}).

Given the complexity of identifying the whole multinomial model, one could resort to a simpler *Binomial logistic model*, which can be seen as a particular case of the multinomial logistic one when just two categories are considered, namely $ASH = 0$ (thermal satisfaction) and $ASH \neq 0$ (thermal dissatisfaction). Of course, it is not difficult to consider a less stringent definition of satisfaction by letting the probabilistic neutral temperature T_N^* be the one that maximizes

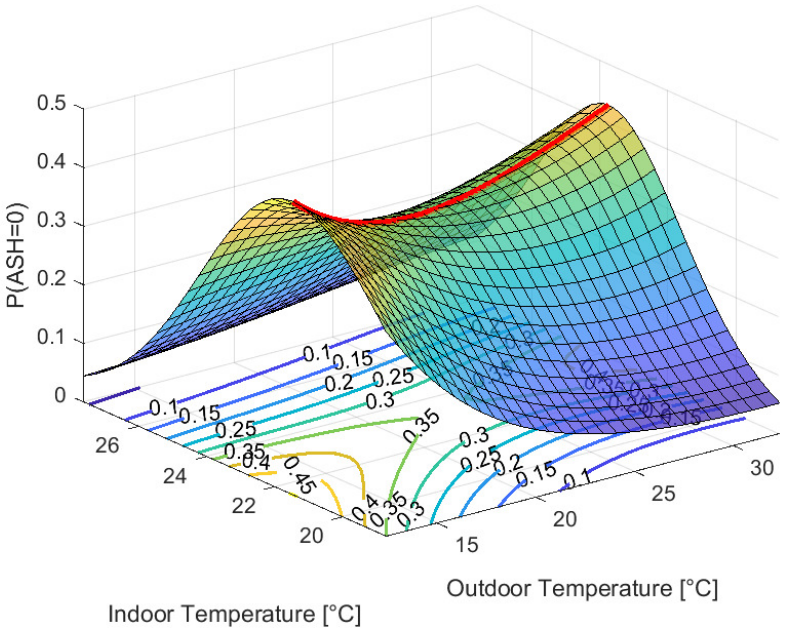

FIGURE 6.7: Ordinal logit model. Surface: $\Pr(ASH = 0 \mid T_{out}, T_{in})$; The projection of the red curve onto the plane (T_{out}, T_{in}) represents the neutral temperature model $\hat{T}_N^* = g(T_{out})$.

the probability $\Pr(|ASH| \leq 1)$, in which case the categories of the binomial model would be $|ASH| \leq 1$ (thermal satisfaction) and $|ASH| > 1$ (thermal dissatisfaction).

Considering polynomials up to the third order, the following model was obtained via stepwise regression using the *stepwiseglm* function of the Matlab Statistics and Machine Learning Toolbox:

$$\Pr(ASH = 0 \mid T_{out}, T_{in}) = \frac{\exp\left(\alpha_v + X^T \beta_v\right)}{1 + \left[\exp\left(\alpha_v + X^T \beta_v\right)\right]}$$

$$X_i = \begin{bmatrix} T_{out} & T_{in} & T_{in}T_{out} & T_{out}^2 & T_{in}^2 \end{bmatrix}^T$$

The function $\Pr(ASH = 0 \mid T_{out}, T_{in})$ is displayed as a 3D surface in Fig. 6.8. Again, the projection of the red curve onto the plane (T_{out}, T_{in}) represents the neutral temperature model $\hat{T}_N^* = g(T_{out})$.

6.6 Comparison and discussion

The main objective of the chapter was the derivation of models of the neutral temperature that improve on the standard literature approach that relies on a preprocessing step at building level. Three approaches have been considered. The exploratory one is based on the regression of temperatures regarded neutral by the occupants against the outdoor temperature. Three regressions methods were tested: linear, Lowess and MLP ANN.

