published in Computer Vision and Image Understanding

A minimum barrier distance for multivariate images with applications

Minh Ôn Vũ Ngọc[a], Nicolas Boutry[a], Jonathan Fabrizio[a,**], Thierry Géraud[a]

[a]*EPITA Research and Development Laboratory (LRDE), Le Kremlin-Bicêtre, France*

ABSTRACT

Distance transforms and the saliency maps they induce are widely used in image processing, computer vision, and pattern recognition. The minimum barrier distance (MBD) has proved to provide accurate results in this context. Recently, Géraud *et al.* have presented a fast-to-compute alternative definition of this distance, called the Dahu pseudo-distance. This distance is efficient, powerful, and have many important applications. However, it is restricted to grayscale images. In this article we revisit this pseudo-distance. First, we offer an extension to multivariate image. We call this extension the *vectorial Dahu pseudo-distance*. We provide an efficient way to compute it. This new version is not only able to deal with color images but also multi-spectral and multi-modal ones. Besides, through our benchmarks, we demonstrate how robust and competitive the vectorial Dahu pseudo-distance is, compared to other MB-based distances. This shows that this distance is promising for salient object detection, shortest path finding, and object segmentation. Secondly, we combine the Dahu pseudo-distance with the geodesic distance to take into account spatial information from the image. This combination of distances provides efficient results in many applications such as segmentation of thin elements or path finding in images.

1. Introduction

Over the past decades, distance transforms have been widely used in computer vision, image processing, and pattern recognition (see Wei et al. (2012); Ciesielski et al. (2014); Zhang et al. (2015); Tu et al. (2016); Huang and Zhang (2018)). In general, distances can be classified into two categories: point-wise and path-wise. Point-wise distances are computed relatively to the domain of an image, while path-wise distances take into account the topographical view of the image. In this paper, we focus on path-wise distances, where images can also be seen as graphs (the vertices are the pixels of the image). The usual method to find the path-wise distance between two pixels is to compute the length of the shortest path in the graph that goes from one of these pixels to the other. The most used path-wise distance in image processing is the geodesic distance (see Toivanen (1996)). More recently, a pseudo-distance, called minimum barrier distance (MBD) has been proposed in Strand et al. (2013).

The barrier "strengths" of a path is the difference between the *altitude* of the highest point of the path and the *altitude* of the lowest point of this path. The minimum barrier distance between two points is the smallest barrier "strengths" among the set of all possible paths between these two points. This distance is studied in Ciesielski et al. (2014) and in Strand et al. (2017). The MBD has many interesting theoretical properties and is an effective tool in image processing and computer vision applications, especially to proceed to salient object detection (see Zhang et al. (2015); Tu et al. (2016); Wang and Wang (2017); Wang et al. (2017); Yang et al. (2017); Huang and Zhang (2018)), interactive segmentation (see Grand-Brochier et al. (2014); Malmberg et al. (2017)) and object localization (see Bharati et al. (2016)). Litterature shows that the minimum barrier distance outperforms the geodesic one on noisy and blurred images (see Strand et al. (2013); Zhang et al. (2015)).

The MBD is powerful, but its computation expensive. Several approximations of this distance have then been proposed in Zhang et al. (2015), in Tu et al. (2016) and in Huang and Zhang (2018).

Recently, the Dahu pseudo-distance has been introduced from the point of view of Mathematical Morphology

**Corresponding author: Tel.: +33-153145940
e-mail: jonathan.fabrizio@lrde.epita.fr (Jonathan Fabrizio)

(see Géraud et al. (2017)) in order to approximate the MBD. This Dahu pseudo-distance is computed by considering an image as a landscape (we also speak about its topographical view). Unlike the approach of Zhang et al. (2015) and of Huang and Zhang (2018) which computes the MBD directly in the image space, the Dahu pseudo-distance can efficiently be computed on a tree-based representation of the image; the tree of shapes of Géraud et al. (2013). Thanks to this approach, the computation of the Dahu pseudo-distance is very fast. However, this distance was initially developed for gray-level images and does not handle color images very well. Therefore, we take here into account the color information to improve it.

The main contributions of our paper are the following.

- We provide a method to efficiently compute the Dahu saliency map while constructing the tree of shapes.
- We offer an extension of the Dahu pseudo-distance to multivariate images and we explain how to compute it fast. We call it *vectorial Dahu pseudo-distance*.
- We extend the Dahu pseudo-distance to a more "clever" version which combines the Dahu pseudo-distance computed on the tree and the geodesic distance computed in the image to refine results (especially to find the shortest path between two points in the image space).

To demonstrate the robustness of the Dahu pseudo-distance, we analyze it in several experiments and applications as follows.

- We explore the properties of the Dahu pseudo-distance: we compare our vectorial Dahu pseudo-distance with the Dahu pseudo-distance computed on separate channels, we analyze the noise stability and the contrast of the vectorial Dahu pseudo-distance.
- We demonstrate the robustness of the vectorial Dahu pseudo-distance in some applications such as salient object detection and shortest path finding by comparing it with other MB-based distances; especially for color images.
- We illustrate the usability of our vectorial Dahu pseudo-distance on multi-spectral images by successfully segmenting objects in satellite multi-spectral images.
- We also demonstrate the usability of our vectorial Dahu pseudo-distance on multi-modal images by segmenting white matter regions in the brain on multi-modal medical images.

The paper is organized as follows. Section 2 contains the state-of-the-art related to the MB-based distance. The Dahu pseudo-distance and the way to compute the saliency map are presented in Section 3. In Section 4, we present an efficient way to compute it. Then we provide an extension to multivariate images. Additionally, the Dahu pseudo-distance is improved and a more "clever" version is provided, using at the same time the spatial and hierarchical information of the image. In Section 5, we investigate the properties of the vectorial Dahu pseudo-distance and we compare it with state-of-the-art results. Some applications are presented in Section 6 to demonstrate the efficiency of our distance. The conclusions and perspectives are discussed in Section 7.

2. State-of-the-art

The MBD was originally introduced by Strand et al. (2013) as a minimum value of the barrier strength among the set of possible paths between two pixels in an image. The MBD has been used in several applications in image processing and computer vision, for instance, in salient object detection (see Zhang et al. (2015); Tu et al. (2016); Yang et al. (2017); Wang and Wang (2017); Wang et al. (2017); Huang and Zhang (2018)), in object localization (see Bharati et al. (2016)), in superpixel segmentation (see Hu et al. (2018)), in interactive segmentation (see Kårsnäs et al. (2012); Ciesielski et al. (2014); Grand-Brochier et al. (2014); Malmberg et al. (2017)), refocusing (see Liu et al. (2016)), object proposals generation (see Huang et al. (2018)) and in object segmentation (see Zhang and Shen (2017); Xiao et al. (2018)).

In salient object detection, the goal is to compute a saliency map that highlights the most important objects in an image. To proceed, the *boundary connectivity prior*, which is presented in Wei et al. (2012), assumes that background regions are usually large, homogeneous, and that the image boundary is mostly background. The MBD estimates a distance from every pixel in the image to the image boundary while considering that image boundary is regarded as the background seeds (see Zhang et al. (2015); Tu et al. (2016); Yang et al. (2017); Wang and Wang (2017); Wang et al. (2017); Huang and Zhang (2018)).

Many applications take advantage of the relevance of the saliency map computed by the MBD. The classical usage of this saliency map is object segmentation. For example, in Zhang and Shen (2017), an object segmentation method is proposed by using an affinity model based on the MBD. Object segmentation is also a starting point for multiple other applications. For example, in Bharati et al. (2016), a tracking method based on the MBD is presented. Another example, exposed in Liu et al. (2016), relies on object segmentation to perform a refocusing. Additionally, the relevance of the saliency map computed by the MBD has also been used in object proposal generation as presented in Huang et al. (2018). This method aims to generate a certain amount of candidate bounding boxes to provide potential object locations for further tasks such as object detection and segmentation.

Besides, the MBD has also been used for interactive segmentation (see Kårsnäs et al. (2012); Strand et al. (2013)). In this application, the user tags a small set F of pixels belonging to the object to set it as foreground and a small set B of pixels outside of the object to set it as background. Interactive segmentation is the binary classification of the object with respect to F and B. Each pixel in the image is classified as foreground or background by comparing the MBD between the pixel itself and the two sets of seeds F and B. In Strand et al. (2013), the MBD is

computed on grayscale images, and its extended color version is presented in Kårsnäs et al. (2012). These articles show that this process is robust to noise, blurring and seed point position.

Another application is the computation of superpixels (see in Hu et al. (2018)). The authors propose a method for superpixel segmentation relying on the MBD. Superpixels are determined around them according to "compact-aware MBD", which is a combination of the original MBD and the (spatial) Euclidean distance.

The MBD is very powerful, but difficult to compute efficiently on images of reasonable size. Because computing the exact MBD usually takes too long, approximate but faster methods have been proposed (see Zhang et al. (2015); Tu et al. (2016); Huang and Zhang (2018)).

In Zhang et al. (2015), the authors presented an approximation (Fast-MBD) with a raster scan algorithm to update the MBD thanks to its neighbors. This salient object detection method runs at about 80 FPS and achieves competitive performance with state-of-the-art saliency detection methods. Despite the fact that it provides good results, the raster scan method has difficulties when the exact path between two pixels is in a direction between the bottom left and the top right of the image (see Huang and Zhang (2018) for details).

Tu et al. (2016) have developed another approximation of the MBD. In their approximation, the input image is represented by its minimum spanning tree (MST). Paths between pixels correspond to paths between the nodes of the tree. The MST highly reduces the size of the space we look for to find the shortest path between two pixels of the image. However, the "simple" structure property of MST can lead to some approximation errors, especially when noise appears in the image.

Recently, a new algorithm to approximate the MBD has been presented in Huang and Zhang (2018), which is inspired from the natural phenomena of water flow. The seed pixels which are usually put on the boundary of the image, are assumed to be sources of water. Then, the water spreads from the sources to the neighboring pixels (with different flow MBD) until all the pixels are flooded. The Waterflow-MBD computation achieves a high-speed performance and shows comparable results with other methods.

These methods based on the MBD achieve state-of-the-art results with other bottom-up methods on saliency map computation. They can also process an image in real-time, which is relevant for applications with speed requirements. On the other hand, they also suffer from a number of limitations. In particular, color images (or more generally multivariate images) are not handled very well (or not handled at all). A multivariate version of the MBD needs to keep the advantages of the MBD and has to be efficient as well. For this reason, we propose a new version of the MBD based on the Dahu pseudo-distance. The proposed method in this paper computes distances taking into account multivariate data which can be made of different color images as well as multimodal or multispectral images. We also demonstrate the robustness of the Dahu pseudo-distance in several applications such as salient object detection, shortest path finding and interactive segmentation. The Dahu pseudo-distance, which inherits the properties from the Tree of Shapes (ToS) (see Caselles and Monasse (2009)), has been shown to be robust to noise and blur effects in the image, and it gives competitive results compared to state-of-the-art methods.

3. The Dahu pseudo-distance

In this section, we give the mathematical background necessary to define the MBD in details and we show how to derive a distance map using the MBD, before addressing a new discrete version of the minimum barrier distance, called the Dahu pseudo-distance and an efficient way to compute it.

3.1. The Minimum Barrier Distance

In image processing applications, an image domain is associated with a graph in which vertices represent discrete pixels on the image and edges represent connections between pixels. A gray-level image (Fig. 1(a)) is then represented as a vertex-valued graph (Fig. 1(b)).

A path in a graph X is a sequence $\pi = \langle ..., p_i, p_{i+1} ... \rangle$, with $p_i \in X$ and $p_{i+1} \in \mathcal{N}_X(p_i)$, where $\mathcal{N}_X$ is the adjacency relations between pixels. Also, the set of paths going from the vertex x to the vertex x' is denoted by $\Pi(x, x')$. The *barrier strength* (also called *barrier distance* or *cost*) τ of a path π in the given gray-level image u is defined as:

$$\tau_u(\pi) = \max_{p_i \in \pi} u(p_i) - \min_{p_i \in \pi} u(p_i). \quad (1)$$

The *barrier strength* of a path is the difference between the highest and lowest pixel values along this path. The *minimum barrier distance* d^{MB} between two vertices x and x' in u is then defined as the minimum of the barrier strengths of all the paths between two given vertices:

$$d_u^{\text{MB}}(x, x') = \min_{\pi \in \Pi(x, x')} \tau_u(\pi), \quad (2)$$

In Fig. 1(b), the blue path, which corresponds to a sequence $\langle 1, 0, 0, 0, 2 \rangle$, is considered as the shortest path between these two red vertices. The corresponding MBD is then equal to 2.

Note that, despite its name, the MBD is not a distance, because it can exist some x, y such that $x \neq y$ and $d_u^{\text{MB}}(x, y) = 0$.

3.2. Distance map based on the MBD

It is common to derive a distance map from the MBD. Given a minimum barrier strength function and a set X' of seed points, a distance map S^{MBD} can be computed by:

$$S_u^{\text{MBD}}(x, X') = \min_{x' \in X'} d_u^{\text{MB}}(x, x'). \quad (3)$$

A distance map is then the MBD from every point of the image to the set X' of seed points. For every point, the MBD looks for the smallest distance between x and any pixel x' that belongs to X'.

The next section presents a variant of the MBD, which is also based on the notion of barrier (Eq. (1)).

