

2017-08-22

The boolean map distance: theory and efficient computation

Malmberg F., Strand R., Zhang J., Sclaroff S. (2017) The Boolean Map Distance: Theory and Efficient Computation. In: Kropatsch W., Artner N., Janusch I. (eds) Discrete Geometry for Computer Imagery. DGCI 2017. Lecture Notes in Computer Science, vol 10502. Springer, Cham
<https://hdl.handle.net/2144/26693>

"Downloaded from OpenBU. Boston University's institutional repository."

The Boolean Map Distance: Theory and Efficient Computation

Filip Malmberg¹, Robin Strand¹, Jianming Zhang², Stan Sclaroff³

¹ Centre for Image Analysis, Department of Information Technology, Uppsala University, Sweden

² Adobe Research, San Jose, USA

³ Department of Computer Science, Boston University, USA

Abstract. We propose a novel distance function, the *boolean map distance* (BMD), that defines the distance between two elements in an image based on the probability that they belong to different components after thresholding the image by a randomly selected threshold value. This concept has been explored in a number of recent publications, and has been proposed as an approximation of another distance function, the *minimum barrier distance* (MBD). The purpose of this paper is to introduce the BMD as a useful distance function in its own right. As such it shares many of the favorable properties of the MBD, while offering some additional advantages such as more efficient distance transform computation and straightforward extension to multi-channel images.

1 Introduction

Distance functions and their transforms (DTs, where each pixel is assigned the distance to a set of *seed pixels*) are used extensively in many image processing applications. Here, we introduce a novel distance function, the *boolean map distance* (BMD), that defines the distance between two elements in an image based on the probability that they belong to different components after thresholding the image by a randomly selected threshold value. The idea of considering connectivity with respect to randomly selected thresholds was first introduced by Zhang and Sclaroff [5], who considered the probability that a pixel does not belong to a component that touches the image border after thresholding with a randomly selected value. This probability was used to identify salient objects in an image, and the concept was named *boolean map saliency*. Here, we consider the probability that *any* pair of pixels are not connected with respect to a random threshold, and thus use the term *boolean map distance* (BMD). Ideas related to this concept have further been explored in a number of recent publications [6, 4, 1], and have been used as an approximation of another distance function, the *minimum barrier distance* (MBD) first proposed by Strand et al. [4].

The purpose of this paper is to introduce the BMD as a useful distance function in its own right. As such the BMD shares many of the favorable properties of the MBD, while offering some additional advantages such as more efficient distance transform computation and straightforward extension to multi-channel images. Specifically, our contributions are as follows:

- We provide a formal definition of the BMD, for both continuous and discrete images.
- We show that the (continuous and discrete domain) BMD is a pseudo-metric, thus motivating the name boolean map distance.
- We prove the equivalence between the (continuous and discrete domain) BMD and the φ mapping introduced by Strand et al. [4] as an approximation to the MBM. Thereby, we strengthen the connection between the BMD and the MBM. Previously, this equivalence was only established in the discrete case [6].
- We summarize available algorithms for computing distance transforms for the discrete BMD. Specifically, we note that the equivalence between the discrete BMD and φ functions allows the BMD to be expressed as the difference between two well-known distance functions, whose distance transforms can be computed using the *image foresting transform* [2], a generalization of Dijkstra’s algorithm. We demonstrate empirically that the resulting algorithm is an order of magnitude faster than previously reported algorithms, while still producing exact results.

2 The Boolean Map Distance in \mathbb{R}^n

We define a n -dimensional gray-scale image I as a pair $I = (D, f)$, where $D \subset \mathbb{R}^n$ and $f : D \rightarrow [0, 1]$ is a continuous function. The restriction of the image values to the range $[0, 1]$ does, for the purposes considered here, not imply a loss of generality [6].