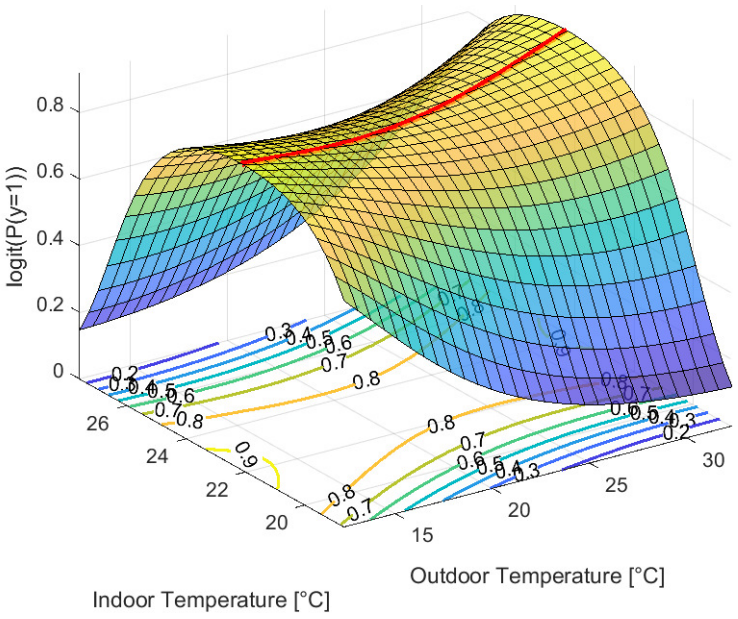

FIGURE 6.8: Binomial logit model. Surface: $\Pr(ASH = 0 \mid T_{out}, T_{in})$; The projection of the red curve onto the plane (T_{out}, T_{in}) represents the neutral temperature model $\hat{T}_N^* = g(T_{out})$.

The second approach, the so-called direct one, is rooted in the very definition of neutral temperature, seen as the temperature in correspondence of which the mean occupants' vote is zero. Two direct approaches have been proposed: a polynomial regression of the ASH score against (T_{out}, T_{in}) and a multinomial logistic regression that provides all the seven probabilities of the ASH scores as functions of (T_{out} and T_{in}). In both cases the obtained model implicitly defines the function $T_N(T_{out})$.

The third approach relies on an alternative definition of the neutral temperature which, rather than being based on the mean vote, maximizes the percentage of satisfied occupants. Two methods have been proposed to obtain this probabilistic neutral temperature: the same multinomial logistic regression already used for the direct approach and a simpler binomial logistic model.

A first finding of our work is the confirmation of the adaptive effect of the outdoor temperature. This is somewhat remarkable because the curves were obtained from the individual raw data, avoiding the intermediate steps at building level that are typically used to separate the useful information from the background noise. The adaptive effect is of particular interest for feedback control and efficient energy management, because higher indoor temperatures might be used as set points without increasing discomfort but with significant energy savings.

Concerning the differences between the neutral temperature models, they are almost always less than 1 and often much smaller. In particular, the models of the exploratory analysis (black, blue and red) are in very good agreement between each other. A similar consistence is observed between the two direct models (dash dot black and magenta). It is of interest that

