

1 Article

# 2 Accelerating the K-Nearest Neighbors Filtering 3 Algorithm to Optimize the Real-Time Classification 4 of Human Brain Tumor in Hyperspectral Images

5 Giordana Florimbi <sup>1,\*</sup>, Himar Fabelo <sup>2</sup>, Emanuele Torti <sup>1</sup>, Raquel Lazcano <sup>3</sup>, Daniel Madroñal <sup>3</sup>,  
6 Samuel Ortega <sup>2</sup>, Ruben Salvador <sup>3</sup>, Francesco Leporati <sup>1</sup>, Giovanni Danese <sup>1</sup>, Abelardo  
7 Báez-Quevedo <sup>2</sup>, Gustavo M. Callicó <sup>2</sup>, Eduardo Juárez <sup>3</sup>, César Sanz <sup>3</sup>, Roberto Sarmiento <sup>2</sup>

8 <sup>1</sup> Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Pavia 27100, Italy;  
9 giordana.florimbi01@universitadipavia.it (G.F.); emanuele.torti@unipv.it (E.T.); leporati@unipv.it (F.L.);  
10 gianni.danese@unipv.it (G.D.)

11 <sup>2</sup> Institute for Applied Microelectronics (IUMA), University of Las Palmas de Gran Canaria (ULPGC), Las  
12 Palmas de Gran Canaria 35017, Spain; hfabelo@iuma.ulpgc.es (H.F.); sortega@iuma.ulpgc.es (S.O.);  
13 abaez@iuma.ulpgc.es (A.B.-Q.); gustavo@iuma.ulpgc.es (G.M.C.); roberto@iuma.ulpgc.es (R.Sar.);

14 <sup>3</sup> Centre of Software Technologies and Multimedia Systems (CITSEM), Technical University of Madrid  
15 (UPM), Madrid 28031, Spain; raquel.lazcano@upm.es (R.L.); daniel.madronal@upm.es (D.M.);  
16 ruben.salvador@upm.es (R.Sal.); ejuarez@sec.upm.es (E.J.); cesar.sanz@upm.es (C.S.)

17 \* Correspondence: giordana.florimbi01@universitadipavia.it; Tel.: +39 0382 985678

18 Academic Editor: name

19 Received: date; Accepted: date; Published: date

20 **Abstract:** The use of hyperspectral imaging (HSI) in the medical field is an emerging approach to  
21 assist physicians in diagnostic or surgical guidance tasks. However, HSI data processing involves  
22 very high computational requirements due to the huge amount of information captured by the  
23 sensor. One of the stages with higher computational load is the K-Nearest Neighbors (KNN)  
24 filtering algorithm. The main goal of this study is to optimize and parallelize the (KNN) algorithm  
25 by exploiting the GPU technology to obtain real-time processing during brain cancer surgical  
26 procedures. This parallel version of the KNN performs the neighboring filtering of a classification  
27 map (obtained from a supervised classifier), evaluating the different classes simultaneously. The  
28 undertaken optimizations and the computational capabilities of the GPU device throw a speedup  
29 of up to 66.18x when compared to a sequential implementation.

30 **Keywords:** K-nearest neighbors filtering; hyperspectral imaging instrumentation; brain cancer  
31 detection; image processing; graphics processing units.

---

## 33 1. Introduction

34 Hyperspectral imaging (HSI) is a non-contact, non-ionizing and non-invasive imaging  
35 technique suitable for medical applications [1,2]. HSI combines traditional imaging and  
36 spectroscopy methods to obtain both spatial and spectral information of the captured scene [3].  
37 Currently, HSI is becoming more popular in surgery applications as a guidance tool for surgeons,  
38 since it can provide more information than traditional imaging techniques, such as Magnetic  
39 Resonance (MR), Computed Tomography (CT), Ultrasound (US) and Positron Emission  
40 Tomography (PET), being a non-invasive and risk-free technique [4–7]. However, one of the main  
41 handicaps of this technology is the high computational requirements needed to process the large  
42 amount of data acquired by the sensor. The use of High Performance Computing (HPC) and highly  
43 parallelized algorithms is mandatory to obtain HSI intra-operative real-time processing [8].

44 The work presented in this paper is focused on the optimization, parallelization and  
45 implementation onto a Graphics Processing Unit (GPU) of the K-Nearest Neighbors (KNN) filtering  
46 algorithm to obtain real-time performance. This work has been done taking into account the results  
47 and intermediate data obtained during the deployment of the HELICoiD (HypErspectraL Imaging  
48 Cancer Detection) European FET project [9–12]. The goal of this project was to apply the HSI  
49 technique to discriminate between tumor and normal brain tissue during neurosurgical operations  
50 in real-time. Employing the developed system, surgeons will have a guidance tool to assist them  
51 during the brain tumor resection, avoiding unintentionally leaving behind small remnants of tumors  
52 and the excessive extraction of normal tissue. This accurate delimitation of the tumor boundaries  
53 will improve surgery outcomes, therefore improving the patient's quality of life.

54 The KNN algorithm is a classifier widely used in several research areas and also in the field of  
55 HSI, where a pixel-wise classification is performed [13]. The most relevant part of this method is the  
56 K-Nearest Neighbors searching, which is a highly demanding task in terms of computational time.  
57 Since the main goal of the majority of the works in the state of the art is to execute this algorithm in  
58 real-time or at least with reduced execution times, it becomes necessary to exploit high performance  
59 devices, being GPUs a highly appealing option. As a massively parallel architecture, this kind of  
60 devices has been widely used for exploiting data parallelism in several applications from different  
61 scientific fields [14–16] and also in HSI [17,18].

62 Concerning the KNN algorithm, several parallel CUDA implementations have been proposed  
63 in the scientific literature. The results reported by these previous works point out that this  
64 technology is able to provide very high speedups compared to serial codes. For example, in [19],  
65 authors provide two CUDA versions of this algorithm, one characterized by custom kernels and the  
66 other exploiting the CUBLAS library [20]. Using synthetic data, they obtain speedups equal to 64x  
67 and 189x, respectively, compared to the highly optimized ANN (Approximate Nearest Neighbors)  
68 C++ library. They also apply the two parallel versions of the KNN algorithm to the high dimensional  
69 SIFT (Scale-Invariant Feature Transform) feature matching, obtaining speedups of 25x and 62x,  
70 respectively. In addition, in [21], the authors exploit the GPU technology to implement a new  
71 version of the KNN algorithm called *Sweet KNN*. This new algorithm is based on a Triangle  
72 Inequality (TI) approach, which tries to reduce the number of distance computations, since the goal  
73 of the work is to find a good balance between redundancy minimization and regularity preservation  
74 for various datasets. This work presents a speedup with an average value equal to 11x compared to  
75 the existing GPU implementations on KNN, with a maximum speedup of 120x.

76 Recent uses of the KNN algorithm show that it is not restricted to only a classification role. In  
77 the last few years it has also been used as a filtering technique, able to improve the results of spectral  
78 classifications by adding spatial domain information [22]. In this work it is presented a parallel  
79 implementation of the KNN filtering algorithm, which can integrate the output of the SVM (Support  
80 Vector Machines) classifier with the one-band representation of a hyperspectral (HS) cube generated  
81 using the PCA (Principal Component Analysis) algorithm. The goal of the entire system is to  
82 perform the real-time classification of brain cancer, where real-time restrictions for this application  
83 are considered to be met when the processing time is lower than the time elapsed between the  
84 consecutive acquisitions of two images of the exposed brain (~1 minute).

## 85 2. Materials and Methods

86 This section presents the HSI instrumentation employed to obtain the in-vivo HS brain cancer  
87 image database, the serial implementation of the algorithm as well as the optimizations and the  
88 parallelization analysis applied to the KNN filtering algorithm.

### 89 2.1. Hyperspectral Acquisition System

90 To obtain the HS in-vivo brain image dataset used in this study, a customized HS acquisition  
91 system was employed [12]. The acquisition system is composed by a VNIR (Visual and Near  
92 Infra-Red) *pushbroom* camera (Hyperspec® VNIR A-Series, Headwall Photonics Inc., Fitchburg, MA,  
93 USA). This camera is able to capture images within the spectral range between 400 and 1000 nm,

94 obtaining 826 spectral bands with a spectral resolution of 2-3 nm and a pixel dispersion of 0.74 nm.  
95 In order to capture the complete HS cube, the camera uses a pushbroom scanning technique, which  
96 allows the 2-D detector to capture the complete spectral dimensions and one spatial dimension of  
97 the scene. By shifting the camera's field of view relative to the scene, the second spatial dimension is  
98 acquired. Figure 1 shows the HS acquisition system while capturing an image of the exposed brain  
99 surface during a surgical operation at the University Hospital Doctor Negrin of Las Palmas de Gran  
100 Canaria (Spain).



101 **Figure 1.** Hyperspectral acquisition system capturing an image during a neurosurgical  
102 operation.

## 103 2.2. Hyperspectral brain cancer image database


104 In this study, a set of five in-vivo brain surface HS images were employed to evaluate the  
105 performance of the KNN filtering implementation. These images were captured using the HS  
106 acquisition system, belonging to adult patients undergoing craniotomy for resection of intra-axial  
107 brain tumor. Images were obtained at the University Hospital Doctor Negrin of Las Palmas de Gran  
108 Canaria (Spain) from four different patients with a confirmed grade IV glioblastoma tumor by  
109 histopathology. The study protocol and consent procedures were approved by the *Comité Ético de*  
110 *Investigación Clínica-Comité de Ética en la Investigación* (CEIC/CEI) of University Hospital Doctor  
111 Negrin and written informed consent was obtained from all subjects.

112 HS images were obtained intra-operatively after craniotomy and resection of the dura. Before  
113 the images were captured, the operating surgeon initially identified the approximate location of  
114 normal brain and tumor. Rubber ring markers were then placed on these locations and the images  
115 were taken with markers in situ. At that point, tissue samples were resected from the marked areas  
116 and sent to pathology for tissue diagnosis. Depending on the location of the tumor, images were  
117 acquired at various stages of the operation. In those cases with superficial tumors, some images were  
118 obtained immediately after the dura was removed, while in the cases with deep laying tumors,  
119 images were obtained during the actual tumor resection.