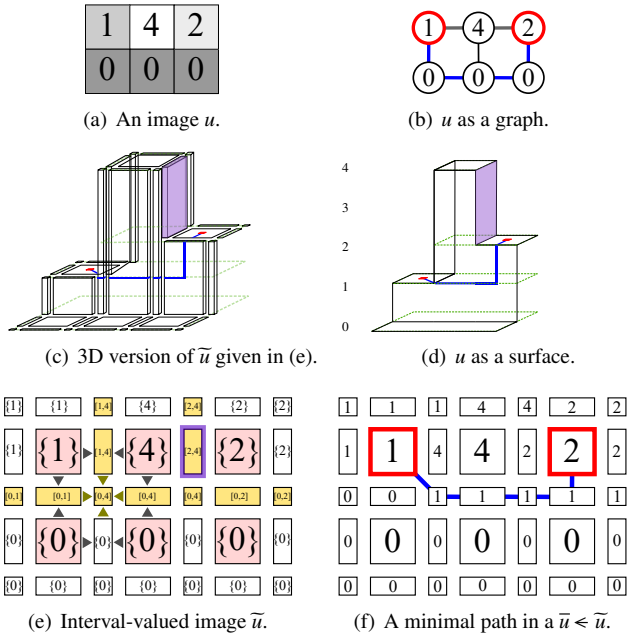

(a) An image u. (b) u as a graph.

(c) 3D version of $\widetilde{u}$ given in (e). (d) u as a surface.

(e) Interval-valued image $\widetilde{u}$. (f) A minimal path in a $\overline{u} \Subset \widetilde{u}$.

Fig. 1. Image representations used to compute barrier distances (see Géraud et al. (2017)).

3.3. The Dahu pseudo-distance

A new discrete version of the MBD, named the Dahu pseudo-distance, has been defined in Géraud et al. (2017). It considers an image (see Fig. 1(a)) as a continuous surface in the set-valued sense (see Fig. 1(d)) on a discrete topological domain called the Khalimsky grid. Details about set-valued continuity and about Khalimsky grids can be found in Kovalevsky (1986) and in Aubin and Frankowska (2009) respectively. The optimal blue path between the two red points is depicted in the image, and has a distance equal to 1. It is slightly different from the original MBD. Let us briefly present this Dahu pseudo-distance.

A gray-level image can be seen as a function $u\colon \mathbb{Z}^2 \to \mathbb{N}$. When we represent an image using a surface, we cannot use scalar functions; we have to use set-valued functions. More exactly, in Géraud et al. (2017), the authors proposed to replace the domain $\mathbb{Z}^2$ by the topological discrete space $\mathbb{H}^2$ of the 2D Khalimsky grid (also known as *2D cubical complex*), and the value domain $\mathbb{N}$ is replaced with the set $\mathbb{I}_{\mathbb{N}}$ of intervals of natural numbers. The 2D cubical complex, which is illustrated in Fig. 1(e) is a set of 2D, 1D, and 0D elements. The 2D elements are the original pixels represented by the big squares in Fig. 1(e). The 1D elements are the rectangles (see Fig. 1(e)) located between the big squares. They are valued by the interval whose minimum and maximum are computed from the two big squares near to it. For example, if the two nearby squares are set at $\{1\}$ and $\{4\}$, the in-between 1D element will be set at $[1, 4]$, and all the level lines between 1 and 4 will cross this 1D element. The 0D elements are the little squares (see Fig. 1(e)) surrounded by 4 squares; they are set at the span value computed from the values of these 4 squares.

Note that the 1D yellow element in Fig. 1(e) which is bounded by a purple border corresponds to the vertical purple part in Fig. 1(c). This 1D element is a way to get a discrete topology and to represent what lies between the pixels.

Thanks to this topology, from a scalar image u, we can construct an interval-valued image $\widetilde{u}$ on the Khalimsky grid, which *really* represents the surface corresponding to u.

Let us introduce the *inclusion relationship*. We say that the real-valued image $\overline{u}$ (a single-valued function) is *included* in the interval-valued image $\widetilde{u}$ when for any element x of the cubical complexes, we have $\overline{u}(x) \in \widetilde{u}(x)$. This inclusion relationship between a scalar image $\overline{u}$ and an interval-valued image $\widetilde{u}$ is denoted by $\Subset$: we write then $\overline{u} \Subset \widetilde{u}$. The Fig. 1(f) depicts an example of a scalar image $\overline{u}$ which is "included" in the interval-valued image $\widetilde{u}$ depicted in Fig. 1(e).

The adaptation of the MBD on the interval-valued image, called the Dahu pseudo-distance (see Géraud et al. (2017)), is noted d^{DAHU}. This Dahu pseudo-distance between two pixels x and x' on the original image u is defined as:

$$d_u^{\text{DAHU}}(x,\, x') = \min_{\overline{u} \Subset \widetilde{u}} d_{\overline{u}}^{\text{MB}}(h_x,\, h_{x'}) \tag{4}$$

$$= \min_{\overline{u} \Subset \widetilde{u}} \min_{\pi \in \Pi(h_x, h_{x'})} \tau_{\overline{u}}(\pi), \tag{5}$$

where h_x and $h_{x'}$ are the 2D elements of the cubical complex corresponding to x and x' respectively. It means that we look for a minimal path in the cubical complex, with the classical definition of the MBD, and consider all the possible scalar functions $\overline{u}$ that are "included" in the interval-valued map $\widetilde{u}$. Returning to the earlier example (Section 3.1, Fig. 1(b)), the shortest path between the two red points in Fig. 1(c), depicted as a blue path in Fig. 1(f) (image $\overline{u}$ is included in the interval-valued image $\widetilde{u}$ that provides the minimal path), has a length of 1. The Dahu pseudo-distance can be interpreted as the *best minimum barrier distance that we can have considering that the input function is continuous in the set valued sense* (see Najman and Géraud (2013)).

Note that, as the MBD, the Dahu pseudo-distance is not a distance, because it can exist some x, y such that $x \neq y$ and $d_u^{\text{DAHU}}(x,\, y) = 0$.

3.4. Efficient Dahu pseudo-distance computation using the tree of shapes

The Dahu pseudo-distance can be computed easily and efficiently thanks to the tree-based representation of the given image. A tree of shapes (see Monasse and Guichard (2000); Caselles and Monasse (2009)) is a morphological self-dual representation of an image. This tree is a decomposition of a gray-level image into connected components, called *shapes*, which can be arranged into a tree encoding an inclusion relationship. A shape is a filled-in connected component without any hole inside (its boundary is then an iso-level line). Two iso-level lines cannot cross each other. A very strong consequence is that shapes are either disjoint or nested, which explains that the tree of shapes is a tree and not a graph with cycles.

The tree of shapes is used to facilitate the computation of the Dahu pseudo-distance. On Fig. 2(a), the path between two points $(x,\, x')$ indicated by red bullets in u is depicted by a blue line, which starts from region B, then goes through A and C, and finally ends in region F. Such a path is minimal because every

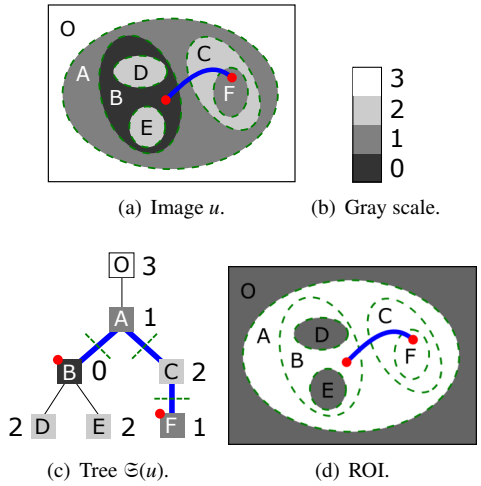

(a) Image u. (b) Gray scale. (c) Tree $\mathfrak{S}(u)$. (d) ROI.

Fig. 2. The tree of shapes of an image allows to easily express and compute the Dahu pseudo-distance and distance maps (see Géraud et al. (2017)).

path in $\Pi(x,\, x')$ should at least cross this same set of level lines to go from x to x'; thus the Dahu pseudo-distance corresponds to the level dynamics of this set of lines. Actually, this path in the image space is exactly *the* (shortest in number of nodes) path in the tree of shapes between the nodes t_x and $t_{x'}$:

$$\dot{\pi}(t_x, t_{x'}) = \langle t_x, \ldots, \mathrm{lca}(t_x, t_{x'}), \ldots, t_{x'} \rangle,$$

where $\mathrm{lca}(t_x, t_{x'})$ is the lowest common ancestor of the pair $(t_x, t_{x'})$ (see the blue path on the tree depicted in Fig. 2(c)). Note that a path in a tree is denoted by $\dot{\pi}$ to distinguish it from paths in the image space.

The Dahu pseudo-distance in the image space between two points x and x' can be written as the minimum barrier distance between the two nodes t_x and $t_{x'}$ representing the components in the tree of shape containing respectively x and x':

$$d_u^{\mathrm{DAHU}}(x,\, x') \;=\; d^{\mathrm{MB}}_{\mathfrak{S}(u)}(t_x,\, t_{x'}) \tag{6}$$

$$\;=\; \max_{t \in \dot{\pi}(t_x, t_{x'})} \mu_u(t) \;-\; \min_{t \in \dot{\pi}(t_x, t_{x'})} \mu_u(t), \tag{7}$$

where $\mu_u(t)$ denotes the gray-level associated with the node t of the tree of shapes $\mathfrak{S}(u)$ of the image u. For instance, in Fig. 2(c), the blue path gives the sequence of node values $\langle 0, 1, 2, 1 \rangle$, so the Dahu pseudo-distance is 2. There is *no need* to find the best scalar image $\overline{u} \lessdot \widetilde{u}$, nor to find the best path $\pi \in \Pi(x,\, x')$ in the image space; it thus means that the seminal definition of the Dahu pseudo-distance (see Eq. (5)) is the best choice to be fast in time. The new expression of this distance (see Eq. (7)) is just a barrier strength computation (such as Eq. (1)) on the trivial path $\dot{\pi}(t_x,\, t_{x'})$ of nodes in the space of the tree of shapes.

3.5. Saliency map based on the Dahu pseudo-distance

A distance map function of an image u can be derived from the MBD as we have seen in Eq. (3). Indeed, we can define the distance map based on the Dahu pseudo-distance as follows:

$$S_u^{\mathrm{DAHU}}(x, X') \;=\; \min_{x' \in X'} d_u^{\mathrm{DAHU}}(x,\, x'),$$

where X' is some set of points of the domain of the image u.

Now, let us define the corresponding set of nodes on $\mathfrak{S}(u)$ of X':

$$T_{X'} \;=\; \{\, t_{x'} \,;\, x' \in X' \,\}. \tag{8}$$

Then, we obtain using Eq. 6 and then Eq. 3 that:

$$S_u^{\mathrm{DAHU}}(x, X') \;=\; S^{\mathrm{MBD}}_{\mathfrak{S}(u)}(t_x, T_{X'}), \tag{9}$$

which shows how the distance map induced by the Dahu pseudo-distance is related to the distance map induced by the MBD. As a consequence, a Dahu distance map is the Dahu pseudo-distance from every node in the tree to the set $T_{X'}$ of seed nodes.

4. Going further with the Dahu pseudo-distance

The Dahu pseudo-distance, which inherits its properties from the tree of shapes, is shown to be efficient for some applications (see Géraud et al. (2017)). For this reason, we increase its computation speed and propose an extension to color and multivariate images. We also propose an improvement of it using a two-steps procedure taking into account the domains of the tree of shape and of the initial image. This last measure is related to the topographical representation of the image.

4.1. Simultaneous computations of the Dahu pseudo-distance and the tree of shapes

In natural images, the border of the image is mostly background (see Wei et al. (2012)). Similar to previous works (see Zhang et al. (2015); Tu et al. (2016); Huang and Zhang (2018)), we compute the distance map, which is the Dahu pseudo-distance of every pixel in the image to the border of the image. In particular, the Dahu pseudo-distance can be computed while constructing the tree of shape. The construction of the tree of shapes is mentioned in Géraud et al. (2013). Our algorithm (see Algo. 1) is a modification of the sorting procedure used to compute the tree of shapes: we add some operations (see the blue lines) to the pixel sorting procedure during the tree construction.

Our algorithm computes the Dahu pseudo distance from seed points (the border of the image) to every other point in the domain of the image. The process follows two steps. During the first step (lines 2 - 18), it crosses all points in the domain using a propagation front. Every pixel is crossed only once (thanks to *deja_vu* variable). This propagation front is managed by a hierarchical queue (q). Then, the algorithm computes two structures *min_im* and *max_im*; *min_im* and *max_im* represent the lower and higher levels arisen during the propagation respectively. In the second step (lines 19 - 20), the Dahu pseudo distance is computed from the two structures *min_im* and *max_im*. All points are crossed (whatever is the order) and the Dahu distance is simply the difference between *max_im* and *min_im* at the considered point.

Our algorithm can be explained thoroughly as follows. Initially, we add an artificial border surrounding the image domain with the unique value l_∞. p_∞ is one point from the border. Only

Algorithm 1: Modification of the sorting procedure of the tree of shapes to compute the Dahu pseudo-distance.