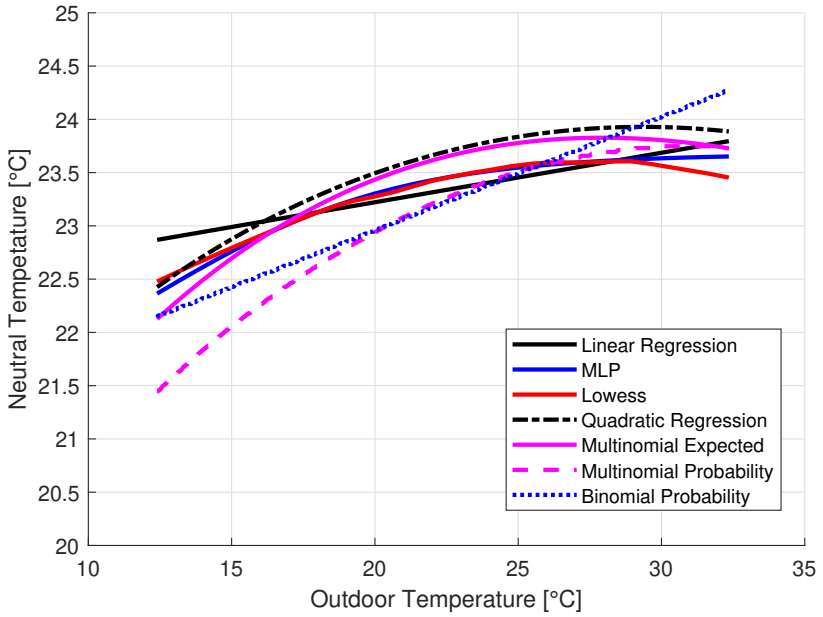

FIGURE 6.9: Comparison of neutral temperature models.

the introduction of a new probabilistic definition of neutral temperature (dashed magenta and dotted blue) leads to smaller predicted values for outdoor temperatures below $25°C$.

As far as the possible developments are concerned, three research directions can be envisioned. More flexible modeling approaches such as nonparametric and machine learning ones could be used to further validate the models presented herein. A second development regards the estimation of adaptive models also on winter data. Third, while the whole ASHRAE RP-884 dataset was included in the analysis, it would be of great interest to perform separate analyses for datasets coming from different countries and or climates in order to establish the possible existence of country or climate effects that should be accounted for.

7

Conclusions

Heating Ventilation and Air Conditioning is a field that is both critical in view of environmental priorities and extremely promising for what regards the potential improvements associated with the Industry 4.0 paradigm. In this thesis, the issue of energy efficient management of HVAC problem has been investigated along two main directions. On the production side, the continuous collection of monitoring data enables the identification of accurate and robust chiller models that can be used for the on line optimization of chillers' operation. On the demand side, the appropriate modelling of users' surveys can yield advanced thermal comfort models, usable to minimize energy consumption while guaranteeing the desired comfort.

Concerning chiller efficiency models, the availability of extensive monitoring data opens the way to new modeling approaches, e.g machine learning ones, but also poses new challenges as the distribution of operating conditions may be very disuniform, which may adversely impact on the generalization capabilities of the models. The main finding is that some traditional models are outperformed by machine learning ones, i.e. Artificial Neural Networks and Gaussian Processes, but also that the classical Gordon-Ng semi-empirical model can prove still useful, in terms of robustness and interpretability of the parameters.

Efficiency and power consumption models find application in the design of optimized management strategies for multiple chiller systems, a problem that due to its NP-hardness is typically solved via heuristic methods. In the thesis, assuming a quadratic power consumption model, a new exact solver, X-OCL, is proposed that, hinging on a partition of the solution space, reduces the problem to the closed form solution of elementary equality-constrained quadratic problems. In the typical industrial applications, the new exact approach appears viable under both the computational aspects and the adequacy of the quadratic assumption. The applicability of X-OCL was demonstrated on both literature benchmarks and an experimental benchmark based

on a extensive 2-year dataset collected in a semiconductor plant. The comparison with historical energy consumption data shows that the optimized management would achieve about 9% in savings.

On the demand size, the main issue was the validation of the so-called adaptive model and the derivation of advanced thermal comfort models. Concerning the former issue, different model have been trained and validated, confirming that the occupants' neutral temperature is indeed affected not only by the indoor temperature, but also from the outdoor one. Moreover, a new and more rigorous definition of neutral temperature has been proposed, based not on the average of the ASHRAE scores, but on the percentage of occupants that express a prescribed degree of satisfaction. The new definition relies on a novel statistical characterization of comfort scores based on multilogistic regression.

In conclusion, the thesis offers an overview of some relevant opportunities that arise from the introduction of innovative learning and optimization approaches in the field of HVAC systems, suggesting that both the production and demand side could strategically exploit the associated benefits.

Appendix A

Derivation of power consumption models from the COP ones

Assuming that the available COP-PLR curves represent the true model of chillers' efficiency, quadratic approximated power consumption models can be easily derived as follows:

1. Sample N data points from the COP-PLR curve of the $i-$th chiller, uniformly in the range $[PLR_{min,i},1]$ obtaining the pairs $\{PLR_i(k), COP_i(k)\},\ k = 1, \dots N$;
2. Compute the chiller's power consumptions as

$$P_i(k) = \frac{PLR_i(k) \cdot Q_{100\%,i}}{COP_i(k)},\ k = 1, \dots N;$$

3. Using the training set made of input-output paired samples $\{PLR_i(k), P_i(k)\},\ k = 1, \dots N$ estimate the parameters of the quadratic model [] via Ordinary Least Squares (OLS).