120 The selected HS images were pre-processed following the pre-processing chain presented in  
121 [12]. The pre-processing chain is composed by four steps: image calibration, noise filtering, band  
122 averaging and pixel normalization. In the first step, after the image acquisition, the HS raw data are  
123 calibrated using a white reference image (captured from a certified white reference tile in the same  
124 illumination conditions in which the images were captured) and a dark reference image (obtained by  
125 keeping the camera shutter closed). This calibration is performed to avoid the problem of the  
126 spectral non-uniformity of the illumination device and the dark currents of the camera sensor. Next,  
127 due to the high spectral noise generated by the camera sensor, a set of steps are applied with the goal  
128 of removing this noise from the spectral signatures and to reduce the number of bands of the  
129 samples without losing the main spectral information. Finally, a normalization step is performed in  
130 order to homogenize the spectral signatures in terms of the reflectance level. The final HS cube is  
131 formed by 128 spectral bands, covering the range between 450 and 900 nm [12].

132 Figure 2a shows an example of the synthetic RGB representation of an HS cube from the HS  
 133 brain cancer image database used in this study. Furthermore, each one of these HS images was  
 134 processed employing a supervised SVM classifier (Figure 2b), and a one-band representation was  
 135 obtained using a PCA algorithm (Figure 2c). Table 1 details the characteristics of each HS image,  
 136 where PXC<sub>Y</sub> stands for Patient X and Capture Y.

137 Four different classes were labeled in the images for the supervised classification: tumor tissue,  
 138 normal tissue, hypervascularized tissue (mainly blood vessels) and background (other materials or  
 139 substances that can be presented in the surgical scene that are not relevant for the tumor resection  
 140 process). These classes were represented in the classification maps with the following colors: red,  
 141 green, blue and black, respectively.



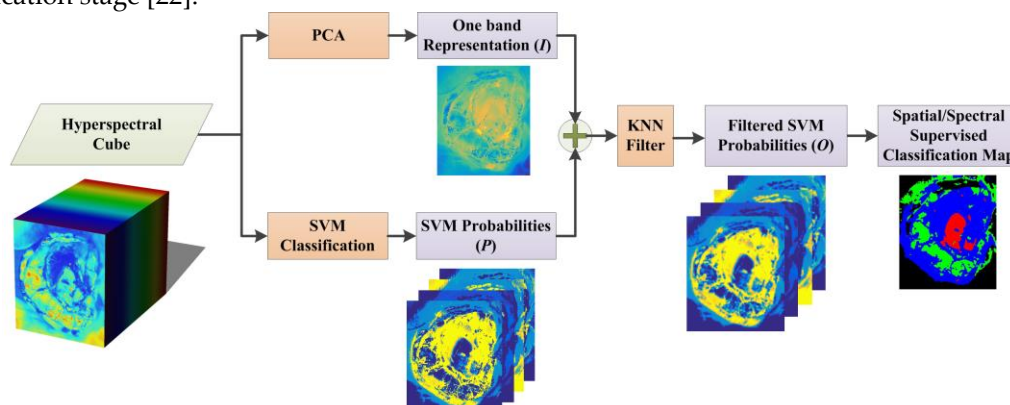
142 **Figure 2.** Example of an in-vivo HS human brain image dataset employed in the study (P2C1).  
 143 (a) Synthetic RGB representation of the HS cube; (b) Supervised classification map obtained  
 144 using the SVM classifier; (c) One-band representation of the HS cube obtained employing  
 145 PCA algorithm.

146 **Table 1.** HS brain cancer image database.

Image ID	#Pixels	Dimensions (width x height x bands)
P1C1	251,532	548 x 459 x 128
P1C2	264,408	552 x 479 x 128
P2C1	219,232	496 x 442 x 128
P3C1	185,368	493 x 376 x 128
P4C1	124,691	329 x 379 x 128

### 147 2.3. K-Nearest Neighbors filtering algorithm

148 As introduced in Section 1, the KNN algorithm has recently been widely used in the field of HSI  
 149 as a filtering technique [22] to refine outputs from classifiers, in this case a SVM, with the spectral  
 150 information computed through the PCA algorithm. As shown in Figure 3, the KNN-based filtering  
 151 algorithm [22][23] receives an input image  $P$ , which is composed of the probability maps estimated  
 152 by the SVM classifier, and a guidance image  $I$ , that is the one-band representation of the HS cube,  
 153 generated using a dimensional reduction algorithm, such as PCA. The output of this algorithm is a  
 154 filtered classification map based on the highest probability assigned to each pixel in the previous  
 155 classification stage [22].



156 **Figure 3.** Block diagram of the KNN based spatial-spectral classification.

157 In this method, the nearest neighbors of a certain pixel are searched in a feature space, which  
 158 contains both the pixel value and the spatial coordinates. This space is defined by a feature vector  $F$   
 159 as shown in Equation (1):

$$F(q) = (I(q), \lambda \cdot l(q), \lambda \cdot h(q)) \quad (1)$$

160 where  $I(q)$  is the normalized pixel value of the guidance image and  $l(q)$  and  $h(q)$  refer to the  
 161 normalized coordinates of pixel  $q$ . The spectral value of the pixel and its spatial coordinates are  
 162 weighted with  $\lambda$ , which is a balance parameter to weigh the importance of the spatial information  
 163 in the searching of the neighbors. If  $\lambda$  is zero, the spatial information will not be considered. If its  
 164 value is higher than zero, more influence is given to the local neighborhood in the filtering process  
 165 [22][23].

166 The KNN searching requires the computation of the distances between pixels on the base of  
 167 the data contained in the feature vector. The distance from a given pixel located at  $(r, c)$  coordinates  
 168 of the image to any other pixel at  $(i, j)$  can be computed using the Euclidean distance, i.e. the  
 169 2-norm:

$$d(I(rc), I(ij)) = \sqrt{(I_{rc} - I_{ij})^2 + (r - i)^2 + (c - j)^2} \quad (2)$$

170 where  $I_{rc}$  is the normalized pixel value of the guidance image  $I$  at row  $r$  and column  $c$  and  $I_{ij}$  is  
 171 the value of every other pixel at row  $i$  and column  $j$ .

172 In this work, also the Manhattan metric (Equation (3)) has been used to compute the distances,  
 173 considering always all the data contained in the feature vector. In section 3, it will be presented a  
 174 comparison between the implementations performed using these two metrics, underlining both the  
 175 differences in terms of processing time and classification results.

$$d(I(rc), I(ij)) = |I_{rc} - I_{ij}| + |r - i| + |c - j| \quad (3)$$

176 Once the distances for each pixel are computed, the algorithm has to sort them to select the  $K$   
 177 nearest neighbors. After the KNN searching is concluded, the algorithm continues with the filtering  
 178 step whose output is the optimized probability  $O(q)$ . For each pixel, it computes a number of  
 179 outputs equal to the number of SVM classes. In particular, for each pixel  $q$  and each SVM class, it  
 180 computes the optimized probability  $O(q)$ , defined as follows:

$$O(q) = \frac{\sum P(s)}{K}, s \in \omega_q \quad (4)$$

181 where  $P$  is the original probability map (one per class) generated by the SVM classifier,  $\omega_q$   
 182 indicates the set of  $K$  nearest neighbors of the pixel  $q$  and  $s$  is the index related to each neighbor of  
 183 the previous set [23].

184 The last step of the algorithm consists of assigning a label to each pixel to generate a new final  
 185 classification map. The label that is assigned to each pixel of the image is the class with the highest  
 186 optimized probability.

#### 187 2.4. KNN filtering algorithm implementation

188 After describing all the steps of the KNN algorithm, this section will introduce the  
 189 optimizations and the parallelization analysis performed to the algorithm in order to reduce its  
 190 computational cost.

##### 191 2.4.1. KNN Search optimization

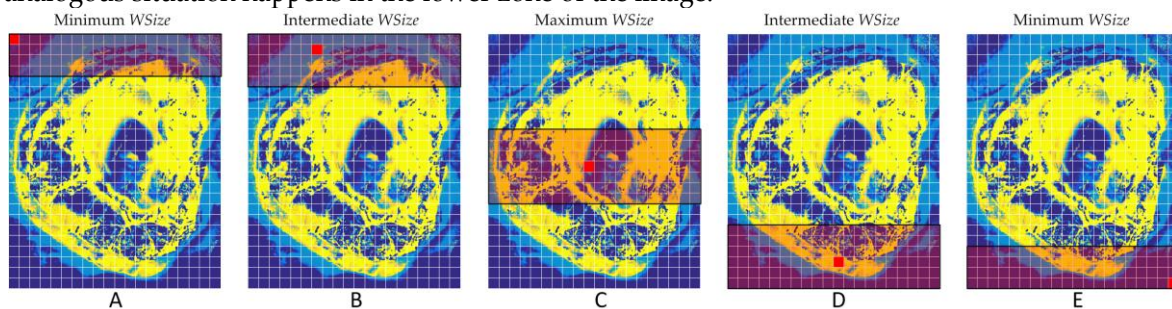
192 After an extensive analysis of the computational cost of the KNN filtering algorithm, it is  
 193 possible to determine that the neighbors searching phase is the most consuming part of the code. For  
 194 this reason, the first optimization proposed is the definition of a *search window* in the neighbors'  
 195 selection. This *search window* is a region close enough to the pixel whose neighbors are going to be  
 196 chosen. In the original algorithm, this step consisted on computing, for each pixel, a number of  
 197 distances equal to  $N_{pixels} - 1$ , where  $N_{pixels}$  is the number of pixels in the image. Our approach is to  
 198 search the  $K$  nearest neighbors of a pixel within this window, not considering the entire image, in

199 order to reduce the number of computed distances, as the probability to find smaller distances in  
 200 further zones of the image tend to 0.

201 Concerning the parameter setting, in [23] it is asserted that  $\lambda=1$  and  $K=40$  are a good  
 202 compromise for this medical imaging application. This value of  $\lambda$  gives a high importance to the  
 203 spatial information, in particular to the local neighborhood. Taking into account this value, and the  
 204 fact that the values of the guidance image  $I$  (i.e. the first term in Equation (2)) are normalized to 1, it  
 205 is easy to foresee the behavior of the sorting algorithm in the neighbors selection. This allows  
 206 introducing heuristic considerations that will help reducing the execution time. For any given pixel  
 207  $I_{rc}$  in the image, the distance computation will follow a pattern determined by the spatial distance,  
 208 i.e. the last two terms in Equation (2), which are related to the spatial coordinates of the pixel. These  
 209 two values will hence dominate the equation once they overcome the spectral value (i.e. the first  
 210 term in Equation (2)), since they will contribute to a distance value for any other pixel that will be  
 211 always predominant if such pixels are far enough from the location of the pixel under consideration.  
 212 In other words, it is a sufficient condition to sort only a certain subset of pixels in a region close  
 213 enough to the pixel whose neighbors are being searched. This effectively reduces the search space  
 214 and the computational cost.