Data: Interval-valued image U, Image domain D, Seed Point p_∞
Result: Dahu pseudo-distance

```
1  begin
       /* q, a priority queue                    */
       /* l, the current level                   */
       /* N(p), the set of neighbors of p        */
2      for all h ∈ D do
3          deja_vu(h) ⟵ false
4      PUSH(q[l_∞], p_∞);
5      deja_vu(p_∞) ⟵ true;
6      l ⟵ l_∞;
7      Image2d min_im, max_im, Dahu;
8      min_im (p_∞) ⟵ l ;
9      max_im (p_∞) ⟵ l ;
10     Dahu (p_∞, p_∞) ⟵ 0 ;
12     while q is not empty do
13         p ⟵ PRIORITY_POP (q, l);
14         for all n ∈ N(p) such as deja_vu(n)== false do
15             l' ⟵ PRIORITY_PUSH(q, n, U, l) ;
16             deja_vu(n) ⟵ true ;
17             min_im (n) ⟵ min(min_im (p),l') ;
18             max_im (n) ⟵ max(max_im (p),l') ;
19     for all p ∈ D do
20         Dahu (p_∞, p) ⟵ max_im (p) - min_im (p) ;
21     return(Dahu)
```

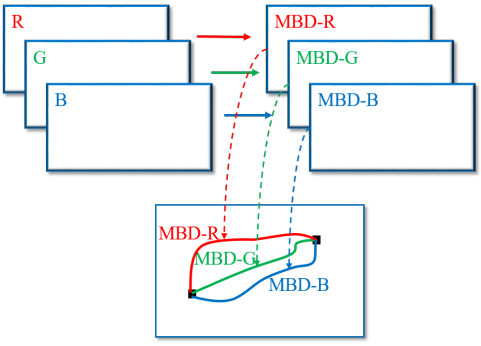

(a) A procedure to compute the MBD and their shortest paths in the color image when processing separately each channel.

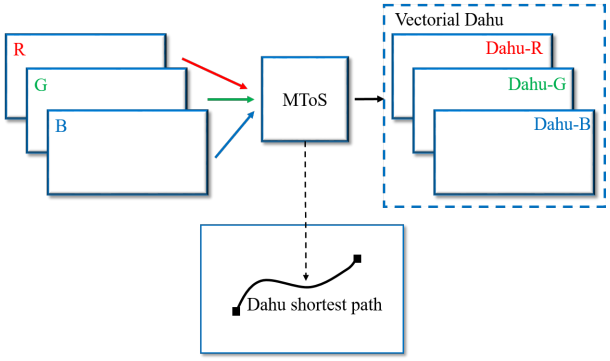

(b) A procedure able to compute the vectorial Dahu pseudo-distance. Even with color images, our method is able to obtain a coherent shortest path between two pixels in the image.

Fig. 3. The computation of the MBD and of the vectorial Dahu pseudo-distance in a color image. Contrary to the MBD computed on color images, which may find a different path in the image for each channel, the Dahu pseudo-distance finds a same path in the image minimizing the sum of the barriers in all channels simultaneously.

one step remains to be able to proceed to the front propagation: we must input the set-valued map U computed thanks to a span-based interpolation on the image u. Then, we call the sorting procedure described in Géraud et al. (2013), which is employed by using a hierarchical queue q; the current level is denoted by l. The Dahu pseudo-distance of the starting point is set at the value 0. Since we use interval-valued maps, we have to decide at which level to enqueue those elements. The face p is enqueued at the value of the interval $U(p)$ which is the closest to l, denoted l' (see the procedure PRIORITY_PUSH). The value l' is compared with the minimum and maximum values of its neighbors to update the Dahu pseudo-distance. When the queue $q(l)$ at the current level is empty, the procedure PRIORITY_POP decides whether the next level to be processed is less or greater than l. This loop continues until all of the pixels have been visited. The resulting pseudo-distance is then obtained. More information about the PRIORITY_PUSH and PRIORITY_POP procedures can be found in Géraud et al. (2013). Note also that to finally obtain the tree of shapes, three procedures must be executed (see Algo. 3 in Géraud et al. (2013)), but we will not go into any further detail since this is not the subject of our article.

When the seed pixels are not placed in the outer border of the image (for example, if they are placed at the center of the image), we need to build the tree of shapes first, and then we can compute the Dahu pseudo-distance. The major difference with a classical saliency map, defined in the image space (such as the one of Eq. (3)), is that the tree structure is one-dimensional. Since the Dahu pseudo-distance on the tree (given by Eq. (7)) has the form of a barrier "max - min", the saliency map $S_{\Xi(u)}^{\mathrm{MBD}}$ at each node t_x can be easily computed by a propagation method on the tree using a priority queue. Afterwards, getting the 2D saliency map S_u^{DAHU} means reading for each x the value of $S_{\Xi(u)}^{\mathrm{MBD}}$ at t_x. Eventually, once we have computed the tree of shapes $\Xi(u)$, the computation of a saliency map $x \mapsto S_u^{\mathrm{DAHU}}(x, X')$ is immediate (whatever the set X').

Last, let us mention that the representation of an image into a tree of connected components is easy to handle (see Carlinet and Géraud (2014)). Furthermore, the tree of shapes of an image can be computed in quasi-linear time w.r.t. the number of pixels (see Géraud et al. (2013)), and can be parallelized (see Crozet and Géraud (2014)).

4.2. Extending the Dahu pseudo-distance to multivariate images

As mentioned before, the previous MBD methods (see Zhang et al. (2015); Tu et al. (2016)) are only defined on grayscale images or on separate channels of color images. In this last case, they compute the mean or the maximal value of the distances obtained on each separate channel (see Tu et al. (2016) for details). This approach is not satisfying for the purpose of image segmentation: we generally obtain different paths for each color, and then computing the mean or the max value of the distances makes no sense and cannot be used for image segmentation. An example of the computation of the MBD is illustrated in Fig. 3(a).

In Kårsnäs et al. (2012), a vectorial minimum barrier distance (VMBD) is proposed to compute the MBD on a multivariate image. However, this VMBD is not easy to compute directly on the image. Moreover, the VMBD is not effective when computing multiple distances between multiple points in images. To solve this problem, in this section, we present a Dahu pseudo-distance extended to multivariate images based on the tree space. In Ôn Vũ Ngọc et al. (2018), this color Dahu pseudo-distance is proposed to detect automatically documents in images.

The tree of shapes, primarily defined on gray-level images, has been recently extended to multivariate data (see Carlinet and Géraud (2015)); this extension is called the *Multivariate Tree of Shapes* (MToS). It yields a tree mapping the inclusion relationship of shapes in the image. Such a representation is of prime importance for computer vision (see Cao et al. (2008)) because it satisfies strong invariance properties featured by natural images, such as self-duality and local contrast changes (see Caselles et al. (1999)).

However, the definition of the Dahu pseudo-distance on the tree of shapes (see Eq. (7)) cannot be used without modification/improvement. In the work of Kårsnäs et al. (2012), four different path costs (linear and non-linear) have been presented: the diameter, maximum diameter, city-block diameter and volume of the bounding box. Using their conclusion and also thanks to our experiments, we chose here to employ the city-block diameter to compute the distance. The choice of the path cost function is debatable but in practice it has a very low impact. Changing the underlying distance changes the magnitude of the result. As long as the underlying choice does not change the order of pixels, all applications relying on the Dahu pseudo-distance will mostly not be impacted.

Let us now consider that $\boldsymbol{u}$ is a multivariate image, t is a node of the MToS of $\boldsymbol{u}$, and $\mu_{\boldsymbol{u}}(t)$ is the vector value associated with the node t. The superscript i indicates which one of the N components of the vector is taken into account. We can then extend the Dahu pseudo-distance like this:

$$d_{\boldsymbol{u}}^{\text{DAHU}}(x,\, x') \;=\; \textstyle\sum_{i\in\{1..N\}} \alpha_i \; \tau_{\boldsymbol{u}}^{(i)}(\,\dot{\pi}(t_x,\, t_{x'})\,). \tag{10}$$

with:

$$\tau_{\boldsymbol{u}}^{(i)}(\dot{\pi}) \;=\; \max_{t\in\dot{\pi}} \mu_{\boldsymbol{u}}^{(i)}(t) \;-\; \min_{t\in\dot{\pi}} \mu_{\boldsymbol{u}}^{(i)}(t), \tag{11}$$

where α_i is the coefficient weighting each channel, thereby representing the importance of the channel.

Algorithm 2: Computation the Dahu pseudo-distance between two pixels in the image.

Data: Image U, Image domain D, Point x, x'
Result: Dahu pseudo-distance

Compute($MToS(u)$);
Compute(t_x), *Compute*($t_{x'}$);
Compute($\text{lca}(t_x, t_{x'})$);
Compute($\dot{\pi}(t_x, t_{x'})$);
for $i \in [1, N]$ **do**
  Compute($\min_{t\in\dot{\pi}(t_x,t_{x'})} \mu_{\boldsymbol{u}}^{(i)}(t)$);
  Compute($\max_{t\in\dot{\pi}(t_x,t_{x'})} \mu_{\boldsymbol{u}}^{(i)}(t)$);
  Compute($\tau_{\boldsymbol{u}}^{(i)}(\dot{\pi}(t_x, t_{x'}))$)(Eq. (11));
end
Compute($d_{\boldsymbol{u}}^{\text{DAHU}}(x,\, x')$)(Eq. (10));
return($d_{\boldsymbol{u}}^{\text{DAHU}}(x,\, x')$)

The vectorial Dahu pseudo-distance between two points x and x' in the domain of the image u can be computed using Algo. 2. After the computation of the tree of shapes, we find the nodes t_x and $t_{x'}$ which correspond to two points x and x' respectively. Then the shortest path $\dot{\pi}(t_x, t_{x'})$ between these two nodes is computed. Therefore, we are able to compute the Dahu pseudo-distance on each channel (see Eq. (11)) and sum up to get the vectorial Dahu pseudo-distance (see Eq. (10)). Please be advised that, the MToS is computed from the ToS of each image channel by merging some marginal shapes. Due to its tree properties, it is not a complete representation of an image. The node of the final tree is associated with multiple values of the image. Therefore, a node has to be assigned to a single value computed from the set of values it contains. In our case, we set each node in the MToS using the median value of its pixels. As a result, the vectorial Dahu pseudo-distance computed on the color image is an approximation of the distance between two points in the image. The whole process to compute the vectorial Dahu distance is illustrated in Fig. 3(b). This way, we obtain a "coherent" shortest path between two pixels in the image (see Fig. 3(b)). As a consequence, we also solve the problem of the different paths of the previous MBD methods that we mentioned at the beginning of this section.

Relying on the presentation of the vectorial Dahu pseudo-distance on multivariate images in the previous paragraphs, we apply it here on RGB color images. To be rigorous, the coefficient, which is the gamma correction in this case should be applied to get linear ranges. Obviously, for many color spaces (like H.L.S.), these coefficients are not valid. Instead of looking for correct coefficients, it is always possible to convert color in

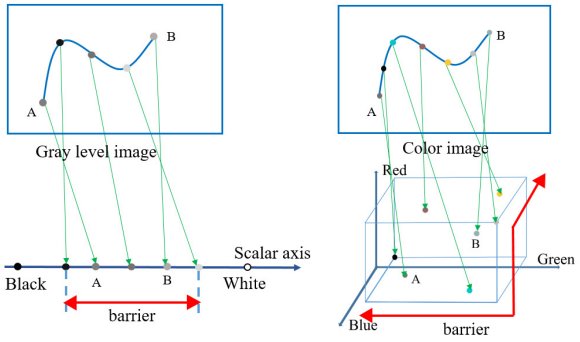

(a) Dahu in the grayscale image. (b) Dahu in the color image.

Fig. 4. The Dahu pseudo-distance in the grayscale image and in the color image.

a color space where each channel is *comparable*. Since the importance of each channel is considered equally, we propose to fix:

$$\alpha_i = 1/N. \tag{12}$$

Then, for RGB-color images, our equation becomes:

$$d_{\boldsymbol{u}}^{\text{DAHU}}(x,\, x') \;=\; \frac{1}{3}\sum_{i\in\{R,G,B\}} \tau_{\boldsymbol{u}}^{(i)}(\dot{\pi}(t_x,\, t_{x'})). \tag{13}$$

Please note that, although Eq. (13) looks simple, we have here a strong result. The Dahu pseudo-distance is one of the *optimal* paths between two points in the image space; this path is such that the set of colors on the path has the smallest 3D bounding box in the color space. This is a highly combinatorial problem which cannot be solved efficiently in the image space. Our contribution here is to turn this problem into an efficient and straightforward computation in a tree space. The Dahu pseudo-distance in the gray-scale image is illustrated in Fig. 4(a), the Dahu pseudo-distance on the color image is illustrated in Fig. 4(b) as the size of the bounding box or the length of the red line.

As presented in Eq. (13), the input of the process is a multivariate image when the output is a (scalar) distance. However, in Fig. 3(b), the output of the process can also be a multivariate image (one distance map by channel). In the experimental section, we will show some examples of what we call abusively "vectorial distance maps". Note that we do not use the vectorial distance map for an evaluation purpose but for visualization only. It is actually a multivariate image, which is computed from a multivariate input based on the vectorial Dahu pseudo-distance. To avoid ambiguities, we will refer in the sequel to *viso* for *vectorial-input-scalar-output*, to *vivo* for *vectorial-input-vectorial-output*, and to *siso* for *scalar-input-scalar-output* Dahu pseudo-distances.

Additionally, our vectorial Dahu pseudo-distance is not restricted to 3 channels and is fully usable on any kind of multichannel images because it relies on the MToS. It means that, we are able, without any additional effort, to compute our vectorial Dahu pseudo-distance on multi-modal images or hyper-spectral images according to Eq. (10). The coefficient on each channel in this equation has to be revisited. A simple idea to refine these weights, is to compute a Principal Component Analysis (abbreviated as P.C.A.) and use eigenvalues to weight the sum.

This extension is a major one, as many existing algorithms, previously restricted to grayscale images, can now be applied on color, multi-spectral, or even hyper-spectral images at low cost. Most of the time, these algorithms work as-is by simply changing the underlying distance (substituting the classical MBD by our vectorial Dahu pseudo-distance). We will illustrate this further, on satellite multi-spectral images, and even on medical multimodal images.