The above procedure was applied to the case studies presented in Section 5.5.1. For both the case studies, 50 data points were sampled from the COP-PLR curves reported in Tables A.1 to obtain the training datasets represented in the left panels of Figs. A.1 and A.2 as red dots. The estimated P-kW curves are shown in the right panels of Figs. A.1 and A.2 and their parameters are reported in Table A.2.

Systems	Chiller	α_i	β_i	γ_i	Q_{nom}
Case 1	1	0.1561	3.9023	-2.5909	450
	2	0.9000	1.8432	-1.4188	450
	3	0.2932	3.0419	-2.0054	1000
	4	0.1415	3.6376	-2.2469	1000
Case 2	1	1.5652	1.8094	-0.9803	1250
	2	1.0519	4.1471	-2.4173	1250
	3	0.5703	3.1602	-2.0912	1250
	4	0.3257	2.3513	-1.4265	1250
	5	0.5438	1.8668	-1.2361	1250
	6	1.5271	1.0634	-0.7238	1250
	7	0.7865	1.8473	-1.1633	1250
	8	0.8499	3.7768	-2.2859	1250
	9	1.1191	1.0228	-0.7542	1250

TABLE A.1: Coefficients of chillers' COP-PLR curves

Systems	Chiller	a_i	c_i	q_i	Q_{nom}
Case 1	1	243.58	-398.01	504.00	450
	2	130.81	-103.53	309.65	450
	3	417.51	-444.57	771.99	1000
	4	383.79	-347.84	611.64	1000
Case 2	1	95.54	321.92	103.60	1250
	2	170.68	41.43	235.46	1250
	3	371.09	-307.99	693.76	1250
	4	477.85	-217.09	733.35	1250
	5	433.17	-186-20	810.89	1250
	6	104.21	358.47	205.50	1250
	7	272.33	116.22	457.96	1250
	8	218.68	-20-94	333.72	1250
	9	191.69	276.56	429.51	1250

TABLE A.2: Estimated parameters of chillers' P-PLR curves

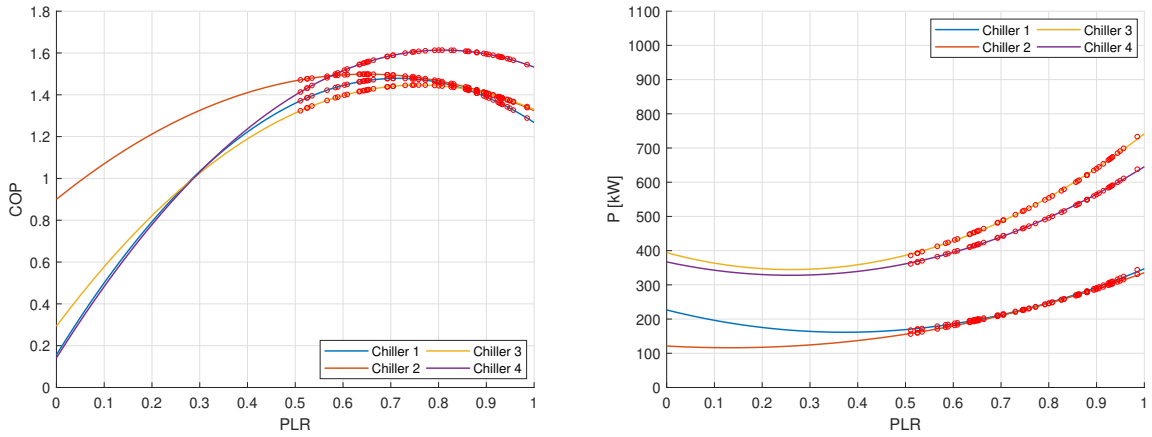

FIGURE A.1: Left: COP-PLR curves of chillers in case study n° 1. Red dots are the 50 samples used to identify the quadratic approximate power consumption model. Right: Identified quadratic P-PLR curves of chillers in case study n° 1. Red dots are the energy consumption data used for the training.