215 Section 3 compares the computational time and classification results for both the serial and the  
 216 parallel implementations varying the window sizes ( $WSize$ ). After several analysis, a window with  
 217  $Wsize=14$ , i.e. 14 rows of the image, has been selected, so the search space contains a number of  
 218 pixels equal to 14 rows  $\times$  total number of columns. A row-wise window has been selected instead of a  
 219 column-wise one in order to have all the data stored sequentially. The version characterized by this  
 220 window size has been chosen as the *reference result* because the classification results are the same  
 221 compared to the implementation that considers the entire window.

222 The window is considered in a symmetric way relatively to the pixel that is being processed, so  
 223 one half of the window is evaluated over the pixel and the other half below it (Figure 4). In order to  
 224 avoid the effect of the borders, those pixels near them are treated separately. In this case, to  
 225 maintain a certain spatial coherence, the size of the window for the pixels in the top-most rows is  
 226 smaller at first, so as not to search further than  $Wsize/2$  down in the image. This way, the band  
 227 grows with each further pixel being processed until the steady state is reached. This happens when  
 228 the number of pixels above the one being processed reaches  $Wsize/2$  and it is kept until an  
 229 analogous situation happens in the lower zone of the image.



230 **Figure 4.** KNN window searching method example. (A) Minimum window size of the first  
 231 pixel; (B) Intermediate window size of a pixel near the upper border; (C) Maximum window  
 232 size of a pixel in the center of the image; (D) Intermediate window size of a pixel near the  
 233 bottom border; (E) Minimum window size of the last pixel.

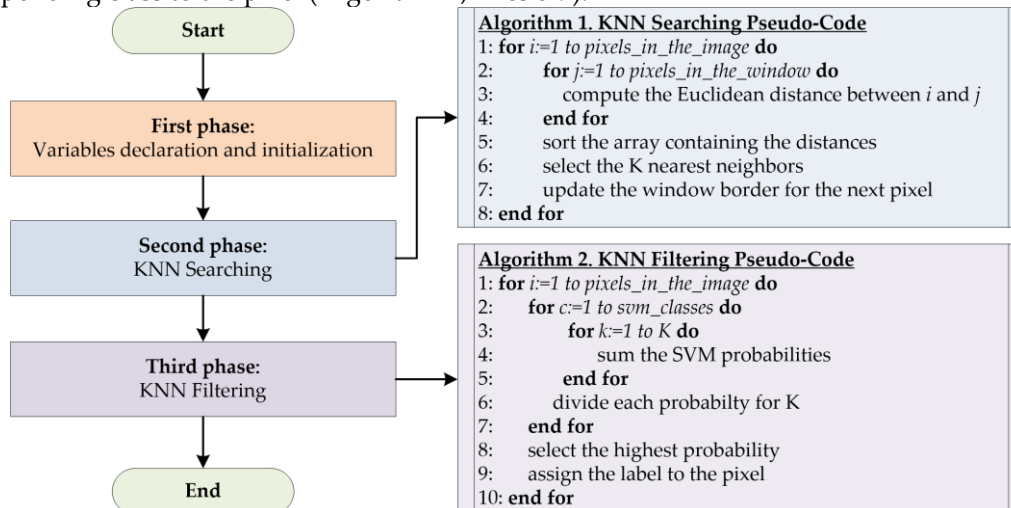
#### 234 2.4.2. Serial implementation

235 The serial implementation of the KNN algorithm is written in C language and presents three  
 236 main phases, as shown in Figure 5. The first one concerns the declaration and initialization of all the  
 237 variables, arrays and structures needed in the computation. For example, for each pixel, two types  
 238 of structures are defined: the former, *featureMatrixNode*, contains all the parameters needed in the  
 239 computation of the Euclidean distance, shown in Equation (2). The latter, *featureDistance*, contains

240 the distances (between the considered pixel and the pixels within its window) and the indexes of  
 241 these pixels.

242 The second phase refers to the  $K$  nearest neighbors searching. Considering every pixel in the  
 243 image (Algorithm 1, line 1), the algorithm computes the distances between it and all the pixels  
 244 inside its window (Algorithm 1, lines 2-4), exploiting the Euclidean metric in Equation (2). After  
 245 storing all the distances in the *featureDistance* structure, the algorithm sorts them in ascending order  
 246 through the Merge Sort algorithm (Algorithm 1, line 5) and selects the indexes of the  $K$  pixels  
 247 characterized by the lower distances (Algorithm 1, line 6). At the end of this phase, the parameters  
 248 related to the window sizes are updated (Algorithm 1, line 7) on the base of the pixel location, as  
 249 described in section 2.4.1.

250 Once all the neighbors of each pixel have been computed, the KNN filtering phase starts  
 251 (Figure 5, Algorithm 2). Its goal is to assign a label to each pixel considering the probability maps  
 252 generated by the SVM algorithm. In this phase, the algorithm computes, for each pixel, a number of  
 253 optimized probabilities  $O(q)$ , described in Equation (4), equal to the number of the SVM classes,  
 254 which are four in this work. In particular, for each class, the SVM probabilities of all the neighbors  
 255 of the pixel that is going to be processed are added (Algorithm 2, lines 1-5). Then, the result is  
 256 divided by the number of neighbors ( $K$ ) (Algorithm 2, line 6). After computing the four optimized  
 257 probabilities  $O(q)$  for each pixel, the algorithm selects the highest value and assigns the label of the  
 258 corresponding class to the pixel (Algorithm 2, lines 8-9).



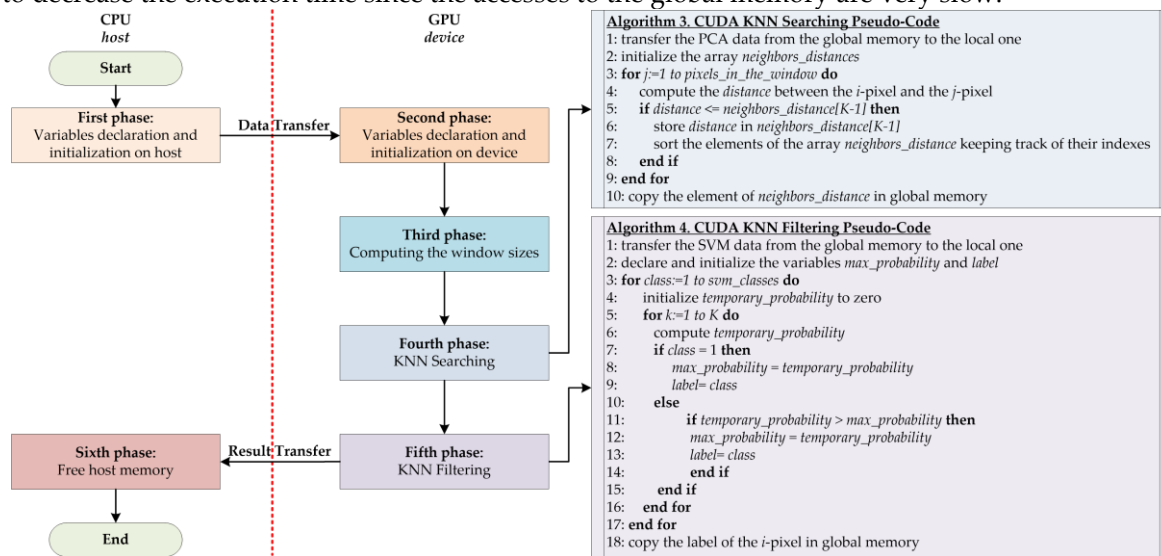
259 **Figure 5.** Flow diagram of the serial implementation of the KNN filtering algorithm.

### 260 2.4.3. Parallel implementation

261 A parallel version of the KNN filtering algorithm has been developed in CUDA in order to  
 262 exploit the NVIDIA GPU technology. The basic idea followed in this approach is that each CUDA  
 263 core has to assign a label to each pixel in parallel. Figure 6 shows the main phases of the parallel  
 264 implementation. The flow starts on the host with the declaration and initialization of all the  
 265 variables (*First phase* in Figure 6). The main difference between this first phase and the  
 266 corresponding one of the serial code is that, in this parallel implementation, the number of arrays,  
 267 structures and variables allocations is decreased in order to save memory.

268 After the first phase, the algorithm transfers to the device the guidance image  $I$  generated by the  
 269 PCA algorithm and the probability maps generated by the SVM classifier. The flow proceeds with  
 270 the resources allocation on the device (*Second phase* in Figure 6). The first step of the KNN filtering  
 271 algorithm on the GPU device concerns the execution of a kernel that evaluates the borders and the  
 272 size of the windows in parallel through the pixels (*Third phase* in Figure 6). Contrary to the serial  
 273 code execution, where the parameters related to the window dimensions are updated at the end of  
 274 the neighbors' selection for each pixel, in the parallel version the algorithm needs to know these  
 275 variables before starting the KNN filtering computation. In fact, in the following steps, it is  
 276 important to copy the PCA and SVM data (already transferred to the device) from the global to the

277 local memory of the GPU, shared by the threads within a block. For this reason, each thread copies  
 278 the part of the data (delimited by the window parameters) needed in the computation. Then, the  
 279 results are copied to the global memory only at the end of the kernel execution. This step is crucial  
 280 to decrease the execution time since the accesses to the global memory are very slow.

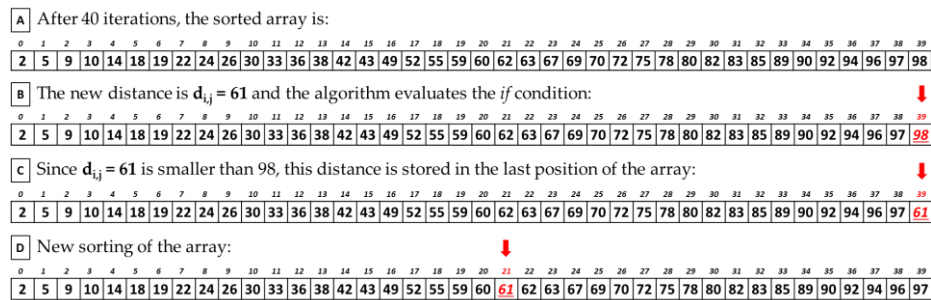


281 **Figure 6.** Flow diagram of the parallel implementation of the KNN filtering algorithm.