4.3. Extending the Dahu pseudo-distance with spatial information

Our vectorial Dahu pseudo-distance is only defined in the tree space, not in the image space. It is ambiguous for us to visualize this distance on the image. Additionally, in the previous section, our distance is proved to solve the problem of the different paths of the previous MBD methods, but we have not discussed the way to find this coherent path in the image space. Therefore, in this section, we describe our method for computing the Dahu shortest path. This proposed improvement of the Dahu pseudo-distance is used in competition with the commonly used geodesic distance.

The goal of MBD computation is to find optimal path connecting seed pixels and every other pixel (Huang and Zhang, 2018). In case of the Waterflow-MBD method (Huang and Zhang, 2018), the parenthood relation between two neighbor pixels is recorded during the propagation process. The shortest path problem is simply tracking back the relation from the destination pixel until the seed pixel. On the other hand, the MST-MBD finds the candidate path from seed pixel to the others relying on the Minimum spanning tree. This tree largely reduces the search space of the shortest path. However, this simple structure is sensitive to noise and blur, thereby leading to some important deviation from the shortest path.

We present here an extension of the Dahu pseudo-distance by taking into account the spatial information between two pixels in the image. In other words, it is a combination between the Dahu pseudo-distance computed on the tree and the geodesic distance computed in the image restricted to all paths minimizing the Dahu pseudo-distance. This improvement is a "two-steps" procedure, illustrated in Fig. 5, in which we look for the minimal path between the two given pixels x and x' (the two red points in Fig. 5 on the left) and we find the red path in Fig. 5 on the right.

In the first step, we denote $par(t_x)$ as the parent node of node t_x in the tree, and $lca(t_x, t_{x'})$ as the lowest common ancestor of the nodes t_x and $t_{x'}$. The shortest path $\dot{\pi}(t_x, t_{x'})$ between two nodes t_x and $t_{x'}$ is the sequence of nodes that begins at node t_x, goes through the lowest common ancestor $lca(t_x, t_{x'})$, and ends

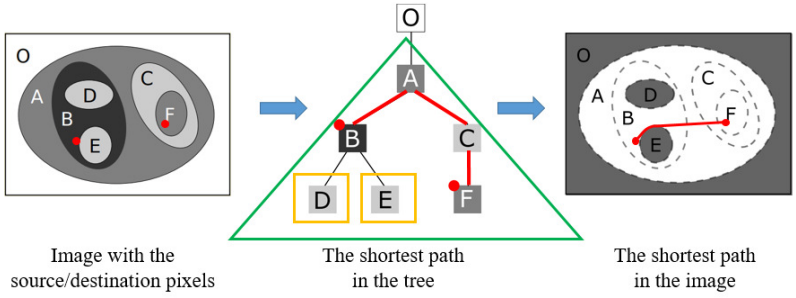

Fig. 5. Extension of the Dahu pseudo-distance: the Dahu pseudo-distance is combined with the geodesic distance.

at node $t_{x'}$. When we have $t_x \neq t_{x'}$, the shortest path $\dot{\pi}(t_x, t_{x'})$ can be formulated as follows:

$$\langle t_x, par(t_x), \ldots, lca(t_x, t_{x'}), ..., par(t_{x'}), t_{x'}, \rangle \qquad (14)$$

otherwise it is the trivial path $\langle t_x \rangle$. This shortest path $\dot{\pi}(t_x, t_{x'})$ is illustrated in red in Fig. 5 in the middle.

The shortest path $\dot{\pi}(t_x, t_{x'})$ in the tree corresponds to a region on the image. We call this region the *shortest path region*. In Fig. 5, the shortest path between nodes B and F is illustrated as the red path. This path goes through regions B, A, C and F. It does not traverse regions O, D and E. Therefore, the *shortest path region* in this case is the white region in the image (on the right in Fig. 5). The *shortest path region* is actually the set of all the possible paths between the two given points in the image space minimizing the Dahu pseudo-distance. This region is connected according to the properties of the connected component trees. Therefore, it ensures to generate a coherent path between the two given pixels in the multivariate image. As a consequence, this extended Dahu pseudo-distance solves the problem that we presented at the beginning of Section 4.2, in which the MBD is computed separately on each channel (but it does not provide a unique path in the image domain).

In the second step, we consider here the spatial information between two points in the image. We want to find a path between the two given pixels x and x', which belongs to the *shortest path region*, so that it has the shortest length in the image space (or more precisely, the geodesic distance in the *shortest path region*). The optimal path is depicted in Fig. 5 as the red line. This path is the shortest path in the sense of the Dahu pseudo-distance between two given pixels x and x'. The shortest path is found in this region by using the heuristic A^* algorithm (see Hart et al. (1968)). This algorithm is a popular technique used in path-finding and graph traversals, especially in games and web-based maps. It is based on the movement cost to move from the seed pixel to a given pixel, and the estimated movement cost to move from that given pixel on the image to the destination. This optimal path has different meanings. It is not only the shortest path in the "color space" but also the shortest path in the image space.

This computation would not have been possible with the color MBD. As seen in Fig. 3, the color MBD may provide different paths on the different channels. On the contrary, the Dahu pseudo-distance makes this combination possible because it provides a unique path in the image, regardless of its number of channels.

This property of the Dahu pseudo-distance has applications related to the shortest path, as will be illustrated with several experiments in the next section.

5. Experimental Results

In this section, we explore the properties of the vectorial Dahu pseudo-distance via some experiments related to visual saliency detection, to noise stability and to the contrast of the Dahu pseudo-distance. Finally, we provide a comparison between the complexities (in time) of the Dahu pseudo-distance vs. some other MB-based distances.

5.1. Visual saliency detection

To show the robustness of the vectorial Dahu pseudo-distance, we start with visual saliency detection applications (see Zhang et al. (2015); Tu et al. (2016); Huang and Zhang (2018)). We remind that visual saliency detection has been widely used in computer vision to obtain visual attention areas in the image.

First, we compare the vectorial Dahu pseudo-distance with the Dahu pseudo-distance on separate channels. Then, we compare the vectorial Dahu pseudo-distance with state-of-the-art MB-based distances.

Datasets. To perform this evaluation, we use the following four large benchmark datasets.

1. MSRA-10K (see Cheng et al. (2015)), which contains 10000 images with pixel accurate salient object labeling for each image.

2. DUTOMRON (see Zhang et al. (2017)), which consists in 5166 challenging images, each of which has one or more salient objects and complex background.

3. ECSSD (see Shi et al. (2016)), which contains 1000 images along with pixel-wise ground truth masks, and includes more salient objects under complex scenes.

4. PASCAL-S (see Li et al. (2014)), which contains 850 images and 1296 object instances. This one is designed to eliminate the center bias and color contrast bias.

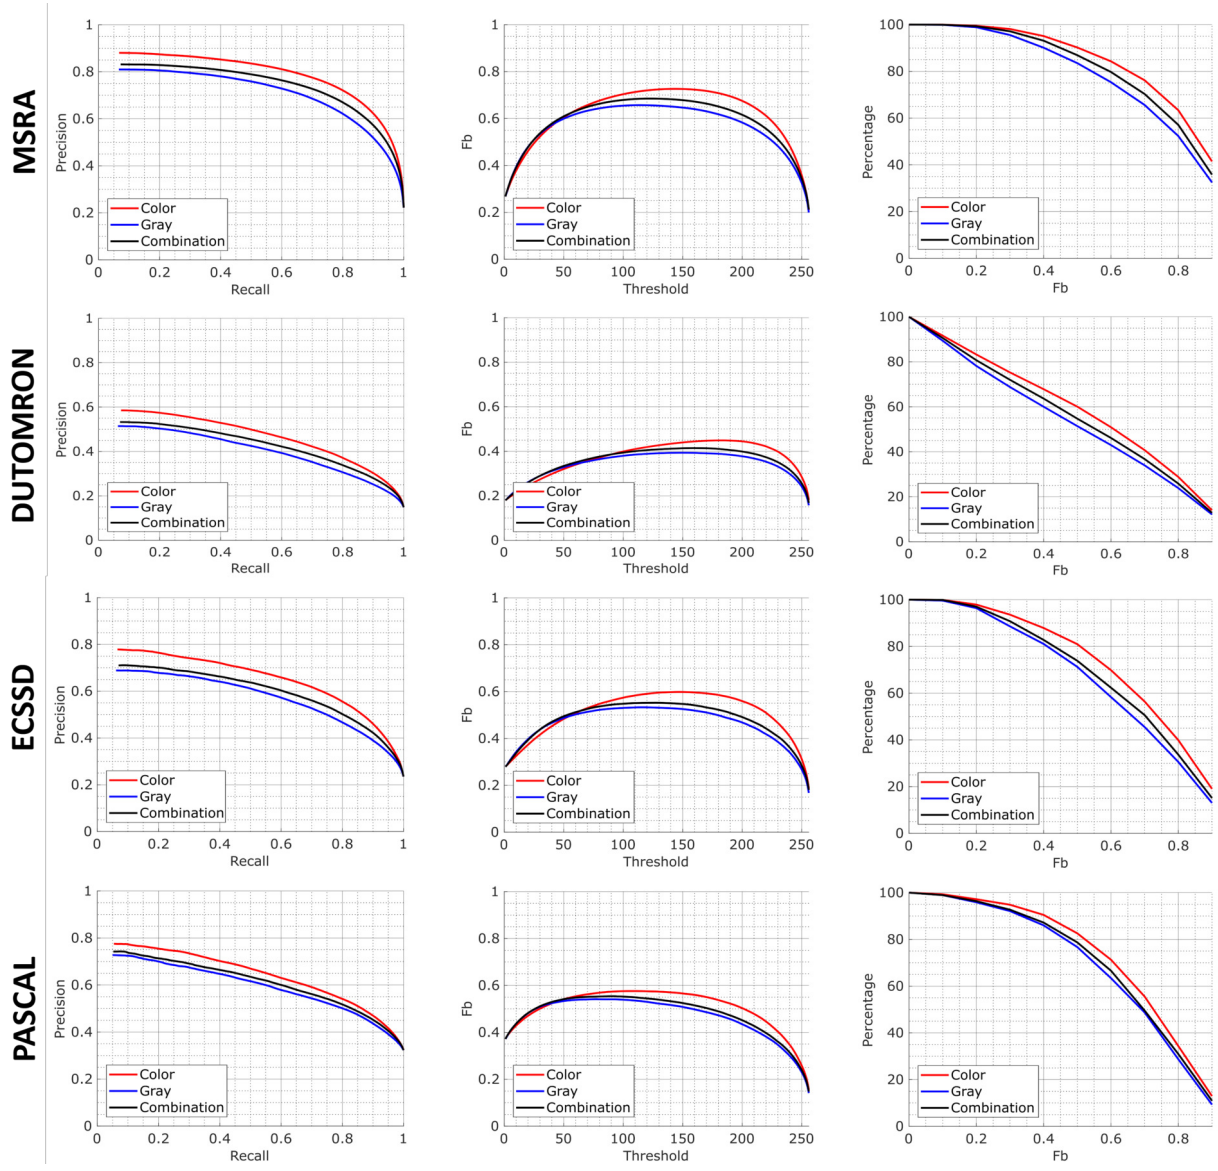

(a) PR (b) F_β-measure (c) Percentage

Fig. 6. Comparison between saliency maps obtained using the vectorial Dahu pseudo-distance and using the Dahu pseudo-distance on separate channels. From top to bottom: the four datasets (MSRA-10K, DUTOMRON, ECSSD, PASCAL-S). From left to right: the three evaluation metrics: (a) Precision-recall curves, (b) F_β-measure, (c) Percentage curves. "Color" is the *color* saliency map computed using our vectorial Dahu pseudo-distance applied directly on color image, "Gray" is the saliency map obtained using the Dahu pseudo-distance computed on the grayscale image and "Combination" is the saliency map obtained by averaging saliency maps computed on separate red, green and blue channels. The three different measures show that our vectorial Dahu pseudo-distance leads to a much better saliency map.

(a) ECSSD

Method	MAE	F_β^{max}	EMD
Color	**0.21**	**0.69**	**0.29**
Gray	0.22	0.6	0.33
R	0.22	0.62	0.34
G	0.22	0.6	0.33
B	0.23	0.62	0.35
Combination	0.22	0.62	0.33

(b) DUTOMRON

Method	MAE	F_β^{max}	EMD
Color	**0.17**	**0.57**	**0.41**
Gray	0.18	0.50	0.43
R	0.18	0.52	0.45
G	0.18	0.50	0.43
B	0.19	0.52	0.45
Combination	0.18	0.52	0.43

(c) PASCAL

Method	MAE	F_β^{max}	EMD
Color	**0.22**	**0.69**	**0.28**
Gray	0.24	0.63	0.3
R	0.23	0.65	0.31
G	0.23	0.64	0.3
B	0.24	0.65	0.31
Combination	0.23	0.65	0.3

(d) MSRA

Method	MAE	F_β^{max}	EMD
Color	**0.16**	**0.79**	**0.17**
Gray	0.19	0.72	0.21
R	0.18	0.75	0.22
G	0.18	0.73	0.21
B	0.18	0.74	0.21
Combination	0.18	0.75	0.23

Table 1. Comparison between saliency maps obtained using the vectorial Dahu pseudo-distance and using the Dahu pseudo-distance on separate channels using MAE, F_β^{max} measure and EMD score. "Color" is the *color* saliency map computed using our vectorial Dahu pseudo-distance applied directly on color image, "Gray" is the saliency map deduced from the Dahu pseudo-distance computed on the grayscale image, R, G and B are the saliency maps deduced from the Dahu pseudo-distance computed on each channel separately and "Combination" is the saliency map obtained by averaging the three saliency maps R, G and B. The best result is highlighted in bold and the worst is underlined. The three different measures show that our vectorial Dahu pseudo-distance leads to a much better saliency map.