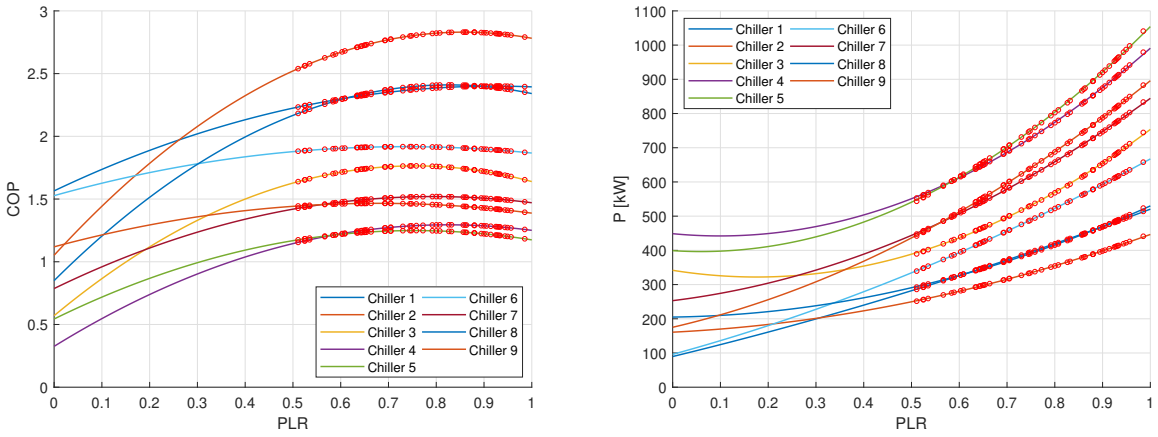

FIGURE A.2: Left: COP-PLR curves of chillers in case study n° 2. Red dots are the efficiency sampled data. Right: P-PLR curves of chillers in case study n° 2. Red dots are the energy consumption data used for the training.

Bibliography

[1] International Energy Agency (IEA), "The future of cooling opportunities for energy-efficient air conditioning," 2018.

[2] X. Guan, Z. Xu, and Q.-S. Jia, "Energy-efficient buildings facilitated by microgrid," *IEEE Transactions on smart grid*, vol. 1, no. 3, pp. 243–252, 2010.

[3] J. F. Nicol, M. Humphreys, *et al.*, "Understanding the adaptive approach to thermal comfort," *ASHRAE transactions*, vol. 104, pp. 991–1004, 1998.

[4] F. Nicol, M. Humphreys, and S. Roaf, *Adaptive thermal comfort: principles and practice*. Routledge, 2012.

[5] R. De Dear and G. S. Brager, "The adaptive model of thermal comfort and energy conservation in the built environment," *International journal of biometeorology*, vol. 45, no. 2, pp. 100–108, 2001.

[6] E. Halawa and J. Van Hoof, "The adaptive approach to thermal comfort: A critical overview," *energy and buildings*, vol. 51, pp. 101–110, 2012.

[7] J. F. Nicol and M. A. Humphreys, "Adaptive thermal comfort and sustainable thermal standards for buildings," *Energy and buildings*, vol. 34, no. 6, pp. 563–572, 2002.

[8] R. Yao, B. Li, and J. Liu, "A theoretical adaptive model of thermal comfort–adaptive predicted mean vote (apmv)," *Building and environment*, vol. 44, no. 10, pp. 2089–2096, 2009.

[9] J. Gordon and K. C. Ng, "Thermodynamic modeling of reciprocating chillers," *Journal of Applied Physics*, vol. 75, no. 6, pp. 2769–2774, 1994.

[10] T. E. Mull, *HVAC principles and applications manual*. McGraw-Hill New York, 1998.