282 In the fourth phase, which corresponds to the second phase of the serial code, each thread  
 283 evaluates the  $K$  nearest neighbors of a pixel in parallel (Figure 6, Algorithm 3). First, the PCA data  
 284 required by each block are copied from global to the local memory (Algorithm 3, line 1). Then, each  
 285 thread of the block declares an array called *neighbors\_distances*, whose dimension is equal to the  
 286 number of  $K$  neighbors, which is set to 40 in this work. This array is initialized with large values  
 287 and will contain the 40 lowest distances computed between the pixels (Algorithm 3, line 2). The  
 288 implementation proceeds computing the distance between pixel  $i$ , represented by the thread, and all  
 289 the pixels within its window (Algorithm 3, lines 3-4). If the distance between the pixels  $i$  and  $j$  is  
 290 smaller or equal to the last element of the *neighbors\_distances* array, this distance will be stored in the  
 291 last position of the array (Algorithm 3, lines 5-6). It is important to highlight that this *if condition* is  
 292 always verified considering the first 40 pixels in the window (i.e. first 40 *for* loop iterations). Once the  
 293 first 40 iterations are executed, in the last position of the array there will be a *real* distance (not the  
 294 initialization value) and it will be the highest value among those already present in the array. This is  
 295 verified because every time that a new distance is stored in the array, the algorithm calls a *sort*  
 296 *function* in order to sort the elements of the array in ascending order, keeping track of their indexes  
 297 (Algorithm 3, line 7). The  $K$  indexes of the selected neighbors are the output of the kernel and will be  
 298 copied to the global memory (Algorithm 3, line 10).

299 Figure 7 shows an example of the evaluation of a new distance by the KNN searching  
 300 algorithm. After the first 40 iterations, the array contains 40 distances stored in ascending order  
 301 (Figure 7A). When a new distance is computed (in this example its value is 61), it is compared with  
 302 the last element of the array, in this case located in position 39 and whose value is 98 (Figure 7B).  
 303 Since the new distance is lower than 98, it is stored in the last position of the array (Figure 7C). At  
 304 this point, the array is sorted again (Figure 7D). Due to the reduced dimension of the array to be  
 305 sorted, the sort function implemented in this work is the *shell sort* algorithm.





306 **Figure 7.** KNN searching new distance evaluation example.

307 After computing all the neighbors for all the pixels, the fifth step of the KNN algorithm starts. In  
 308 this phase, the KNN filtering is computed by every thread of each block (Figure 6, Algorithm 4).  
 309 First, each thread copies the SVM probabilities of their corresponding neighbors from the global to  
 310 the local memory (Algorithm 4, line 1). For each class, the algorithm computes the  
 311 *temporary\_probability* value of each pixel, which is the sum of the SVM probabilities of all the  
 312 neighbors of the reference pixel (Algorithm 4, lines 3-6). If the algorithm is executing the first  
 313 iteration of the first *for* loop (i.e. if it is considering the first class), the variable *max\_probability*  
 314 assumes the value of *temporary\_probability* variable and the index of the class is stored in *label*  
 315 (Algorithm 4, lines 7-10). In the following iterations, after computing the *temporary\_probability*, its  
 316 value is stored only if it is higher than the *max\_probability value* (which represents the highest  
 317 probability value of the previous classes). In this case, the index of the class is also stored (Algorithm  
 318 4, line 11-14). At the end of the *for* loop that iterates on the number of classes, the algorithm selects  
 319 the label of the pixel corresponding to the highest sum of probabilities among the four classes. It is  
 320 worth noting that the algorithm evaluates immediately if the sum of probabilities could be the  
 321 highest among the classes or not. This fact means that some arrays declared in the serial version can  
 322 be replaced with a few variables, thus saving memory. At the end of this phase, the label of the pixel,  
 323 which is the output of this step, is stored in the global memory (Algorithm 4, line 18).

324 Once the KNN algorithm execution ends on the GPU device, an array containing the labels of  
 325 all the pixels is transferred from the GPU device to the CPU host. At this point, the memory can be  
 326 released (*Sixth phase* in Figure 6).

### 327 3. Experimental Results and Discussion

328 This section presents the results of the implementations of the KNN-based filtering algorithm  
 329 by evaluating different sets of parameters in order to evaluate them, analyzing both the  
 330 computational times and the classification accuracy.

#### 331 3.1. KNN window search optimization results with Euclidean distance

332 In section 2.4.1, it was described an important optimization introduced in the serial and parallel  
 333 implementations concerning the computation of the distances between pixels inside a window and  
 334 not within the entire image. Reducing the space where the algorithm evaluates the distances ensures  
 335 a significant decrease of the computational time, as shown in Table 2. In particular, the table  
 336 provides the execution times for all the images, considering both the case in which the neighbors are  
 337 searched within the entire image (*EI*) and within a window with 14 rows (*WSize14*). The speedup  
 338 obtained with the optimization has been also included. In addition, this table shows the total  
 339 number of pixels of each image and the number of pixels inside the smallest and the biggest window  
 340 in the *WSize14* implementation. The times refer to tests where the Euclidean distance has been  
 341 considered. The simulations of the serial code have been carried out on an Intel i7 processor,  
 342 working at 3.50 GHz, equipped with 16 GB RAM.

343

344  
345**Table 2.** Execution times of the serial code considering as search space both the entire image (*EI Time*) and a window with 14 rows (*WSize14 Time*).

Image ID	#Pixels	EI Time [s]	WSize14 Time [s]	Speedup	Min WSize14 [#Pixels]	Max WSize14 [#Pixels]
P1C1	251,532	17,173.74	503.89	34.08x	3,836	7,672
P1C2	264,408	19,135.58	509.16	37.58x	3,864	7,728
P2C1	219,232	15,630.77	374.67	41.72x	3,472	6,944
P3C1	185,368	9,788.58	322.86	30.32x	3,451	6,902
P4C1	124,691	4,015.89	139.30	28.83x	2,303	4,606

346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358

Data presented in Table 2 show that this optimization allows a huge decrease in the execution times. For example, considering the biggest image of the dataset, P1C2, the time of the implementation that considers the entire image is 19,135.58 seconds (about 5 hours and 30 minutes). Considering a window of 14 rows as neighbors search space, this time decreases to 509.16 seconds (about 8 minutes). The reason of this huge time difference is that, when the algorithm has to consider the entire image, it has to compute a number of distances equal to  $(264,408-1)$  for each pixel, where 264,408 is the number of pixels of the P1C2 image. Considering the window technique, the algorithm computes a number of distances that, for the same image, varies from  $(3,864-1)$  to  $(7,728-1)$ , where 3,864 and 7,728 are the number of pixels inside the windows with the minimum and the maximum sizes respectively (depending whether the pixel is in the borders or in the center of the image). Concerning to the classification results, it is important to notice that there are no differences in the results, and therefore all the pixels are classified with the same labels, using either the entire image or a window.

359  
360  
361  
362  
363  
364  
365  
366  
367  
368

Considering this significant result, the computational time variations were evaluated when the window size was reduced. Furthermore, since the main goal of the work was to reach real-time execution, a parallel version of the algorithm was developed in CUDA language to exploit the GPU technology. The GPU device used during the tests was an NVIDIA Tesla K40 GPU. This board is based on the Kepler architecture (working at 875 MHz) and it is equipped with 2,880 CUDA cores and 12 GB GDDR5 memory with a peak bandwidth of 288 GB/s [24]. The board is connected to the CPU host through a PCI Express 2.0. **Errore. L'origine riferimento non è stata trovata.** shows the execution times of the serial and parallel implementations characterized by window sizes that vary from 14 to 2 with decrements of 2. In addition, the speedups between the serial and the parallel codes are presented.

369  
370**Table 3.** Execution time results of the serial and parallel implementations using the Euclidean distance employing different window sizes.

Image ID	Processing Type	Processing Time [s]						
		WSize14	WSize12	WSize10	Wsize8	Wsize6	Wsize4	WSize2
P1C1	Serial	503.89	406.32	383.71	262.00	221.71	118.25	59.08
	CUDA	12.83	11.49	6.23	5.52	3.85	2.29	1.22
	Speedup	39.25x	35.33x	61.59x	47.42x	57.52x	51.44x	48.10x
P1C2	Serial	509.16	424.22	408.52	276.06	235.76	125.34	62.36
	CUDA	13.53	12.09	6.47	5.73	3.99	2.39	1.26
	Speedup	37.62x	35.08x	63.07x	48.11x	59.04x	52.31x	49.15x
P2C1	Serial	374.67	315.73	302.54	239.58	151.23	95.02	47.06
	CUDA	10.55	5.70	5.18	3.62	2.67	1.70	1.03
	Speedup	35.51x	55.39x	58.30x	66.18x	56.58x	55.76x	45.62x
P3C1	Serial	322.86	263.40	254.30	202.56	122.16	78.47	39.97
	CUDA	9.00	4.92	4.45	3.15	2.30	1.51	0.92
	Speedup	35.85x	53.46x	57.06x	64.17x	52.92x	51.80x	43.16x
P4C1	Serial	139.30	118.94	115.07	90.81	55.29	35.84	18.11
	CUDA	3.21	2.34	2.16	1.63	1.12	0.83	0.60
	Speedup	43.38x	50.63x	53.23x	55.61x	49.01x	42.82x	30.13x


371

372 The reduction of the window size supposes a decrease in the execution times because the  
 373 algorithm has to compute a lower number of distances. For example, considering the P1C2 image,  
 374 the time varies from 509.16 seconds (~8 minutes) to 62.36 seconds (~1 minute) in the serial versions of  
 375 *WSize14* and *WSize2*, respectively. If the parallel implementation of the same image is considered,  
 376 the times present a further decrease. In fact, for the same image the parallel version of *WSize14* is  
 377 ~37x times faster than the serial version, taking only 13.53 seconds instead of ~8 minutes. At the same  
 378 time, the parallel execution of *WSize2* takes only 1.26 seconds instead of ~1 minute (the speedup is  
 379 ~49x). Concerning all the images in the reference implementation *WSize14*, the speedups are always  
 380 higher than 35x and in the best case (P4C1) it reaches 43x. If we consider all the other versions, with  
 381 the decreased windows sizes, the parallel code shows even higher speedups. For example,  
 382 considering the P1C2 image and the window size *WSize10*, the parallel code takes 6.48 seconds while  
 383 the serial version takes 408.52 seconds (~6 minutes), obtaining a speedup of ~63x. Nevertheless, it is  
 384 necessary to examine these times and speedups also taking into account the classification results. It  
 385 is very important to consider if, reducing the window size, there are pixels classified with different  
 386 labels compared to the *reference* version (*WSize14*). Table 4 shows the number of misclassified pixels  
 387 between the reference result and other window sizes. Additionally, the percentage of the difference  
 388 is shown.