For $p, q \in D$, a *path* from p to q (in D) is any continuous function $\pi : [0, 1] \rightarrow D$ with $\pi(0) = p$ and $\pi(1) = q$. We use the symbol $\Pi_{p,q}^D$ to denote the family of all such paths. The *reverse* π^{-1} of a path π is defined as $\pi^{-1}(s) = \pi(1 - s)$ for all $s \in [0, 1]$. Recall that $D \subset \mathbb{R}^n$ is *path connected* if for every $p, q \in D$ there exists a path $\pi : [0, 1] \rightarrow D$ from p to q . For the remainder of this section, we assume that the set D is path connected.

Let $t \in [0, 1]$. Given an image $I = (D, f)$, we define the function $T_{I,t} : D \rightarrow \{0, 1\}$ by

$$T_{I,t}(p) = \begin{cases} 0, & \text{if } f(p) < t \\ 1, & \text{otherwise} \end{cases}, \quad (1)$$

for all $p \in D$. We refer to any function that maps image elements to the set $\{0, 1\}$ as a *boolean map*. The boolean map $T_{I,t}$ represents the *thresholding* of the image I by t . For any $p, q \in D$, we say that p and q belong to the same component of $T_{I,t}$ if there exists a path $\pi \in \Pi_{p,q}^D$ such that either $T_{I,t}(\pi(s)) = 0$ for all $s \in [0, 1]$ or $T_{I,t}(\pi(s)) = 1$ for all $s \in [0, 1]$. A path satisfying either of these criteria is called a *connecting path*. Otherwise, p and q belong to different components of $T_{I,t}$. We use the notation $p \underset{t}{\sim} q$ to indicate that p and q belong to the same component of $T_{I,t}$, while the notation $p \underset{t}{\not\sim} q$ indicates that they belong to different components.

Definition 1. Let t be a random value sampled from a uniform probability distribution over $[0, 1]$. The continuous domain boolean map distance $BMD : D \times D \rightarrow [0, 1]$ is defined as

$$BMD(p, q) = P(p \not\sim_t q) = 1 - P(p \sim_t q) \quad (2)$$

for all $p, q \in D$, where $P(A)$ denotes the probability of the event A .

Definition 2. A function $d : D \times D \rightarrow [0, \infty)$ is a pseudo-metric on a set D if, for every $p, q, r \in D$,

- (i) $d(p, p) = 0$ (identity)
- (ii) $d(p, q) \geq 0$ (non-negativity)
- (iii) $d(p, q) = d(q, p)$ (symmetry)
- (iv) $d(p, r) \leq d(p, q) + d(q, r)$ (triangle inequality)

If additionally it holds that $d(p, q) = 0 \Leftrightarrow p = q$ for all p, q , then d is a metric.

In the proof that BMD is a pseudo-metric, we will use the following notion. The concatenation $\pi_1 \cdot \pi_2$ of the paths π_1 and π_2 such that $\pi_1(1) = \pi_2(0)$ is

$$(\pi_1 \cdot \pi_2)(s) = \begin{cases} \pi_1(2s) & \text{if } s \in [0, 1/2] \\ \pi_2(2s) & \text{otherwise} \end{cases}, \quad (3)$$

Theorem 1. BMD is a pseudo-metric.

Proof. First, we show that BMD obeys property (i). Consider a path $\pi(s)$ such that $\pi(s) = p$ for any $s \in [0, 1]$. This path is obviously connecting p to itself. Thus, $p \sim_t p$ for all t , and so $P(p \not\sim_t p) = 0$.

Since BMD is defined as a probability, we have $BMD(p, q) \in [0, 1]$ for all p, q . Thus, BMD clearly obeys property (ii).

Next, we show that BMD obeys property (iii). If, for a given threshold t , there exists a connecting path π from p to q then the reverse path π^{-1} is a connecting path from q to p . Thus, $p \not\sim_t q \Leftrightarrow q \not\sim_t p$, and so $P(p \not\sim_t q) = P(q \not\sim_t p)$.