[11] N. Nassif, S. Kajl, and R. Sabourin, "Optimization of hvac control system strategy using two-objective genetic algorithm," *HVAC&R Research*, vol. 11, no. 3, pp. 459–486, 2005.

[12] S. Wang and Z. Ma, "Supervisory and optimal control of building hvac systems: A review," *HVAC&R Research*, vol. 14, no. 1, pp. 3–32, 2008.

[13] A. Beghi, L. Cecchinato, and M. Rampazzo, "A multi-phase genetic algorithm for the efficient management of multi-chiller systems," *Energy Conversion and Management*, vol. 52, no. 3, pp. 1650–1661, 2011.

[14] A. Beghi, L. Cecchinato, G. Cosi, and M. Rampazzo, "A pso-based algorithm for optimal multiple chiller systems operation," *Applied Thermal Engineering*, vol. 32, pp. 31–40, 2012.

[15] M. Comstock, J. Braun, and E. Groll, "The sensitivity of chiller performance to common faults," *HVACR Research*, vol. 7, no. 3, pp. 263–279, 2001.

[16] P. Sreedhara and P. Haves, "Comparison of chiller models for use in model-based fault detection," 2001.

[17] S. M. Namburu, M. S. Azam, J. Luo, K. Choi, and K. R. Pattipati, "Data-driven modeling, fault diagnosis and optimal sensor selection for hvac chillers," *IEEE transactions on automation science and engineering*, vol. 4, no. 3, pp. 469–473, 2007.

[18] A. Beghi, L. Cecchinato, F. Peterle, M. Rampazzo, and F. Simmini, "Model-based fault detection and diagnosis for centrifugal chillers," in *Control and Fault-Tolerant Systems (SysTol), 2016 3rd Conference on*, pp. 158–163, IEEE, 2016.

[19] J. Gordon and K. Ng, "Cool thermodynamics, cambridge int," *Science, Cambridge*, 2000.

[20] T.-S. Lee, K.-Y. Liao, and W.-C. Lu, "Evaluation of the suitability of empirically-based models for predicting energy performance of centrifugal water chillers with variable chilled water flow," *Applied energy*, vol. 93, pp. 583–595, 2012.

[21] F. W. Yik and V. K. Lam, "Chiller models for plant design studies," *Building Services Engineering Research and Technology*, vol. 19, no. 4, pp. 233–241, 1998.

[22] H. Horak, *DOE-2 reference manual*. 1979.

[23] D. J. Swider, "A comparison of empirically based steady-state models for vapor-compression liquid chillers," *Applied thermal engineering*, vol. 23, no. 5, pp. 539–556, 2003.

[24] M. Hydeman and K. L. Gillespie Jr, "Tools and techniques to calibrate electric chiller component models/discussion," *ASHRAE transactions*, vol. 108, p. 733, 2002.

[25] A. Standard, "550/590—1998 water chilling packages using the vapor compression cycle," *Air Conditioning & Refrigeration Institute*, 2002.

[26] M. Sugiyama, M. Krauledat, and K.-R. MÃžller, "Covariate shift adaptation by importance weighted cross validation," *Journal of Machine Learning Research*, vol. 8, no. May, pp. 985–1005, 2007.

[27] W. Jiang and T. A. Reddy, "Reevaluation of the gordon-ng performance models for water-cooled chillers," *ASHRAE transactions*, vol. 109, p. 272, 2003.

[28] D. J. MacKay, "Bayesian interpolation," *Neural computation*, vol. 4, no. 3, pp. 415–447, 1992.

[29] C. E. Rasmussen, "Gaussian processes in machine learning," in *Advanced lectures on machine learning*, pp. 63–71, Springer, 2004.

[30] T. Reddy, D. Niebur, J. Gordon, J. Seem, K. Andersen, G. Cabrera, Y. Jia, and P. Pericolo, "Final report: Development and comparison of on-line model training techniques for model-based fdd methods applied to vapor compression chillers," *ASHRAE Research Project*, 2001.

[31] J. Chen and R. J. Patton, *Robust model-based fault diagnosis for dynamic systems*, vol. 3. Springer Science & Business Media, 2012.