389 **Table 4.** Number of pixels with different classification result using the Euclidean distance  
 390 between the different computed windows sizes and the reference one (*WSize14*).

Image ID	#Pixels	#Different pixels (% of difference) compared to the reference ( <i>WSize14</i> )					
		<i>WSize12</i>	<i>WSize10</i>	<i>WSize8</i>	<i>WSize6</i>	<i>WSize4</i>	<i>WSize2</i>
P1C1	251,532	0 (0.000%)	0 (0.000%)	65 (0.026%)	3,672 (1.460%)	9,096 (3.616%)	20,476 (8.141%)
P1C2	264,408	0 (0.000%)	1 (0.000%)	60 (0.023%)	2,845 (1.076%)	7,705 (2.914%)	22,054 (8.341%)
P2C1	219,232	0 (0.000%)	4 (0.002%)	65 (0.030%)	2,606 (1.189%)	6,532 (2.979%)	18,015 (8.217%)
P3C1	185,368	0 (0.000%)	1 (0.000%)	49 (0.026%)	2,273 (1.226%)	5,604 (3.023%)	13,981 (7.542%)
P4C1	124,691	3 (0.002%)	7 (0.005%)	71 (0.057%)	1,498 (1.201%)	3,733 (2.993%)	10,089 (8.091%)

391 Considering the first three windows sizes (*WSize12*, *WSize10*, *WSize8*) for all the images, the  
 392 number of pixels classified with different labels is very low, taking into account the final application  
 393 of the system. In fact, the highest percentage of different pixels is 0.057% and it is related to the P4C1  
 394 image, which, in the version *WSize8*, presents 71 different pixels on a total amount of 124,691 pixels.  
 395 Concerning the other three windows sizes, the highest percentage of different pixels for window  
 396 *WSize6* is 1.46% considering the P1C1 image (3,672 different pixels on 251,532). For the window  
 397 *WSize4*, the percentage of different pixels is ~3.62%, referred also to the P1C1 image (9,096 different  
 398 pixels on 251,532) and for *WSize2*, the highest percentage is ~8.341%, considering the biggest image  
 399 of the database, P1C2 (22,054 different pixels on 264,408). At this point, there is a further evaluation  
 400 that can be made considering that this algorithm is part of a system whose main goal is to  
 401 discriminate between tumor and healthy tissue. Despite this, the classification is made between four  
 402 classes that are normal tissue, tumor tissue, hypervascularized tissue and background [23]. From the  
 403 surgical and medical point of view, it is clear that a wrong discrimination between tumor and  
 404 healthy tissue has much greater and transcendental relevance than just a misclassification issue  
 405 between tumor and any other classes (hypervascularized and background) or between healthy,  
 406 hypervascularized and background classes. It is possible to re-evaluate again the results of Table 4,  
 407 considering that in the different *WSize* executions only a low percentage of different pixel labels are  
 408 exchanged between tumor and normal tissue. Figure 8 shows the percentage of pixels that are  
 409 misclassified between tumor and healthy tissues, tumor and hypervascularized tissues and tumor  
 410 and background, considering all the windows sizes for each image compared to the reference  
 411 version. In addition, the graph presents the classification differences between healthy,  
 412 hypervascularized and background classes (called *Others*).  
 413



**Figure 8.** Percentage of pixels that have been misclassified using the Euclidean distance between tumor and healthy tissues (blue), tumor and hypervascularized tissues (orange), tumor tissue and background (gray) and the other misclassifications between healthy, hypervascularized and background (yellow). The results were obtained per each window size implementations compared to the *WSize14* for each image of the dataset.

414  
415  
416  
417  
418

419 As it can be seen in Figure 8, only in the case of the P3C1 image using the *WSize8*, the algorithm  
420 misclassifies approximately 2% of the pixels (1 out of 49 pixels), exchanging the labels between  
421 tumor and healthy tissues. In all the other implementations of *WSize8*, the classification differences  
422 do not involve the tumor class. Furthermore, in the versions related to the three smallest windows  
423 (*WSize6*, *WSize4*, *WSize2*), the percentage of the pixels exchanged between these two classes is lower  
424 than the percentages of pixels exchanged between the other classes. For example, for the biggest  
425 image of the database (P1C2), in the *WSize6* implementation, the classification difference between  
426 tumor and healthy tissue represents 2.43% out of 2,845 different pixels. Considering the same image  
427 in the *WSize4* and in the *WSize2* implementations, this percentage is of 2.60% out of 7,705 pixels and  
428 1.89% out of 22,054 pixels, respectively. The highest percentage of difference between these two  
429 classes is found in the *WSize6* version regarding the P1C1 image, where it is about 3.57% out of 3,672  
430 pixels. According to these data, it is clear that the algorithm can correctly distinguish the tumor from  
431 the healthy tissue, while it makes more errors in separating the tumor from the hypervascularized  
432 tissue. The highest percentages of misclassified pixels between the tumor and the hypervascularized  
433 classes reach 26.28% of the total number of different pixels (P3C1 image, *WSize4* version). In fact,  
434 according to what it is said in [23], these two classes referred to tissues with similar spectral  
435 signatures that can produce some misclassifications. On the other hand, the spectral signatures of  
436 tumor and healthy tissues present remarkable differences that allow the algorithm to distinguish  
437 these two classes in the classification.

### 438 3.2. KNN window search optimization results using Manhattan distance

439 As it was said before, the neighbors search supposes the heaviest computational load of the  
440 KNN filtering algorithm. Although the distances computation is the most time-consuming task, the  
441 number of evaluated distances has been reduced in this study by considering a window, so a  
442 smallest area is considered instead of the entire image. To further reduce the execution time of this  
443 phase, the Manhattan metric has been tested instead of the Euclidean one, as described in Equation  
444 (3). Table 5 compares the times of the serial code using both the entire image (*EI*) and the reference  
445 window (*WSize14*), employing both the Euclidean and the Manhattan distances. The speedup  
446 obtained using the Manhattan distance and the percentages of pixels that are different in the results  
447 are also presented in this table.

448 As said in the previous paragraph, searching the neighbors inside a window instead of the  
449 entire image allows saving time without changing the results of the classification. A further  
450 reduction of the execution time is obtained using the Manhattan metric in the distance  
451 computations. In fact, for the biggest image of the database (P1C2), the time is reduced from ~5  
452 hours (19,135.58 seconds) using the Euclidean distance to ~2 hours (7,683.44 seconds) in the case of  
453 using the entire image. If the neighbors are searched within the window (for the same image), the

454 time decreases to ~3 minutes (202.42 seconds) using the Manhattan distance. Concerning all the  
 455 images, it is possible to reach speedups from 2.22x to 3.33x, considering the versions with the entire  
 456 image, and from 2.49x to 2.66x in the *WSize14* executions. Comparing the implementations that  
 457 exploit the Manhattan distance and the ones that use the Euclidean metric, the number of pixels  
 458 classified with different labels is quite low: the highest percentage of different pixels is 1.33% in the  
 459 P1C1 image. Furthermore, it is important to highlight that there are no differences in the  
 460 classification results comparing the entire image and the *WSize14* versions, using the Manhattan  
 461 distance.

462 **Table 5.** Comparison of the execution time of the serial versions obtained employing the  
 463 Euclidean and Manhattan distances with the entire image (*EI*) and the *WSize14*. The  
 464 table also presents the classification differences between the Euclidean and  
 465 Manhattan implementations.

Distance Type	P1C1		P1C2		P2C1		P3C1		P4C1	
	<i>EI</i>	<i>WSize14</i>	<i>EI</i>	<i>WSize14</i>	<i>EI</i>	<i>WSize14</i>	<i>EI</i>	<i>WSize14</i>	<i>EI</i>	<i>WSize14</i>
Euclidean [s]	17,173.75	503.89	19,135.58	509.17	15,630.77	374.67	9,788.58	322.87	4,015.89	139.30
Manhattan [s]	7,222.13	190.02	7,683.44	202.42	4,735.87	146.93	3,382.84	121.37	1,807.91	55.91
Speedup	2.38x	2.65x	2.49x	2.51x	3.3x	2.55x	2.89x	2.66x	2.22x	2.49x
Difference	1.33%		0.99%		1.03%		1.10%		1.06%	

466 At this point, it is interesting to evaluate how the execution times can be reduced changing the  
 467 size of the windows using the Manhattan metric in the distances computation. The results shown in  
 468 Table 6 confirm that decreasing the number of distance computations, i.e. the variations of the  
 469 window sizes, allows further reductions of the computational time. The lowest execution times are  
 470 obtained exploiting the GPU technology that can run the parallel algorithm taking ~8 seconds  
 471 (compared to ~3 minutes) if the biggest image (P1C2) with the *WSize14* version is considered. The  
 472 speedups obtained using this device and the optimizations introduced in the code are significant  
 473 and they can reach up to 33.2x (P2C1-*WSize12*). For some images and for some window dimensions,  
 474 the algorithm takes only a few seconds, but what is even more important to consider is the number  
 475 of pixels that are misclassified when the window size decreases (0).

476 **Table 6.** Execution time results of the serial and parallel implementations using the  
 477 Manhattan distance and using different window sizes.