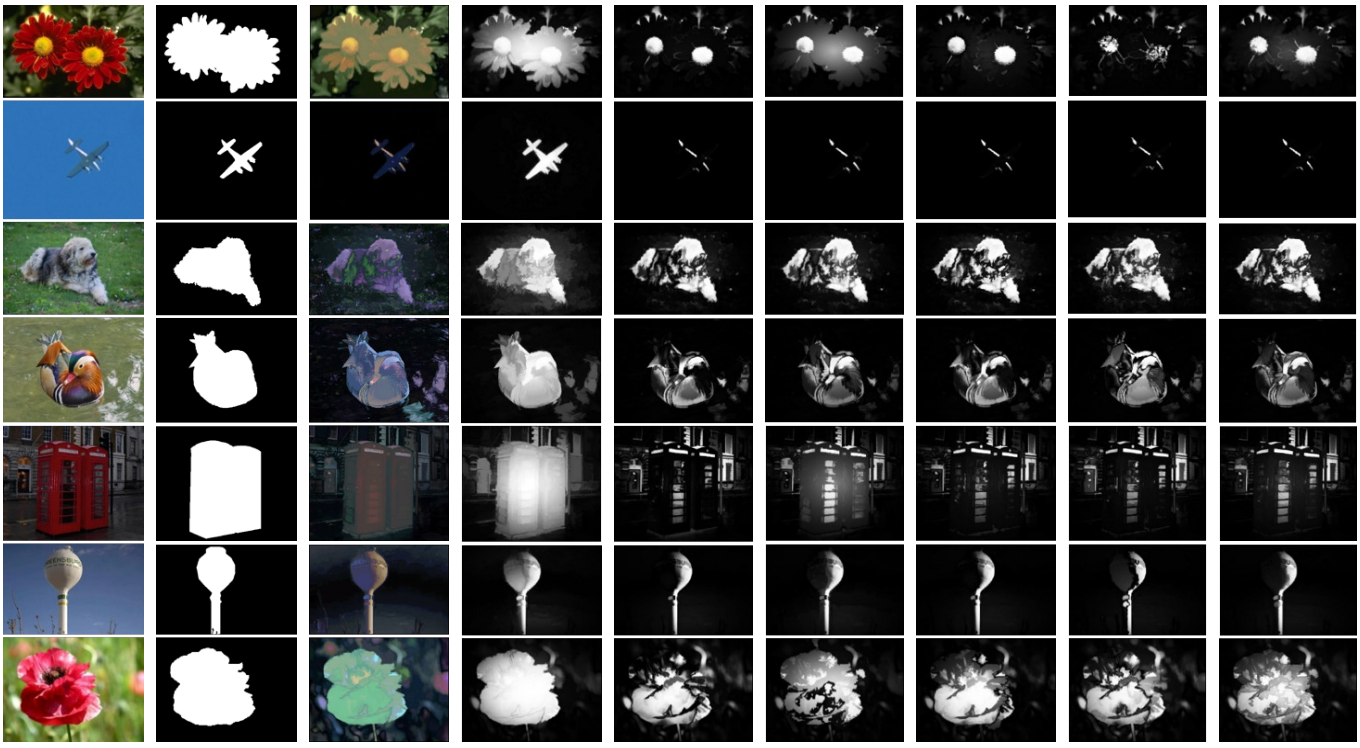

(a) Input (b) GT (c) *vivo* (d) *viso* (e) Gray (f) Red (g) Green (h) Blue (i) Combination

Fig. 7. Several saliency maps of the vectorial Dahu pseudo-distance on color images and the Dahu pseudo-distance on separate channels. Note that image (c) and (d) are respectively the *vivo* and *viso* Dahu pseudo-distances on the color image. The Dahu pseudo-distance on the color image highlights the object over the background, whereas, when only one channel is used, the saliency map only spots a part of the object.

Among these datasets, the DUTOMRON dataset is the most challenging.

Evaluation metrics. We use the following measures.

- The Precision-Recall (PR) curve is used to evaluate the overall performance of a method concerning its trade-off between the precision and recall rates.
- The Mean Absolute Error (MAE), which is the average difference between a saliency map S (gray-level image) and a ground-truth image GT (binary image):

$$\text{MAE} = \frac{\sum_{x\in\mathcal{D}} |GT(x) - S(x)|}{|\mathcal{D}|}, \tag{15}$$

  with $\mathcal{D}$ the domain of the initial image.
- An F_β-measure defined by:

$$F_\beta = (1+\beta^2) \times P \times R \,/\, (\beta^2 \times P + R), \tag{16}$$

  where P and R are respectively the precision and the recall which we mentioned above. We set $\beta^2 = 0.3$ (because it is the classical setting in the visual saliency community).
- The percentage curve, which shows the number of images in the dataset having a F_β score over a specific value. To compute it, we threshold the saliency map at each value between 0 and 255, and we choose the "best" threshold set, that is, the one that gives the highest F_β score (we call this score ${F_\beta}^{max}$). After its computation for each image in the dataset, we compute the corresponding histogram (we choose a number of bins equal to 10), and we finally obtain the percentage curve.
- A score (briefly called EMD) inspired from Calarasanu et al. (2015) relying on the Earth Mover's Distance, which is the cross-bin distance function. It is used as a measure to estimate the dissimilarity between two signatures. In our case, the EMD is computed as the cost between the histogram of F_β score and the histogram of the ground truth image, which is equivalent to one bin at the value $F_\beta = 1$.

5.1.1. Comparison of saliency maps obtained by the usual Dahu pseudo-distance on separate channels and by our vectorial ("color") Dahu pseudo-distance

Experimental setting. We compare our ("*color*") Dahu saliency map (the extension of the Dahu pseudo-distance on the color images which are mentioned in Section 4.2) with the Dahu saliency map computed on separate channels (gray, red, green, blue) and a simple combination of saliency maps computed on each three color channels (pixel-wise average of the three channels).

Initially, input images are resized proportionally so that the maximum dimension is 300 pixels. Then to use the Dahu

pseudo-distance in visual saliency detection, we adopt two priors about the background in natural images, namely *boundary* and *connectivity priors*, which are proposed in Wei et al. (2012). A border with the median value of all of the pixels on the boundary of the image is added to the image. We consider all the pixels in the added border of the image as seed pixels. For the post-processing step, we used the same method as presented in Zhang et al. (2015) to "normalize" the resulting saliency maps.

Evaluation using PR curves. In Fig. 6, we show the PR curves for the saliency maps: directly computed on color images, computed on grayscale images, a pixel-wise combination saliency map of the three channels (as presented in Zhang et al. (2015)). The vectorial Dahu pseudo-distance outperforms the Dahu pseudo-distances on grayscale images and the combination of three channels in all four datasets. On the most challenging dataset (DUTOMRON), the performance of the distance maps deduced from Dahu pseudo-distance are lower than the performance of the one on other datasets. Note that in this dataset, there are multiple objects in images and the color contrasts between the foreground and the background are low.

Evaluation using MAE. The MAE scores of compared methods are shown in Table 1. Note that the lower the MAE is, the better the performance of the method is. The comparison of the saliency maps shows that the Dahu pseudo-distance does not give a better score on the grayscale images compared to the separate channels (R/G/B) while the pixel-wise combination saliency map does improve. This comparison shows also that the vectorial Dahu pseudo-distance achieves better scores than all other methods.

Evaluation using F_β-measure. We adopt the F_β-measure proposed in Margolin et al. (2014) to evaluate saliency maps. In Fig. 6 and in Table 1, the vectorial Dahu pseudo-distance achieves significantly better scores than the Dahu pseudo-distance on grayscale images, and than the combination of the Dahu pseudo-distance across all datasets. We also notice that the F_β-measure curves of the Dahu pseudo-distance have stable and flat curves, which is an advantage because the "best" threshold remains unknown and can vary a lot from an image to another.

Evaluation using percentage curves and EMD. In Fig. 6, the vectorial Dahu pseudo-distance provides better percentage curves than the others. Notably in the MSRA-10K and ECSSD dataset, the number of good saliency maps (F_β-measure > 0.8) of the vectorial Dahu pseudo-distance is higher by around 7% than the Dahu pseudo-distance on separate channels. In the case of the MSRA dataset, the vectorial Dahu pseudo-distance has more than 60% good saliency maps with the only assumption that the boundary is mostly background. Additionally, the EMD results of the vectorial Dahu pseudo-distance is lower than the Dahu pseudo-distance on the separate channel, which proves that our proposed distance improves saliency map computation.

We present here some examples of saliency maps induced by the Dahu pseudo-distance. The saliency map ("*viso*") and the color representation of the saliency map ("*vivo*") are respectively shown in Fig. 7(d) and in Fig. 7(c). The "optimal" visual quality is reached for the vectorial Dahu pseudo-distance (compared to the Dahu pseudo-distances on separate channels or on the grayscale image). Indeed, the main barrier is clearly visible around the objects. The robustness of the vectorial Dahu pseudo-distance is easy to explain: the tree of shapes on the color image contains more information and is more structured than the tree of shape computed on separate channels.

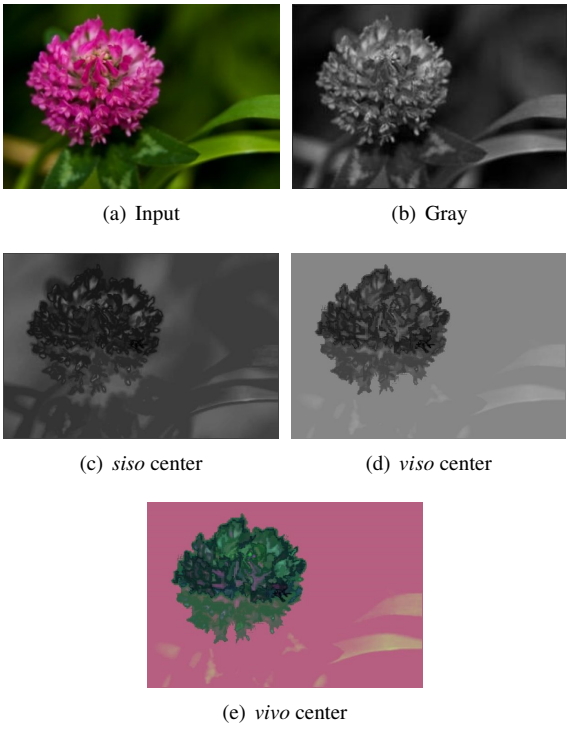

(a) Input (b) Gray

(c) *siso* center (d) *viso* center

(e) *vivo* center

Fig. 8. The saliency map deduced from the Dahu pseudo-distances when the seed point is placed in the center of the image. (a) the color image; (b) the corresponding grayscale image; (c) the "*siso*" saliency map deduced from Dahu pseudo-distance; (d) the "*viso*" and (e) *vivo* saliency map deduced from the vectorial Dahu pseudo-distance.

In another example (see Fig. 8), we compare visually the saliency map deduced from the vectorial Dahu pseudo-distance and the Dahu pseudo-distance when the seed point is placed in the center of the image. The flower zone in the "*viso*" image is spotted and is well-contrasted with the background, whereas "*siso*" image does not well distinguish between the background and the flower. Besides, in the "*viso*" image, similar intensities are obtained on most of the background regions in the distance map. Typically, the more homogeneous the distance map is in the background, the fewer seed points we need to segment the image. This is an advantage of the vectorial Dahu pseudo-distance to reduce the number of seed points for object segmentation.