[32] K. K. Andersen and T. A. Reddy, "The error in variables (eiv) regression approach as a means of identifying unbiased physical parameter estimates: Application to chiller performance data," *HVAC&R Research*, vol. 8, no. 3, pp. 295–309, 2002.

[33] K. Ng, H. Chua, W. Ong, S. Lee, and J. Gordon, "Diagnostics and optimization of reciprocating chillers: theory and experiment," *Applied thermal engineering*, vol. 17, no. 3, pp. 263–276, 1997.

[34] M. Comstock and J. E. Braun, "Fault detection and diagnostic (fdd) requirements and evaluation tools for chillers," *West Lafayette, IN: ASHRAE*, 2002.

[35] G. Casella and R. L. Berger, *Statistical inference*, vol. 2. Duxbury Pacific Grove, CA, 2002.

[36] S. Chatterjee and A. S. Hadi, *Regression analysis by example*. John Wiley & Sons, 2015.

[37] N. R. Draper and H. Smith, *Applied regression analysis*, vol. 326. John Wiley & Sons, 2014.

[38] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*. CRC press, 1994.

[39] A. C. Davison and D. V. Hinkley, *Bootstrap methods and their application*, vol. 1. Cambridge university press, 1997.

[40] Y.-C. Chang, "A novel energy conservation method—optimal chiller loading," *Electric Power Systems Research*, vol. 69, no. 2-3, pp. 221–226, 2004.

[41] Y.-C. Chang, T.-S. Chan, and W.-S. Lee, "Economic dispatch of chiller plant by gradient method for saving energy," *Applied Energy*, vol. 87, no. 4, pp. 1096–1101, 2010.

[42] Z. W. Geem, "Solution quality improvement in chiller loading optimization," *Applied Thermal Engineering*, vol. 31, no. 10, pp. 1848–1851, 2011.

[43] E. Salari and A. Askarzadeh, "A new solution for loading optimization of multi-chiller systems by general algebraic modeling system," *Applied Thermal Engineering*, vol. 84, pp. 429–436, 2015.

[44] Y.-C. Chang, "Genetic algorithm based optimal chiller loading for energy conservation," *Applied Thermal Engineering*, vol. 25, no. 17-18, pp. 2800–2815, 2005.

[45] Y.-C. Chang, F.-A. Lin, and C. H. Lin, "Optimal chiller sequencing by branch and bound method for saving energy," *Energy conversion and management*, vol. 46, no. 13-14, pp. 2158–2172, 2005.

[46] Y.-C. Chang, W.-H. Chen, C.-Y. Lee, and C.-N. Huang, "Simulated annealing based optimal chiller loading for saving energy," *Energy Conversion and Management*, vol. 47, no. 15-16, pp. 2044–2058, 2006.

[47] Y.-C. Chang, "An innovative approach for demand side management—optimal chiller loading by simulated annealing," *Energy*, vol. 31, no. 12, pp. 1883–1896, 2006.

[48] A. J. Ardakani, F. F. Ardakani, and S. Hosseinian, "A novel approach for optimal chiller loading using particle swarm optimization," *Energy and Buildings*, vol. 40, no. 12, pp. 2177–2187, 2008.

[49] W.-S. Lee and L.-C. Lin, "Optimal chiller loading by particle swarm algorithm for reducing energy consumption," *Applied Thermal Engineering*, vol. 29, no. 8-9, pp. 1730–1734, 2009.

[50] Y.-C. Chang, C.-Y. Lee, C.-R. Chen, C.-J. Chou, W.-H. Chen, and W.-H. Chen, "Evolution strategy based optimal chiller loading for saving energy," *Energy Conversion and Management*, vol. 50, no. 1, pp. 132–139, 2009.

[51] W.-S. Lee, Y.-T. Chen, and Y. Kao, "Optimal chiller loading by differential evolution algorithm for reducing energy consumption," *Energy and Buildings*, vol. 43, no. 2-3, pp. 599–604, 2011.

[52] L. dos Santos Coelho, C. E. Klein, S. L. Sabat, and V. C. Mariani, "Optimal chiller loading for energy conservation using a new differential cuckoo search approach," *Energy*, vol. 75, pp. 237–243, 2014.