Image ID	Processing Type	Processing Time [s]						
		<i>WSize14</i>	<i>WSize12</i>	<i>WSize10</i>	<i>Wsize8</i>	<i>Wsize6</i>	<i>Wsize4</i>	<i>WSize2</i>
P1C1	Serial	190.02	192.19	158.15	129.54	81.61	54.14	28.70
	CUDA	7.63	7.07	5.04	4.62	3.38	2.04	1.18
	Speedup	24.90x	27.15x	31.34x	27.99x	24.09x	26.47x	24.13x
P1C2	Serial	202.41	204.65	169.23	138.12	84.98	57.51	29.75
	CUDA	8.01	7.40	5.21	4.84	3.51	2.09	1.22
	Speedup	25.26x	27.64x	32.48x	28.52x	24.20x	27.45x	24.28x
P2C1	Serial	146.92	152.20	125.60	102.94	63.90	42.57	21.81
	CUDA	6.44	4.58	4.27	3.16	2.34	1.52	1.00
	Speedup	22.81x	33.20x	29.35x	32.50x	27.25x	27.86x	21.66x
P3C1	Serial	121.37	126.98	104.79	86.83	55.09	36.24	18.54
	CUDA	5.57	4.03	3.76	2.79	2.04	1.37	0.90
	Speedup	21.75x	31.47x	27.86x	31.11x	26.88x	26.42x	20.54x
P4C1	Serial	55.91	58.06	42.62	39.42	24.39	16.62	8.66
	CUDA	2.81	2.12	1.98	1.49	1.04	0.80	0.60
	Speedup	19.87x	27.27x	21.46x	26.46x	23.31x	20.59x	14.42x

479  
480**Table 7.** Number of pixels with different classification results using the Manhattan distance between the computed window sizes and the *Wsize14* Manhattan.


Image ID	#Pixels	#Different pixels (% of difference) compared to the <i>Wsize14</i> (Manhattan) version					
		<i>WSize12</i>	<i>WSize10</i>	<i>WSize8</i>	<i>WSize6</i>	<i>WSize4</i>	<i>WSize2</i>
P1C1	251,532	2 (0.001%)	65 (0.026%)	1,832 (0.728%)	4,507 (1.792%)	8,889 (3.534%)	19,783 (7.865%)
P1C2	264,408	3 (0.001%)	56 (0.021%)	1,483 (0.561%)	3,891 (1.471%)	7,839 (2.965%)	21,714 (8.812%)
P2C1	219,232	3 (0.001%)	44 (0.020%)	1,320 (0.602%)	3,483 (1.589%)	6,743 (3.075%)	17,653 (8.052%)
P3C1	185,368	2 (0.001%)	48 (0.026%)	1,033 (0.557411%)	2,825 (1.524%)	5,474 (2.953%)	13,436 (7.248%)
P4C1	124,691	2 (0.001%)	35 (0.028%)	699 (0.560%)	2,014 (1.615%)	3,831 (3.072%)	9,613 (7.709%)

481  
482  
483  
484  
485  
486  
487  
488  
489

As it can be seen in the results shown in 0, *WSizes12* and *WSizes10* present a reduced number of different pixels compared to the other implementations. Analyzing the Euclidean distance results presented in Table 4, this consideration can be made for the first three tests (*WSizes12*, *WSizes10* and *WSize8*) but, in this case, the number of different pixels in *WSize8* is higher than the first two versions. Despite this, it is important to highlight that the classification differences shown in 0 are not very relevant for the final application of the system. In this application, a solution with a good compromise between real-time execution and classification accuracy of the results has to be selected. In addition, it is also important to evaluate the percentage of different pixels that are misclassified between tumor and healthy tissues and between tumor and the other classes.

490  
491  
492  
493  
494  
495  
496  
497  
498

Figure 9 shows the percentage of pixels that are misclassified using the Manhattan metric between the different classes. In this figure, it is possible to notice that the algorithm misclassifies more pixels between tumor and hypervascularized classes than between tumor and healthy classes. In fact, the highest percentage of pixels misclassified is 30.77% related to the P1C1 image with *WSize10*, where the algorithm exchanges the labels of 20 pixels (between tumor and hypervascularized tissue) out of a total amount of 65 different pixels compared to the reference version *WSize14* (0). Concerning the comparison between tumor and healthy classes, the number of pixels with an exchanged label is very low: the worst case is always the P1C1 image (*WSize8*), where 66 out of 1,832 different pixels are misclassified, being a 3.60% of pixels.

499  
500  
501  
502  
503

**Figure 9.** Percentage of misclassified pixels using the Manhattan distance between tumor and healthy tissues (blue), tumor and hypervascularized tissues (orange), tumor tissue and background (gray) and the other misclassifications between healthy, hypervascularized and background (yellow). The results were obtained per different window sizes implementation compared to the *WSize14* for each image of the dataset.

504 

### 3.3. Summary

505  
506  
507  
508  
509  
510  
511

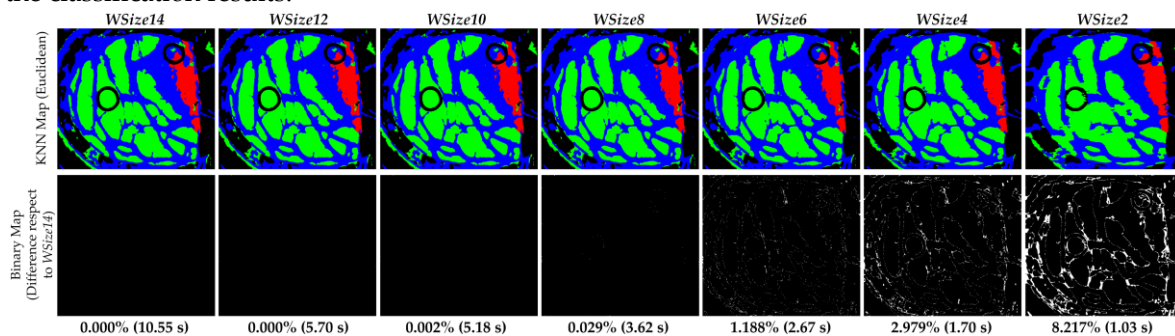
In this study, the results of serial and parallel versions of the KNN filtering algorithm for the classification of in-vivo brain tumor from hyperspectral images are presented. In particular, the importance of reducing the area of the neighbors search in order to decrease the elaboration time is explained. In fact, the results prove that searching the neighbors of a pixel within a window instead of the entire image supposes a significant reduction of the computation time. It is important to notice that introducing a window (characterized by 14 rows with the reference pixel in the center) does not affect the result of the classification. For this reason, this version has been defined as the reference

512 one (*WSize14*). Reducing the window size compared to the reference one, the time of the  
 513 computation drastically decreases but the number of pixels that the algorithm misclassifies  
 514 increases. At this point, it is important to select the best versions that have a good tradeoff between  
 515 performance and number of misclassifications.

516 In the previous sections, the percentages and the number of different pixels between versions  
 517 with different window sizes were analyzed. Concerning the implementations that exploit the  
 518 Euclidean distance, Table 4 demonstrated that by using the window sizes *WSize12*, *WSize10* and  
 519 *WSize8*, the number of different pixels was lower than those obtained using the *WSize6*, *WSize4* and  
 520 *WSize2*, compared with the reference test (*WSize14*).

521 In Figure 10, the KNN filtered maps obtained from the P2C1 image and the binary maps are  
 522 shown, where the differences between the evaluated window size version and the reference version  
 523 are highlighted. Despite the differences between the first and the last three versions shown in Table  
 524 4, it is possible to see that the KNN filtered maps of the implementations for *WSize6* and *WSize8* do  
 525 not present relevant dissimilarities. In fact, it is important to remember that the main goal of these  
 526 maps is to delineate the tumor area, in order to provide a guidance tool to the surgeons during the  
 527 tumor resection. In this context, it is clear that the number of different pixels in *WSize8* and *WSize6*  
 528 versions are not so significant for the final application of the system, since the surgeon always resect  
 529 a security margin around the tumor tissue.

530 In the previous paragraph, the data showed that the algorithm is able to correctly discriminate  
 531 between tumor and healthy classes. This consideration can also be seen in the KNN filtered maps,  
 532 where the area related to the tumor tissue (red) remains roughly the same in the implementations  
 533 *WSize6*, *WSize8*, *WSize10* and *WSize12* compared to the *WSize14* one. Considering the *WSize4* and the  
 534 *WSize2* versions, it is possible to appreciate that the margins of the tumor are not as evident and well  
 535 defined as in the other images, confirming what has been said in the previous paragraphs analyzing  
 536 the classification results.

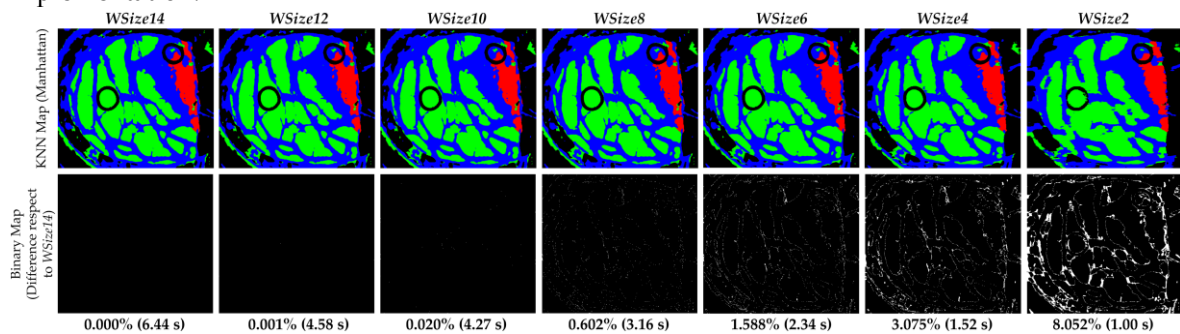


537 **Figure 10.** Results of the KNN filtering algorithm obtained from the P2C1 image using the  
 538 Euclidean distance. The first row shows the filtered classification maps generated using  
 539 different window sizes. The second row presents the binary maps where the pixels  
 540 differences between the current generated map and the reference one (*WSize14*) are shown.  
 541 In addition, the percentage of differences and the execution time results are detailed.