5.1.2. Comparison of saliency maps of the vectorial Dahu pseudo-distance with state-of-the-art methods

Experimental setting: In this section, the saliency map computed by the vectorial Dahu pseudo-distance is compared with some saliency maps deduced from multiple MB-based methods: Fast-MBD (see Zhang et al. (2015)), MST-MBD (see Tu et al. (2016)), and Waterflow-MBD (see Huang and Zhang

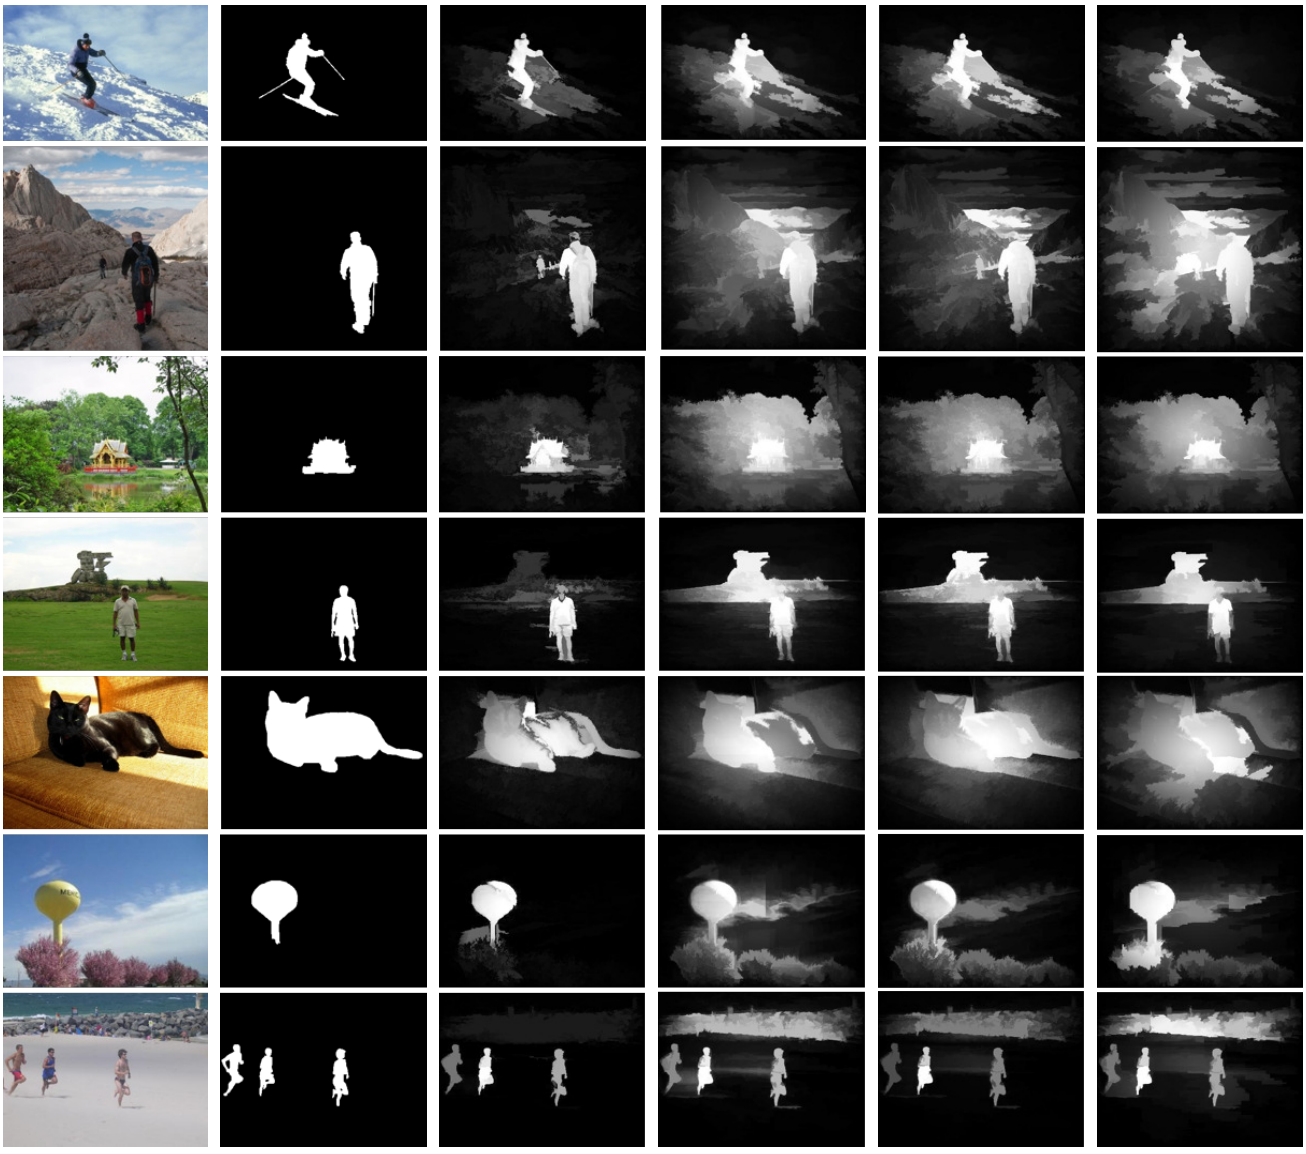

(a) Input (b) GT (c) *viso* Dahu (d) Fast-MBD (e) Waterflow-MBD (f) MST-MBD

Fig. 9. Comparison on color images of saliency maps deduced from our vectorial Dahu pseudo-distance on color images with saliency maps deduced from state-of-the-art methods.

(a) ECSSD

Method	MAE	$F_\beta{}^{max}$	EMD
Dahu	**0.21**	0.73	0.228
Fast-MBD	0.22	**0.74**	0.21
MST-MBD	0.22	0.73	0.227
Waterflow	0.22	**0.74**	**0.205**

(b) DUTOMRON

Method	MAE	$F_\beta{}^{max}$	EMD
Dahu	**0.17**	**0.634**	**0.316**
Fast-MBD	0.21	0.626	0.324
MST-MBD	0.21	0.606	0.344
Waterflow	0.21	**0.634**	**0.316**

(c) PASCAL

Method	MAE	$F_\beta{}^{max}$	EMD
Dahu	**0.22**	0.72	0.23
Fast-MBD	0.24	**0.73**	**0.22**
MST-MBD	0.24	0.72	0.23
Waterflow	0.24	**0.73**	**0.22**

(d) MSRA

Method	MAE	$F_\beta{}^{max}$	EMD
Dahu	**0.17**	0.815	0.14
Fast-MBD	0.18	0.821	0.135
MST-MBD	0.18	0.812	0.143
Waterflow	0.18	**0.824**	**0.132**

Table 2. Numerical comparison of saliency maps deduced from the vectorial Dahu pseudo-distance applied on color images and different MB-based distances adapted to manage color images. The comparison is performed using $\overline{F_\beta}$ measure and EMD score. Best scores are in bold. Results of all methods are comparable and variations among them are negligible.

(2018)). To compare these methods, we modify them, as Huang and Zhang (2018) do, by adding color and computing a color MBD by summing MBD on each channel. For the MST-MBD method, we construct an MST from the color image, then we compute the MBD in this tree. In order to fairly evaluate the performance of these methods, we add an outer border to the image and consider all pixels on the boundary image as the background. Note that, in this experiment, we just want to compare the Dahu pseudo-distance with the MB-based distance, we do not try to achieve the best results of the saliency maps. The same post-processing to normalize the saliency map, as in the previous section, is applied here.

Evaluation using MAE: Our method gives better MAE scores than other MB-based methods across all datasets. However, the

difference is very low. It can be explained by the fact that the vectorial Dahu pseudo-distance tends to give distance values lower than other MB-based distances, especially in the background regions, which constitute the largest part of an image.

Evaluations using the F_β-measure: The F_β-measure is illustrated in Table 2. At a glance, the vectorial Dahu pseudo-distance shows equivalent results to the MST-MBD method and lower results than the Fast-MBD and Waterflow-MBD methods. However, the differences between these methods are minimal. In the DUTOMRON dataset, the Dahu pseudo-distance achieves better F_β-measures than other methods. Especially, in the MSRA dataset, the Dahu pseudo-distance and MB-based methods can achieve a high value of 0.82.

Evaluation Using EM distance. For the EMD, the Fast-MBD and the Waterflow-MBD methods achieve similar results in all datasets, whereas the Dahu pseudo-distance gives comparable results with the MST-MBD method, and slightly lower results than the Fast-MBD and the Waterflow-MBD methods but here again, the difference is rather low.

Some example images are given in Fig. 9. In these images, the backgrounds are not homogeneous like in the scene of the sky, the field of grass or even the sofa image. The Dahu pseudo-distance seems to work better in these cases and achieves better performance than the MB-based distances. The tree of shapes properties and the insertion of the inter-pixels between the neighbor pixels allow the Dahu pseudo-distance to get the lower value compared to the MB-based distances. Additionally, each node on the tree of shapes is set at the median value of all the pixels in the node, which reduces the impact of noise in the color images. In the next section, we will explore this problem in greater detail.

5.2. Efficiency and robustness of the algorithm

In this section, we investigate the ability to distinguish object and background of the Dahu pseudo-distance. We also analyze the noise stability of the vectorial Dahu pseudo-distance when noise in the image increases.

5.2.1. Ability to distinguish object and background

We analyze here the ability to separate the object from the background. To do so, we measure the difference between the Dahu pseudo-distance and the MB-based distances (MST-MBD and Waterflow-MBD) between two random markers in the image by using the ratio between the inter-distance (the distance from a marker outside the object to a marker inside the object) and the intra-distance (the distance from two markers inside the object). We cannot include Fast-MBD in this comparison because the Fast-MBD (see Zhang et al. (2015)) method works only when all the seed pixels are in the boundary of the image.

We randomly create 100 markers in the image and sequentially compute the distance between two markers. The Dahu pseudo-distance between two markers X and X' is computed this way:

$$d_u^{\text{DAHU}}(X, X') = \min_{x' \in X'} \min_{x \in X} d_u^{\text{DAHU}}(x, x'). \tag{17}$$

Using the binary ground truth, the inter- and intra-distances are well defined. The contrast metric is defined by the ratio between the average of the inter-distances and the average of the intra-distances:

$$R = \frac{\frac{1}{N_1} \sum_{N_1} d_{\text{inter}}}{\frac{1}{N_2} \sum_{N_2} d_{\text{intra}}} \tag{18}$$

in which N_1 and N_2 are respectively the numbers of inter- and intra-distances.

Table 3. A comparison of ratio of inter- and intra-distances between the Dahu pseudo-distance and other MB-based methods.

Dataset	MST-MBD	Waterflow-MBD	Dahu
ECSSD	1.28	1.36	**1.404**
PASCALS	1.324	1.398	**1.448**
DUTOMRON	1.341	1.432	**1.483**
MRSA	1.784	**1.997**	**1.992**

In Table 3, the ratio of the Dahu pseudo-distance is higher than the one of the MB-based distances in all datasets. It means that the Dahu pseudo-distance is more contrasted than the MB-based distances. We can give an intuition of this result. During the front propagation process while constructing the tree of shapes, the pixel can pass through the inter-pixels. As a consequence, the Dahu pseudo-distance tends to decrease its path cost between pixels in the same background while retaining the contrast between objects and background. It leads to an increase of the ratio of the inter- and intra-distances of the Dahu pseudo-distance.

5.2.2. Robustness to noise

This section shows the impact of noise on the Dahu pseudo-distance and MB-based distances. An example image is chosen in Fig. 10 where two markers p_1 and p_2 (5×5 pixels) are set in the background and another marker p_3 is placed inside the object. A zero mean Gaussian noise is added to the image with the respective variance values: 0.0001, 0.001, 0.01, 0.1 and 0.5. One hundred noisy images are generated for each value of variance. The three markers are fixed for the entire experiment. We observe here the inter-distance $d(p_1, p_3)$ and intra-distance $d(p_1, p_2)$ during the test of the Dahu pseudo-distance or the MB-based one.

The results of the experiments are presented in Fig. 11 with the mean values as well as the associated confidence intervals. In both Fig. 11(a) and Fig. 11(b), we can see the evolution of the Dahu pseudo-distance and other MB-based distances. The MST-MBD and Waterflow-MBD both increase when the variance of noise increases. Especially when the noise variance is high, the difference between inter- and intra-distances of MST-MBD and Waterflow-MBD is minimal, whereas the ratio of inter- and intra-distances of the Dahu pseudo-distance remains more stable. This experiment shows that the vectorial Dahu pseudo-distance is robust to noise variations. This property is important for many real-world applications.

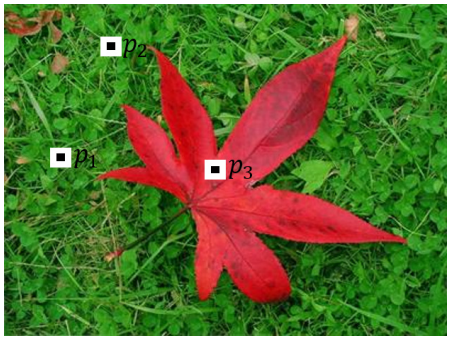

Fig. 10. An example image to investigate noise stability of the Dahu pseudo-distance and MB-based distance. The points p_1 and p_2 belong to the background, when p_3 is inside the object (this picture comes from the MSRA dataset (see Cheng et al. (2015))).

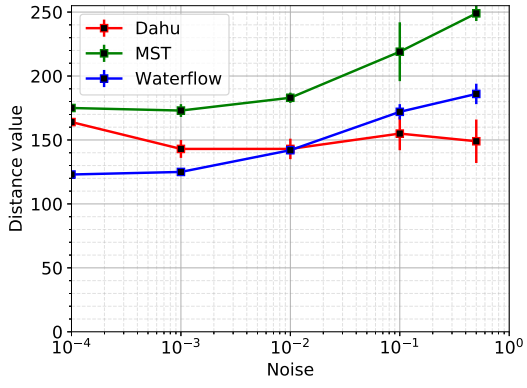

(a) Inter-distance

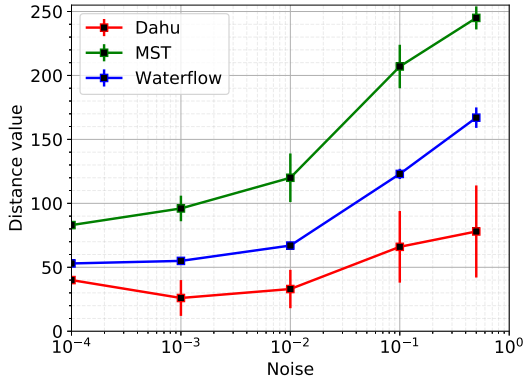

(b) Intra-distance

Fig. 11. Stability of the inter- and intra-distances using the vectorial Dahu pseudo-distance or other MB-based methods against Gaussian noise.

5.3. Speed performance

In this section, we measure the time necessary to compute numerous distances between two points using the Dahu pseudo-distance and other MB-based distances. The experiment is implemented between 100, 1000, 10000 and 100000 pairs of pixels on 20 tested images. The evaluation is conducted using a 2.6 GHz CPU with 8GB of RAM. The size of the test image is the same as used in the previous experiment (the maximum dimension is 300 pixels). Our method is implemented in C++.