[53] M. H. Sulaiman, H. Ibrahim, H. Daniyal, and M. R. Mohamed, "A new swarm intelligence approach for optimal chiller loading for energy conservation," *Procedia-Social and Behavioral Sciences*, vol. 129, pp. 483–488, 2014.

[54] L. dos Santos Coelho and V. C. Mariani, "Improved firefly algorithm approach applied to chiller loading for energy conservation," *Energy and Buildings*, vol. 59, pp. 273–278, 2013.

[55] P.-y. Duan, J.-q. Li, Y. Wang, H.-y. Sang, and B.-x. Jia, "Solving chiller loading optimization problems using an improved teaching-learning-based optimization algorithm," *Optimal Control Applications and Methods*, vol. 39, no. 1, pp. 65–77, 2018.

[56] Z.-x. Zheng and J.-q. Li, "Optimal chiller loading by improved invasive weed optimization algorithm for reducing energy consumption," *Energy and Buildings*, vol. 161, pp. 80–88, 2018.

[57] Y.-C. Chang, "An outstanding method for saving energy-optimal chiller operation," *IEEE transactions on energy conversion*, vol. 21, no. 2, pp. 527–532, 2006.

[58] S. Huang, W. Zuo, and M. D. Sohn, "Amelioration of the cooling load based chiller sequencing control," *Applied Energy*, vol. 168, pp. 204–215, 2016.

[59] EIA, "Commercial buildings energy consumption survey (cbecs)," tech. rep., Energy Information Administration, 2016.

[60] A. Brambilla, G. Salvalai, M. Imperadori, and M. M. Sesana, "Nearly zero energy building renovation: From energy efficiency to environmental efficiency, a pilot case study," *Energy and Buildings*, vol. 166, pp. 271 – 283, 2018.

[61] F. Birol, "Towards a zero-emission, efficient, and resilient buildings and construction sector - global status report," tech. rep., UN environment, 2017.

[62] L. Yang, H. Yan, and J. C. Lam, "Thermal comfort and building energy consumption implications – a review," *Applied Energy*, vol. 115, pp. 164 – 173, 2014.

[63] P. O. Fanger *et al.*, "Thermal comfort. analysis and applications in environmental engineering.," *Thermal comfort. Analysis and applications in environmental engineering.*, 1970.

[64] A. Auliciems, "Towards a psycho-physiological model of thermal perception," *International Journal of Biometeorology*, vol. 25, no. 2, pp. 109–122, 1981.

[65] R. de Dear and M. Fountain, "Field experiments on occupant comfort and office thermal environments in a hot-humid climate," 1994.

[66] M. Humphreys, "Outdoor temperatures and comfort indoors," *Batiment International, Building Research and Practice*, vol. 6, no. 2, pp. 92–92, 1978.

[67] M. A. Humphreys, "Field studies of thermal comfort compared and applied, department of the environment: Building research establishment," *Current paper CP*, vol. 76, no. 75, 1975.

[68] A. Auliciems and R. de Dear, "Airconditioning in australia i—human thermal factors," *Architectural Science Review*, vol. 29, no. 3, pp. 67–75, 1986.

[69] R. De Dear and G. S. Brager, "Developing an adaptive model of thermal comfort and preference," 1998.

[70] R. J. De Dear, "A global database of thermal comfort field experiments," *ASHRAE transactions*, vol. 104, p. 1141, 1998.

[71] M. A. Humphreys and J. F. Nicol, "Outdoor temperature and indoor thermal comfort: Raising the precision of the relationship for the 1998 ashrae database of field studies/discussion," *Ashrae Transactions*, vol. 106, p. 485, 2000.

[72] E. Ballantyne, R. Hill, and J. Spencer, "Probit analysis of thermal sensation assessments," *International Journal of Biometeorology*, vol. 21, no. 1, pp. 29–43, 1977.

[73] R. Williams *et al.*, "Generalized ordered logit/partial proportional odds models for ordinal dependent variables," *Stata Journal*, vol. 6, no. 1, p. 58, 2006.