542 In the second row of Figure 10 are presented the binary maps that show the pixel differences  
 543 between all the window sizes versions and the reference implementation (*WSize14*). In particular, by  
 544 analyzing the binary maps of *WSize4* and *WSize2*, it is possible to identify several differences  
 545 compared to *WSize14*. For this reason, these two versions should not be chosen for the final solution.  
 546 However, in the binary maps of *WSize6* and *WSize8*, there are few differences, and they are barely  
 547 appreciated analyzing the KNN filtered maps. It is important to remember that the suitable version  
 548 for this application is the option that offers a good compromise between accurate classification and  
 549 fast execution. Exploiting the GPU technology, the parallel version of the KNN algorithm with  
 550 *WSize8* employs ~3.62 seconds to filter the P2C1 image, while the *WSize6* implementation is executed  
 551 in ~2.67 seconds. For the biggest image of the database (P1C2), the *WSize6* implementation allows to  
 552 save ~2 seconds compared to the *WSize8* version. According to these results, the *WSize8* version has  
 553 been selected as the best solution, giving priority to the classification accuracy but considering also a

554 fast implementation. On the contrary, the *WSize6* implementation has been chosen as the fastest  
 555 implementation with acceptable accuracy results.

556 Similarly, the same evaluation can be done considering the implementations that exploit the  
 557 Manhattan metric for the computation of the distances. Analyzing the computational times in the  
 558 previous sections, it is evident that this metric leads to faster executions than using the Euclidean  
 559 distance. In the first row of Figure 11 are shown the KNN filtered maps of the P2C1 image using  
 560 different window sizes and employing the Manhattan distance. In the second row are presented the  
 561 binary maps to evaluate the differences between the developed versions compared to the *WSize14*  
 562 implementation.



563 **Figure 11.** Results of the KNN filtering algorithm obtained from the P2C1 image using the  
 564 Manhattan distance. The first row shows the filtered classification maps generated using  
 565 different window sizes. The second row presents the binary maps where the pixels  
 566 differences between the current generated map and the *WSize14*. In addition, the percentage  
 567 of differences and the execution time results are detailed.

568 According to the data shown in 0, in Figure 11 it is possible to see that the KNN filtered maps of  
 569 the versions *WSize12* and *WSize10* are practically identical to the map obtained with the *WSize14*.  
 570 The number of different pixels is low enough to not perceive the differences between the  
 571 classification maps. In Table 7, it is also evident that the number of different pixels from the *WSize8*  
 572 to the *WSize2* implementations drastically increases. Concerning the KNN filtered maps of the  
 573 *WSize4* and, in particular, the *WSize2* implementations, the differences are very clear since the  
 574 margin of the tumor is not as well defined as in the other maps. Instead, in the filtered maps of the  
 575 *WSize8* and *WSize6* implementations, the classification differences are not so evident, especially  
 576 taking into account the tumor tissue area. The differences between all the versions compared to the  
 577 *WSize14* implementation can be evaluated in the binary maps (Figure 11, second row). Even if the  
 578 *WSize4* and *WSize2* are the fastest implementations, their binary maps clearly show that these two  
 579 versions cannot be chosen because the amount of different pixels compared to *WSize14* is too high.  
 580 However, in the binary maps of *WSize12* and *WSize10*, it is evident that these implementations offer  
 581 the highest accuracy but the slowest execution times. Finally, concerning the *WSize8* and *WSize6*  
 582 implementations, it is possible to determine that the *WSize8* version has the highest accuracy but the  
 583 execution time is slower than the *WSize6* version (the former exhibits 3.16 seconds and the latter 2.34  
 584 seconds). Also in this case, the best solution is chosen on the base of the degree of accuracy and the  
 585 time constraints that the application requires.

586 At this point, the best solutions selected between the Manhattan versions (*WSize8* and *WSize6*)  
 587 have to be compared with the reference test *WSize14* that exploits the Euclidean distance. In fact, the  
 588 original algorithm is characterized by the use of the Euclidean metric in the neighbors search within  
 589 the entire image. Since the *WSize14* Euclidean implementation does not have any differences in the  
 590 classification results compared to the original version, the results of the Manhattan best solutions  
 591 have to be compared with the reference results.

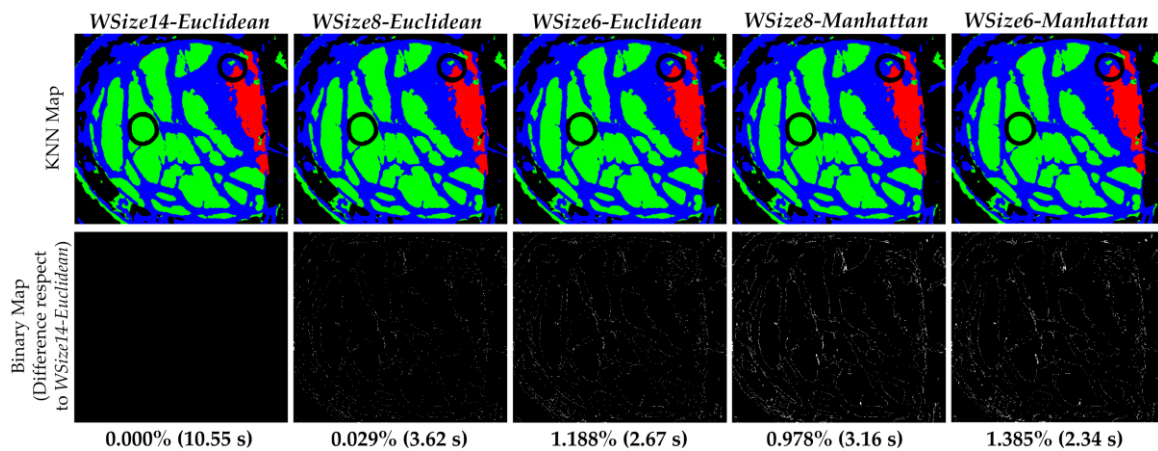
592 By analyzing the results comparison shown in Figure 12, it is possible to see that all these  
 593 versions have a reduced percentage of different pixels compared to the *WSize14-Euclidean*  
 594 implementation. In all the obtained KNN filtered maps, the boundaries of the tumor area are  
 595 accurately defined. The solutions where the results are more similar to the reference implementation



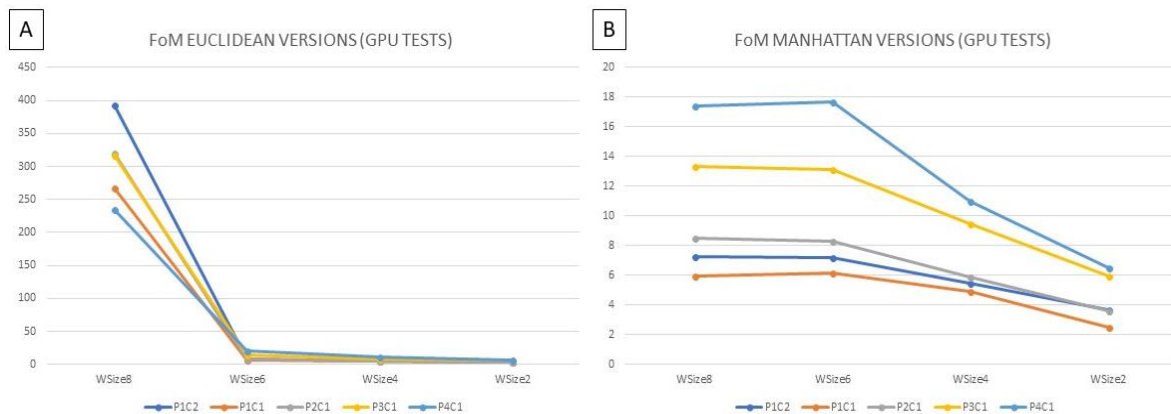
596 are the *WSize8-Euclidean* and *WSize8-Manhattan* versions, which differ 0.029% and 0.978%  
 597 respectively, compared to the *WSize14-Euclidean* reference. The versions characterized by a window  
 598 with 6 rows are less accurate than the previous ones, but they are faster. Concerning the  
 599 computational times, the parallel execution of the reference solution is executed in ~10.55 seconds,  
 600 while the *WSize8-Euclidean* and *WSize8-Manhattan* versions are executed in 3.62 and 3.16 seconds  
 601 respectively. The *WSize6-Euclidean* and *WSize6-Manhattan* implementations require 2.67 and 2.34  
 602 seconds respectively. Finally, a figure of merit (*FoM* in Equation (5)) which relates the execution  
 603 time (*t*) and the classification results (*err*) was considered to select the best solution that offers the  
 604 highest value. The version *WSize8-Euclidean* is chosen as the best solution since it presents the  
 605 highest value of *FoM* (Figure 13). To the best of our knowledge, the state of the art does not provide  
 606 implementations of the KNN filtering algorithm that could be a touchstone for a fair comparison  
 607 with the presented work.

$$FoM = \frac{1}{(t * err)} \tag{5}$$

608



609 **Figure 12.** Results comparison of the KNN filtered maps from the P2C1 image using Manhattan  
 610 and Euclidean distances. The first row shows the filtered classification maps generated  
 611 using different window sizes and distance metrics. The second row presents the binary  
 612 maps where the pixels differences between the current generated map and the reference one  
 613 (*WSize14-Euclidean*) are shown. In addition, the percentage of differences and the execution  
 614 time results are detailed.



615 **Figure 13.** Figure of metric computed comparing (A) the Euclidean versions *WSize8*, *WSize6*,  
 616 *WSize4* and *WSize2* with the reference *WSize14-Euclidean*, (B) the Manhattan versions  
 617 *WSize8*, *WSize6*, *WSize4* and *WSize2* with the reference *WSize14-Euclidean*.

618

#### 619 4. Conclusions

620 This work presents the development of a parallel version of the KNN filtering algorithm  
621 exploiting the NVIDIA Tesla K40 board. The goal of the implementation was to reduce the execution  
622 time of the KNN filter to reach real-time constraints, which is mandatory considering the final  
623 application of the system. This application is related with the detection and identification of in-vivo  
624 brain tumor boundaries during neurosurgical operations by using hyperspectral images. For every  
625 pixel of the image, the parallel version of the algorithm computes each phase of the algorithm  
626 simultaneously. A first optimization was to introduce a *search window* in the K nearest neighbors  
627 search step, which is the most time-consuming part of the algorithm. The selection of the neighbors  
628 within a region close enough to the pixel, instead of the entire image, allows to significantly reduce  
629 the computational time of the algorithm. Furthermore, variations of the window sizes have been  
630 explored in order to evaluate the accuracy of the results and a possible reduction of the  
631 computational time. All the variations were considered exploiting both the Euclidean and the  
632 Manhattan metrics for the distance computation. The results obtained in this analysis show that, for  
633 the proposed final application, the implementation characterized by a search window of 8 rows  
634 using Euclidean distance is the best solution. This version performs the classification of the  
635 considered images in less than 6 seconds, with speedups up to 102.5x and 4317.9x compared with  
636 the *Wsize14-Euclidean* and the *entire image* versions, respectively. Further developments must be  
637 carried out to integrate this parallel version of the KNN filtering algorithm with the other parts of  
638 the brain cancer detection algorithm (i.e., SVM classifier and PCA) in a single system capable of  
639 computing the classification maps of the hyperspectral brain cancer images in surgical-time to assist  
640 neurosurgeons during the resection of the tumor tissues.