The execution time is illustrated in Fig. 12 with means and confidence intervals. The construction of our tree of shapes is based on the max-tree algorithm which is designed in Carlinet et al. (2018). The whole process is linear on average (and quasi-linear at worst). The computation of the ToS runs at about 20 FPS when used on grayscale images, whereas it takes about 1 second to construct the MToS of the color image. Although the computation of the MToS is longer than the ToS, the vectorial Dahu pseudo-distance achieves better performances as we presented in Section 5.1.1. Depending on the application, we can choose either the ToS or MToS to compute the Dahu pseudo-distance. On the other hand, the construction of the MST is fast (30 FPS) and easy to implement. However, this method is sensitive to the impact of noise and usually does not provide good results in this case.

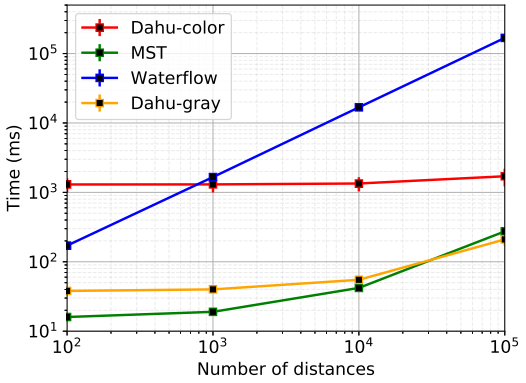

Fig. 12. Execution time (in milliseconds) to compute numerous distances between two points using the (pseudo-)distances presented in this paper.

As we can see in this figure, there is another convenient point of the Dahu pseudo-distance. For a small number of distances, the Waterflow-MBD has an advantage compared to the vectorial Dahu pseudo-distance. However, when the number of distances increases, the Dahu pseudo-distance and the MST-MBD are much faster than the Waterflow-MBD. It can be explained by the fact that the Dahu pseudo-distance and the MST-MBD take a fixed time to construct the tree, but when the tree is computed, the time to compute the distances is extremely fast thanks to the fast search of the nodes corresponding to the points in this tree. This is a huge advantage for some applications.

6. Applications

The main use of the Dahu pseudo-distance is visual saliency detection, which is considered as an intermediary step for various applications such as object detection, object segmentation and tracking. The visual saliency detection is carefully investigated in the previous section. In this section, we demonstrate the ability of the Dahu pseudo-distance in other applications. First, we present the shortest path finding application, which is a direct application of the extension of the Dahu pseudo-distance taking into account the spatial information in the image. Secondly, the Dahu pseudo-distance is applied to segment

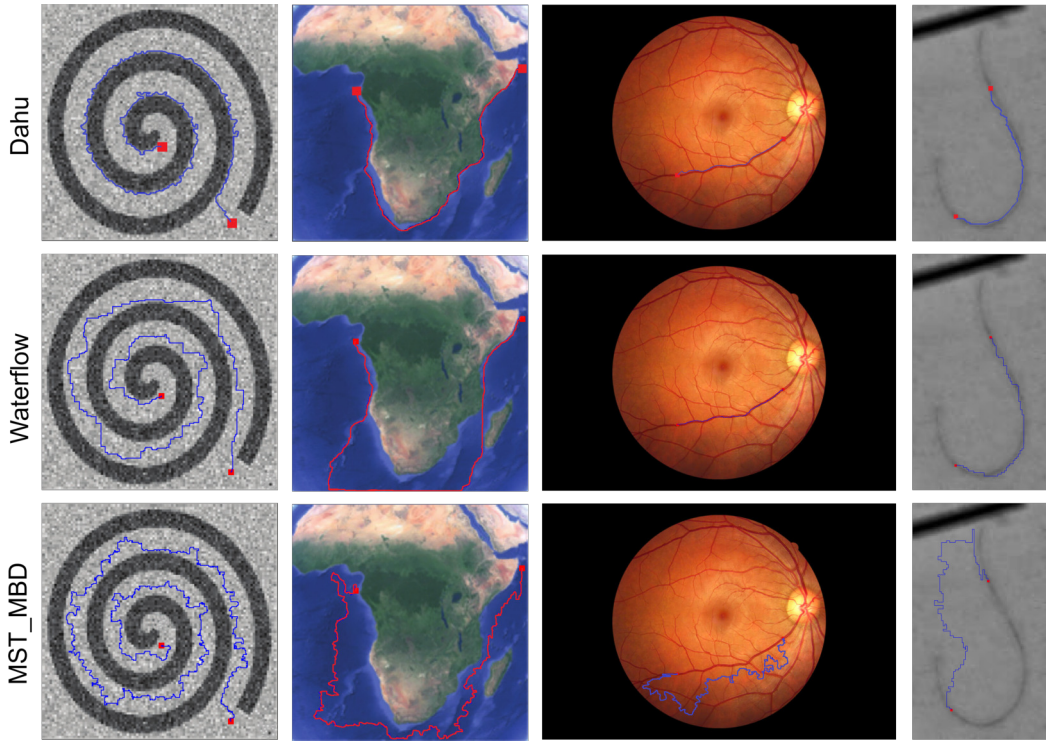

Fig. 13. Shortest path finding in images. The input images and the end points (depicted in red) of the path we want to find are shown on each picture. Results are given for Dahu pseudo-distance, Waterflow-MBD and MST-MBD. Images are extracted from Holuša and Sojka (2017) and from Vincent (1998).

the white matter region in multi-modal medical images. In the last application, we exploit the Dahu pseudo-distance to segment objects in satellite images to validate the ability of the Dahu pseudo-distance on multi-spectral images.

6.1. Shortest path finding

In this section, we validate the shortest path finding application which is presented in Section 4.3. To do that, we compare the shortest paths found by the Dahu pseudo-distance and by the other MB-based distances. Tested images, which are extracted from Vincent (1998) and from Holuša and Sojka (2017) such as a noisy synthetic image, a map image, a retinal photography and a thin glass fiber are illustrated in Fig. 13.

In the synthetic spiral image (see Fig. 13, column 1), there are two parts: the spiral and the background. We can see that the shortest path provided by the Dahu pseudo-distance is "shorter" than the ones provided by the other MB-based distances. The two chosen markers are in the background, and the shortest path between them based on our distance, follows the shape of the spiral as we expected.

Similarly to the map image (Fig. 13, column 2), the goal is to find the shortest path connecting two points located on the sea near the coast. The shortest path based on the Dahu pseudo-distance is still better than the ones using other MB-based pseudo-distances.

In the retinal image (Fig. 13, column 3), the two chosen markers are placed on a blood vessel. As demonstrated, the Dahu pseudo-distance and Waterflow-MBD give satisfying results while the MST-MBD is sensitive to noise and to blurring (its shortest path is deviated from the blood vessel).

Similarly, in the last example (see Fig. 13, column 4), the markers are placed on the glass fiber. The image is quite blurred, and the intensities of pixels along the fiber are varying, some parts of the fiber are darker than other parts. However, both the Waterflow-MBD and the Dahu pseudo-distance still find the shortest path that follows the fiber.

To conclude, the Dahu pseudo-distance achieves a better performance than the other MB-based pseudo-distances in this context.

6.2. Dahu pseudo-distance on multimodal and multispectral images

Multivariate images are widely used in various applications, ranging from medical imagery to satellite remote sensing. Multivariate can designate a multi-spectral, multi-modal or multi-source image which corresponds to a set of image channels. A color image is just a special case of multivariate image. In this section, we present the application of the vectorial Dahu pseudo-distance in multi-modal medical and multispectral satellite images. We use the same strategy to handle them, which is illustrated in Fig. 14. The method begins with the construction of the MToS. Then we put markers in the image and compute a distance map from these markers based on the Dahu pseudo-distance. Finally, we use simple thresholding to segment the object in the image.

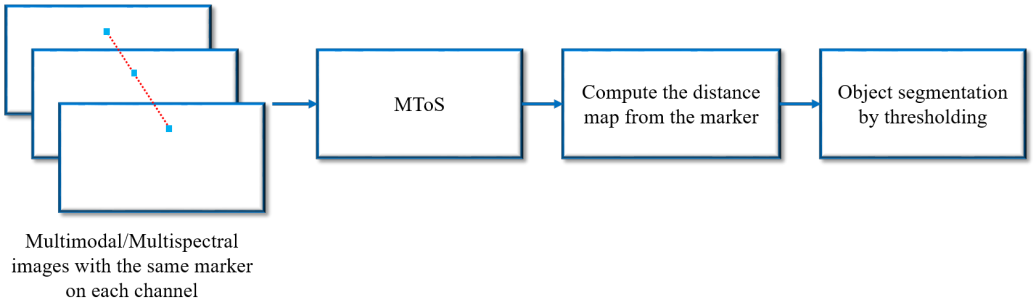

Fig. 14. A scheme for object segmentation on multimodal/multispectral images.

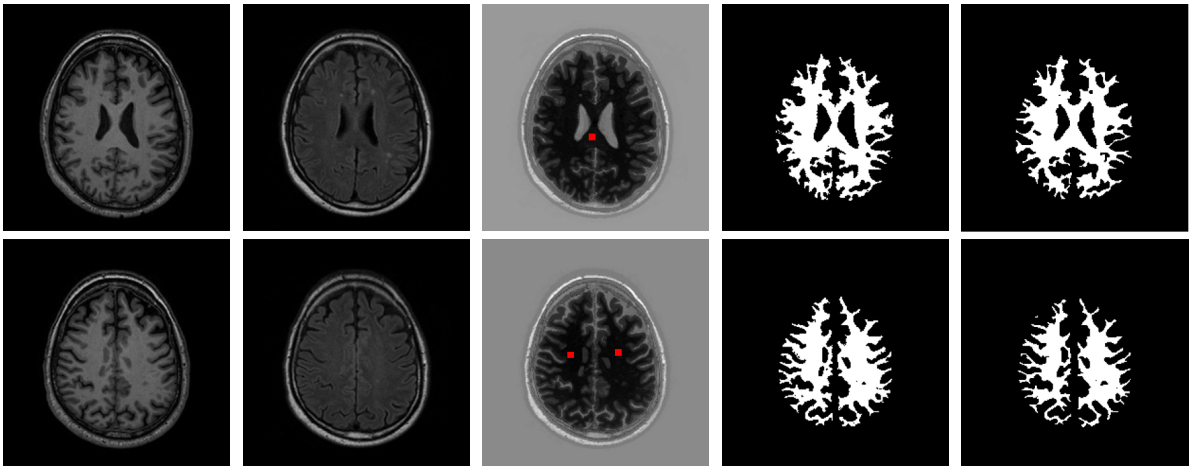

(a) T1 (b) T2-FLAIR (c) Distance map (d) Segmentation (e) GT

Fig. 15. White matter segmentation using the vectorial Dahu pseudo-distance. Images are taken from Mendrik et al. (2015). As we can observe on the segmentations (see Subfigures (*d*)) and on the ground truths (see Subfigures (*e*)), the white matter has been well segmented thanks to the vectorial Dahu pseudo-distance.

6.2.1. Multimodal medical images

Multi-modal images are becoming increasingly common in diagnosis and treatment planning (see Martí-Bonmatí et al. (2010)). They are defined as a combination of imaging modalities, which are acquired using different techniques such as computed tomography (CT), magnetic resonance imaging (MRI), and positron emission tomography (PET). Multi-modal images are also used to overcome the limitations induced by specific activities of each individual technique. In this subsection, we applied the vectorial Dahu pseudo-distance to segment the white matter in 3D brain MR images.

We consider two images: the T1 (Fig. 15(a)) and the T2-FLAIR slice (Fig. 15(b)) as inputs of our experiment. Then, we construct the MToS on these images to get the mutual information from different machines. A marker (5 x 5 pixels) is put on the white matter region to compute a Dahu distance map (Fig. 15(c)). We first remark that the MToS preserves the geometric information of the two channels and mixes them in a sensible way. We further observe that the distance map gives low values to the white matter region. A simple threshold method is used to segment the white matter region in the image (Fig. 15(d)). As can be seen, our method not only achieves good segmentation results compared to the ground truth image, but the vectorial Dahu pseudo-distance proves to be efficient for this experiment.

6.2.2. Satellite multi-spectral images

Over the past few years, the use of multi-spectral images has been increasingly investigated in many applications, especially in target detection and recognition (see Bioucas-Dias et al. (2013)). Multi-spectral images collect information from hundreds of spectrum bands, thus providing a powerful tool to discriminate different objects. Similarly to the usage of the vectorial Dahu pseudo-distance in the previous section, we employ the vectorial Dahu pseudo-distance to segment object regions in the image.

We apply our method on the Pavia University dataset (see Licciardi et al. (2009)). It consists of 103 images which correspond each to a spectral channel. The dataset has a size of 610*340 pixels, contains nine classes which represent trees, meadows, asphalt, etc. The images are pre-processed with a *P.C.A* algorithm (see Jolliffe (1986)) to reduce the correlation among the bands. This algorithm also selects the best bands for object detection. This pre-processing relies on the fact that neighbor bands of multi-spectral images are highly correlated and contain mutual information about the object.