641 **Acknowledgments:** This work has been supported in part by the Canary Islands Government through the  
642 ACIISI (Canarian Agency for Research, Innovation and the Information Society), ITHACA project  
643 “Hyperspectral identification of Brain tumors” under Grant Agreement ProID2017010164. This work has been  
644 also supported in part by the European Commission through the FP7 FET (Future Emerging Technologies)  
645 Open Programme ICT- 2011.9.2, European Project HELICoiD “HypErspectral Imaging Cancer Detection” under  
646 Grant Agreement 618080. Additionally, this work has been supported in part by the 2016 PhD Training  
647 Program for Research Staff of the University of Las Palmas de Gran Canaria and the 2017 PhD Training  
648 Program for Research Staff of the Canary Islands Government through the ACIISI (Canarian Agency for  
649 Research, Innovation and the Information Society). The authors would like to thank NVIDIA Corporation for  
650 the donation of the NVIDIA Tesla K40 GPU used for this research.

651 **Author Contributions:** G.F. performed the GPU implementations, the algorithms optimizations, designed and  
652 performed experiments, and wrote the manuscript. H.F. and E.T. designed and performed experiments and  
653 edited the manuscript. R.L., D.M., S.O. and R.Sal. performed the serial algorithm implementation and  
654 optimizations, performed experiments, and edit the manuscript. F.L., G.D., A.B.-Q., G.M.C., E.J., C.S. and  
655 R.Sar. supervised the project and edited the manuscript.


656 **Conflicts of Interest:** The authors declare no conflict of interest. The founding sponsors had no role in the  
657 design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, and in  
658 the decision to publish the results.

#### 659 5. References

- 660 1. Calin, M. A.; Parasca, S. V.; Savastru, D.; Manea, D. Hyperspectral Imaging in the Medical Field: Present  
661 and Future. *Appl. Spectrosc. Rev.* **2013**, *49*, 435–447, doi:10.1080/05704928.2013.838678.
- 662 2. Lu, G.; Fei, B. Medical hyperspectral imaging: a review. *J. Biomed. Opt.* **2014**, *19*, 10901,  
663 doi:10.1117/1.JBO.19.1.010901.
- 664 3. Chang, C.-I. *Hyperspectral imaging: techniques for spectral detection and classification*; Springer Science &  
665 Business Media, 2003; Vol. 1;
- 666 4. Akbari, H.; Kosugi, Y. Hyperspectral imaging: A new modality in surgery. In *Recent advances in*  
667 *biomedical engineering*; InTech, 2009.

- 668 5. Vo-Dinh, T. A hyperspectral imaging system for in vivo optical diagnostics. *Eng. Med. Biol. Mag. IEEE*  
669 2004, 23, 40–49.
- 670 6. Panasyuk, S. V.; Yang, S.; Faller, D. V.; Ngo, D.; Lew, R. A.; Freeman, J. E.; Rogers, A. E. Medical  
671 hyperspectral imaging to facilitate residual tumor identification during surgery. *Cancer Biol. Ther.* **2007**,  
672 6, 439–446, doi:10.4161/cbt.6.3.4018.
- 673 7. Mori, M.; Chiba, T.; Nakamizo, A.; Kumashiro, R.; Murata, M.; Akahoshi, T.; Tomikawa, M.; Kikkawa,  
674 Y.; Yoshimoto, K.; Mizoguchi, M.; Sasaki, T.; Hashizume, M. Intraoperative visualization of cerebral  
675 oxygenation using hyperspectral image data: a two-dimensional mapping method. *Int. J. Comput. Assist.*  
676 *Radiol. Surg.* **2014**, 9, 1059–1072, doi:10.1007/s11548-014-0989-9.
- 677 8. Plaza, A.; Plaza, J.; Paz, A.; Sanchez, S. Parallel Hyperspectral Image and Signal Processing. *IEEE Signal*  
678 *Process. Mag.* 2011, 28, 119–126.
- 679 9. Kabwama, S.; Bulters, D.; Bulstrode, H.; Fabelo, H.; Ortega, S.; Callico, G. M.; Stanciulescu, B.; Kiran, R.;  
680 Ravi, D.; Szolna, A.; others Intra-operative hyperspectral imaging for brain tumour detection and  
681 delineation: Current progress on the HELICoid project. *Int. J. Surg.* **2016**, 36, S140.
- 682 10. Salvador, R.; Ortega, S.; Madroñal, D.; Fabelo, H.; Lazcano, R.; Marrero, G.; Juárez, E.; Sarmiento, R.;  
683 Sanz, C. HELICoid: Interdisciplinary and collaborative project for real-time brain cancer detection. In  
684 *ACM International Conference on Computing Frontiers 2017, CF 2017*; 2017.
- 685 11. Szolna, A.; Morera, J.; Piñeiro, J. F.; Callicó, G. M.; Fabelo, H.; Ortega, S. Hyperspectral Imaging as A  
686 Novel Instrument for Intraoperative Brain Tumor Detection. *Neurocirugia* **2016**, 27, 166.
- 687 12. Fabelo, H.; Ortega, S.; Lazcano, R.; Madroñal, D.; M. Callicó, G.; Juárez, E.; Salvador, R.; Bulters, D.;  
688 Bulstrode, H.; Szolna, A.; Piñeiro, J. F.; Sosa, C.; J. O’Shanahan, A.; Bisshopp, S.; Hernández, M.; Morera,  
689 J.; Ravi, D.; Kiran, B. R.; Vega, A.; Báez-Quevedo, A.; Yang, G.-Z.; Stanciulescu, B.; Sarmiento, R. An  
690 intraoperative visualization system using hyperspectral imaging to aid in brain tumor delineation.  
691 *Sensors* **2018**, 18, doi:10.3390/s18020430.
- 692 13. Kuo, B. C.; Yang, J. M.; Sheu, T. W.; Yang, S. W. Kernel-based KNN and Gaussian classifiers for  
693 hyperspectral image classification. In *International Geoscience and Remote Sensing Symposium (IGARSS)*;  
694 2008; Vol. 2.
- 695 14. Petaccia, G.; Leporati, F.; Torti, E. OpenMP and CUDA simulations of Sella Zerbino Dam break on  
696 unstructured grids. *Comput. Geosci.* **2016**, 20, 1123–1132, doi:10.1007/s10596-016-9580-5.
- 697 15. Florimbi, G.; Torti, E.; Danese, G.; Leporati, F. High Performant Simulations of Cerebellar Golgi Cells  
698 Activity. In *2017 25th Euromicro International Conference on Parallel, Distributed and Network-based*  
699 *Processing (PDP)*; 2017; pp. 527–534.
- 700 16. Wang, Y.; Li, L.; Wang, J.; Tian, R. GPU Acceleration of Smoothed Particle Hydrodynamics for the  
701 Navier-Stokes Equations. *2016 24th Euromicro Int. Conf. Parallel, Distrib. Network-Based Process.* **2016**, 478–  
702 485, doi:10.1109/PDP.2016.28.
- 703 17. Torti, E.; Fontanella, A.; Plaza, A. Parallel real-time virtual dimensionality estimation for hyperspectral  
704 images. *J. Real-Time Image Process.* **2017**, doi:10.1007/s11554-017-0703-6.
- 705 18. Torti, E.; Cividini, C.; Gatti, A.; Danese, G.; Leporati, F.; Fabelo, H.; Ortega, S.; Callicò, G. M. The  
706 HELICoid Project: Parallel SVM for Brain Cancer Classification. *Proc. - 20th Euromicro Conf. Digit. Syst.*  
707 *Des. DSD 2017* **2017**, 445–450, doi:10.1109/DSD.2017.33.
- 708 19. Garcia, V.; Debreuve, E.; Nielsen, F.; Barlaud, M. K-nearest neighbor search: Fast GPU-based  
709 implementations and application to high-dimensional feature matching. *Proc. - Int. Conf. Image Process.*  
710 *ICIP 2010*, 3757–3760, doi:10.1109/ICIP.2010.5654017.

- 711 20. NVIDIA cuBLAS Available online: <https://docs.nvidia.com/cuda/cublas/index.html> (accessed on May 9,  
712 2018).
- 713 21. Chen, G.; Ding, Y.; Shen, X. Sweet KNN: An efficient KNN on GPU through reconciliation between  
714 redundancy removal and regularity. *Proc. - Int. Conf. Data Eng.* **2017**, 621–632,  
715 doi:10.1109/ICDE.2017.116.
- 716 22. Huang, K.; Li, S.; Kang, X.; Fang, L. Spectral–Spatial Hyperspectral Image Classification Based on KNN.  
717 *Sens. Imaging* **2016**, *17*, 1–13, doi:10.1007/s11220-015-0126-z.
- 718 23. Fabelo, H.; Ortega, S.; Ravi, D.; Kiran, B. R.; Sosa, C.; Bulters, D.; Callicó, G. M.; Bulstrode, H.; Szolna,  
719 A.; Piñeiro, J. F.; Kabwama, S.; Madroñal, D.; Lazcano, R.; J-O’Shanahan, A.; Bisshopp, S.; Hernández,  
720 M.; Báez, A.; Yang, G.-Z.; Stanciulescu, B.; Salvador, R.; Juárez, E.; Sarmiento, R. Spatio-spectral  
721 classification of hyperspectral images for brain cancer detection during surgical operations. *PLoS One*  
722 **2018**, *13*, 1–27, doi:10.1371/journal.pone.0193721.
- 723 24. NVIDIA TESLA K40 GPU ACTIVE ACCELERATOR Board Specification Available online:  
724 [https://www.nvidia.com/content/PDF/kepler/Tesla-K40-Active-Board-Spec-BD-06949-001\\_v03.pdf](https://www.nvidia.com/content/PDF/kepler/Tesla-K40-Active-Board-Spec-BD-06949-001_v03.pdf)  
725 (accessed on May 9, 2018).  
726



© 2018 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).