In our case, we choose the first 5 channel components. As

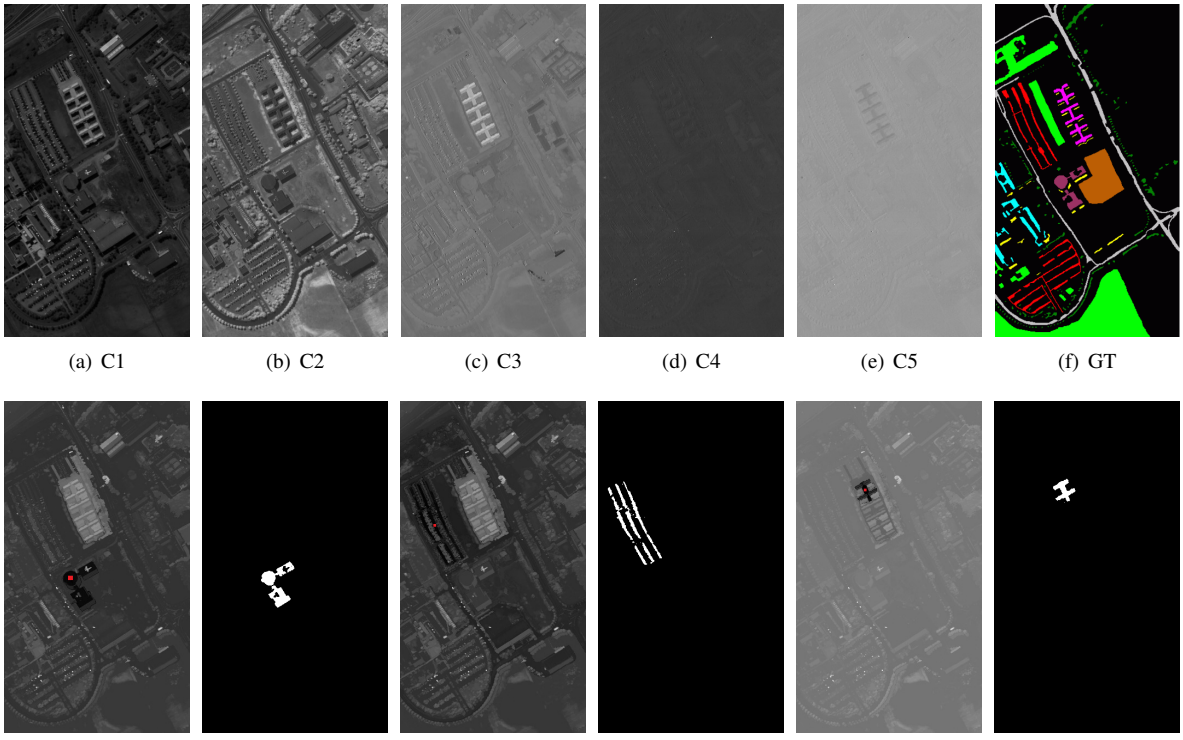

(a) C1 (b) C2 (c) C3 (d) C4 (e) C5 (f) GT

(g) Distance map (h) Segmentation (i) Distance map (j) Segmentation (k) Distance map (l) Segmentation

Fig. 16. Object segmentation on multispectral images. Objects are manually selected with a marker (in red in pictures). Images C1-C5 are extracted by using a principal component analysis (PCA) algorithm.

we can see in Fig. 16, some objects clearly appear in some images but not in the others. The MToS is then constructed on these images. We put some markers in the image to compute the distance map. Then a simple threshold is used to segment the object in the image. As we can see in Fig. 16, our method can segment the objects in the image with high accuracy, for instance, the painted metal sheets, the bitumen, and self-blocking bricks classes. These results demonstrate the robustness of the vectorial Dahu pseudo-distance in this context.

7. Conclusions and perspectives

In this paper, we have studied the Dahu pseudo-distance and have presented multiple improvements. First, we have introduced a vectorial extension capable of dealing with multi-channels images. Obviously, this vectorial Dahu pseudo-distance processes color images which is already a great improvement. However, it is also not restricted to three channels images. Second, we have improved the Dahu pseudo-distance by combining the pseudo-distance with information on the spatial domain of the images. Such an improvement opens new areas of applications, in competition with the commonly used *geodesic distance*.

After having compared our new versions with several MB-based pseudo-distances in many situations and applications, we have proven that taking into account the color of the images brings noticeable improvements. We have also proven that our vectorial Dahu pseudo-distance is less affected by noise in the image than other MB-based pseudo-distances.

We have further demonstrated the improvement induced by this new vectorial Dahu pseudo-distance, since we have shown that it can handle multimodal and multispectral images by testing it on multimodal medical images and multi-spectral satellite images.

Another advantage of our new vectorial version is that it comes at almost no additional cost. Thanks to a clever representation of images, the multivariate tree of shapes, the distance is quasi instantaneous to compute (and the tree can be computed in a quasi linear time with respect to the number of pixels of the images). It is then possible to use it in real time.

In the future, we plan to use the vectorial Dahu pseudo-distance in some applications like automatic object detection and interactive segmentation. Furthermore, we want to investigate the case of embedded environments.

Acknowledgment

This work has been conducted in the context of the MOBIDEM project, part of the “Systematic Paris-Region” and “Images & Network” Clusters (France). This project is partially funded by the French Government and its economic development agencies.

References

Aubin, J.P., Frankowska, H., 2009. Set-valued analysis. Springer Science & Business Media.

Bharati, S.P., Nandi, S., Wu, Y., Sui, Y., Wang, G., 2016. Fast and robust object tracking with adaptive detection, in: Proceedings of the 28th International Conference on Tools with Artificial Intelligence (ICTAI), pp. 706–713.

Bioucas-Dias, J.M., Plaza, A., Camps-Valls, G., Scheunders, P., Nasrabadi, N., Chanussot, J., 2013. Hyperspectral remote sensing data analysis and future challenges. IEEE Geoscience and Remote Sensing Magazine 1, 6–36.

Calarasanu, S., Fabrizio, J., Dubuisson, S., 2015. Using histogram representation and Earth Mover's Distance as an evaluation tool for text detection, in: Proceedings of the International Conference on Document Analysis and Recognition (ICDAR), pp. 221–225.

Cao, F., Lisani, J.L., Morel, J.M., Musé, P., Sur, F., 2008. A Theory of Shape Identification. volume 1948 of *Lecture Notes in Mathematics*. Springer.

Carlinet, E., Crozet, S., Géraud, T., 2018. The tree of shapes turned into a max-tree: A simple and efficient linear algorithm, in: Proceedings of the IEEE International Conference on Image Processing (ICIP), pp. 1488–1492.

Carlinet, E., Géraud, T., 2014. A comparative review of component tree computation algorithms. IEEE Transactions on Image Processing 23, 3885–3895.

Carlinet, E., Géraud, T., 2015. MToS: A tree of shapes for multivariate images. IEEE Transactions on Image Processing 24, 5330–5342.

Caselles, V., Coll, B., Morel, J.M., 1999. Topographic maps and local contrast changes in natural images. International Journal on Computer Vision 33, 5–27.

Caselles, V., Monasse, P., 2009. Geometric Description of Images as Topographic Maps. volume 1984 of *Lecture Notes in Mathematics*. Springer.

Cheng, M.M., Mitra, N.J., Huang, X., Torr, P.H.S., Hu, S.M., 2015. Global contrast based salient region detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 37, 569–582.

Ciesielski, K.C., et al., 2014. Efficient algorithm for finding the exact minimum barrier distance. Computer Vision and Image Understanding 123, 53–64.

Crozet, S., Géraud, T., 2014. A first parallel algorithm to compute the morphological tree of shapes of nD images, in: Proceedings of the IEEE International Conference on Image Processing (ICIP), pp. 2933–2937.

Géraud, T., Carlinet, E., Crozet, S., Najman, L., 2013. A quasi-linear algorithm to compute the tree of shapes of n-D images., in: Proceedings of the International Symposium on Mathematical Morphology (ISMM). Springer. volume 7883 of *Lecture Notes in Computer Science*, pp. 98–110.

Géraud, T., Xu, Y., Carlinet, E., Boutry, N., 2017. Introducing the Dahu pseudo-distance, in: Proceedings of the International Symposium on Mathematical Morphology (ISMM). Springer. volume 10225 of *Lecture Notes in Computer Science*, pp. 55–67.

Grand-Brochier, M., Vacavant, A., Strand, R., Cerutti, G., Tougne, L., 2014. About the impact of pre-processing tools on segmentation methods applied for tree leaves extraction, in: Proceedings of the International Conference on Computer Vision Theory and Applications (VISAPP), pp. 507–514.

Hart, P.E., Nilsson, N.J., Raphael, B., 1968. A formal basis for the heuristic determination of minimum cost paths. IEEE Transactions on Systems Science and Cybernetics 4, 100–107.

Holuša, M., Sojka, E., 2017. The k-max distance in graphs and images. Pattern Recognition Letters 98, 103–109.

Hu, Y., Li, Y., Song, R., Rao, P., Wang, Y., 2018. Minimum barrier superpixel segmentation. Image and Vision Computing 70, 1–10.

Huang, X., Zhang, Y., 2018. Water flow driven salient object detection at 180 FPS. Pattern Recognition Letters 76, 95–107.

Huang, X., Zheng, Y., Huang, J., Zhang, Y.J., 2018. A minimum barrier distance based saliency box for object proposals generation. IEEE Signal Processing Letters 25.

Jolliffe, I.T., 1986. Principal Component Analysis. Springer. chapter Principal components in regression analysis. Series in Statistics, pp. 129–155.

Kårsnäs, A., Strand, R., Saha, P.K., 2012. The vectorial minimum barrier distance, in: Proceedings of the International Conference on Pattern Recognition (ICPR), pp. 792–795.

Kovalevsky, V., 1986. On the topology of discrete spaces. Studientexte, Digitale Bildverarbeitung .

Li, Y., Hou, X., Koch, C., Rehg, J.M., Yuille, A.L., 2014. The secrets of salient object segmentation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 280–287.

Licciardi, G., Pacifici, F., Tuia, D., Prasad, S., West, T., Giacco, F., Thiel, C., Inglada, J., Christophe, E., Chanussot, J., et al., 2009. Decision fusion for the classification of hyperspectral data: Outcome of the 2008 GRS-S data fusion contest. IEEE Transactions on Geoscience and Remote Sensing 47, 3857–3865.

Liu, N., Ju, R., Ren, T., Wu, G., 2016. A saliency-guided method for automatic photo refocusing, in: Proceedings of the International Conference on Internet Multimedia Computing and Service (ICIMCS), pp. 264–267.

Malmberg, F., Nordenskjöld, R., Strand, R., Kullberg, J., 2017. Smartpaint: a tool for interactive segmentation of medical volume images. Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization (CMBBE) 5, 36–44.

Margolin, R., Zelnik-Manor, L., Tal, A., 2014. How to evaluate foreground maps?, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 248–255.

Martí-Bonmatí, L., Sopena, R., Bartumeus, P., Sopena, P., 2010. Multimodality imaging techniques. Contrast Media & Molecular Imaging 5, 180–189.

Mendrik, A.M., Vincken, K.L., Kuijf, H.J., Breeuwer, M., Bouvy, W.H., De Bresser, J., Alansary, A., De Bruijne, M., Carass, A., El-Baz, A., et al., 2015. MRBrainS challenge: Online evaluation framework for brain image segmentation in 3T MRI scans. Computational Intelligence and Neuroscience 2015, 1.

Monasse, P., Guichard, F., 2000. Fast computation of a contrast-invariant image representation. IEEE Transactions on Image Processing 9, 860–872.

Najman, L., Géraud, T., 2013. Discrete set-valued continuity and interpolation, in: Proceedings of the International Symposium on Mathematical Morphology (ISMM). Springer. volume 7883 of *Lecture Notes in Computer Science*, pp. 37–48.

Ôn Vũ Ngọc, M., Fabrizio, J., Géraud, T., 2018. Saliency-based detection of identity documents captured by smartphones, in: Processing of the IAPR International Workshop on Document Analysis Systems (DAS), pp. 387–392.

Shi, J., Yan, Q., Xu, L., Jia, J., 2016. Hierarchical image saliency detection on extended cssd. IEEE Transactions on Pattern Analysis and Machine Intelligence 38, 717–729.

Strand, R., et al., 2013. The minimum barrier distance. Computer Vision and Image Understanding 117, 429–437.

Strand, R., et al., 2017. The minimum barrier distance: A summary of recent advances, in: Proceedings of the International Conference on Discrete Geometry for Computer Imagery (DGCI). Springer. volume 10502 of *LNCS*, pp. 57–68.

Toivanen, P.J., 1996. New geodesic distance transforms for gray-scale images. Pattern Recognition Letters 17, 437–450.

Tu, W.C., He, S., Yang, Q., Chien, S.Y., 2016. Real-time salient object detection with a minimum spanning tree, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2334–2342.

Vincent, L., 1998. Minimal path algorithms for the robust detection of linear features in gray images. Computational Imaging and Vision 12, 331–338.

Wang, A., Wang, M., 2017. RGB-D salient object detection via minimum barrier distance transform and saliency fusion. IEEE Signal Processing Letters 24, 663–667.

Wang, G., Zhang, Y., Li, J., 2017. High-level background prior based salient object detection. Journal of Visual Communication and Image Representation 48, 432–441.

Wei, Y., Wen, F., Zhu, W., Sun, J., 2012. Geodesic saliency using background priors, Springer. volume 7574 of *Lecture Notes in Computer Science*, pp. 29–42.

Xiao, H., Feng, J., Lin, G., Liu, Y., Zhang, M., 2018. Monet: Deep motion exploitation for video object segmentation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1140–1148.

Yang, B., Zhang, X., Chen, L., Yang, H., Gao, Z., 2017. Edge guided salient object detection. Neurocomputing 221, 60–71.

Zhang, J., Sclaroff, S., Lin, Z., Shen, X., Price, B., Mech, R., 2015. Minimum barrier salient object detection at 80 FPS, in: Proceedings of the International Conference on Computer Vision (ICCV), pp. 1404–1412.

Zhang, J.M., Shen, Y.X., 2017. Spectral segmentation via minimum barrier distance. Multimedia Tools and Applications 76, 25713–25729.

Zhang, L., Yang, C., Lu, H., Ruan, X., Yang, M.H., 2017. Ranking saliency. IEEE Transactions on Pattern Analysis and Machine Intelligence 39, 1892–1904.