Finally, we show that BMD obeys property (iv). If, for a given threshold t , there exists a connecting path π_1 from p to q and another connecting path π_2 from q to r , then the concatenation $\pi_1 \cdot \pi_2$ of these two paths is a connecting path from p to r . Thus, the set of thresholds t for which $q \sim_t p$ and $q \sim_t r$ is a subset of the set of thresholds for which $p \sim_t r$, and so $BMD(p, r) \leq BMD(p, q) + BMD(q, r)$. \square

Note that for a constant function f we have $P(p \not\sim_t q) = 0$ for all $p, q \in D$, and thus BMD is not in general a metric.

2.1 Equivalence between BMD and φ

In this section, we prove the equivalence between the proposed BMD mapping and the φ mapping defined by Strand et al. (Definition 3 in [4]). We recall the definition of φ :

Definition 3. Let $I = (D, f)$. The mapping $\varphi : D \times D \rightarrow [0, \infty)$ is defined as

$$\varphi(p, q) = \inf_{\pi_1 \in \Pi_{p,q}^D} \max_s f(\pi_1(s)) - \sup_{\pi_2 \in \Pi_{p,q}^D} \min_s f(\pi_2(s)) \quad (4)$$

for all $p, q \in D$.

As stated by Strand et al. [4] the minimum/maximum over the numbers $s \in [0, 1]$ are attained, while neither the infimum nor supremum operators can be replaced by maximum/minimum.

Theorem 2. The mappings BMD and φ are equal, i.e. $\varphi(p, q) = BMD(p, q)$ for all $p, q \in D$.

Proof. We begin our proof by observing that

$$BMD(p, q) = P(p \underset{t}{\approx} q) = 1 - P(p \underset{t}{\sim} q) = 1 - P(A \vee B) \quad (5)$$

where A is the event $(p \underset{t}{\sim} q \wedge f(p) < t \wedge f(q) < t)$ and B is the event $(p \underset{t}{\sim} q \wedge f(p) \geq t \wedge f(q) \geq t)$. Since the events A and B are mutually exclusive, it follows that

$$BMD(p, q) = 1 - P(A) - P(B). \quad (6)$$

First, we study the probability $P(A)$ that p and q belong to the same component of $T_{I,t}$, and $f(p)$ and $f(q)$ are both less than t . This is true if there exists a path $\pi \in \Pi_{p,q}^D$ such that

$$f(\pi(s)) < t \text{ for all } s \in [0, 1]. \quad (7)$$

Let $c = \inf_{\pi_1 \in \Pi_{p,q}^D} \max_s f(\pi_1(s))$. Then a path π satisfying the condition given in eq. (7) exists if $t > c$, but does not exist if $t < c$. If $t = c$, then the existence of π is not possible to determine in the general case but depends on whether, for the specific f , p and q at hand, there exists a path whose maximum value $\max_s f(\pi(s))$ attains the infimum over all paths in $\Pi_{p,q}^D$. Let $\phi : [0, 1] \rightarrow \{0, 1\}$ be an indicator function for the event A , defined by

$$\phi(t) = \begin{cases} 1 & \text{if } A \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

Depending on the specific values of f , p and q , we have either $\phi = \phi_1$ or $\phi = \phi_2$ where

$$\phi_1(t) = \begin{cases} 1 & \text{if } t \geq c \\ 0 & \text{otherwise} \end{cases} \quad \phi_2(t) = \begin{cases} 1 & \text{if } t > c \\ 0 & \text{otherwise} \end{cases} . \quad (9)$$

We note that

$$P(A) = \int_0^1 \phi(t)dt = \int_0^1 \phi_1(t)dt = \int_0^1 \phi_2(t)dt = 1 - c = 1 - \inf_{\pi_1 \in \Pi_{p,q}^D} \max_s f(\pi_1(s)) . \quad (10)$$

Thus, regardless of the existence of the path π in the case where $t = c$, the probability of the event A occurring is $P(A) = 1 - \inf_{\pi_1 \in \Pi_{p,q}^D} \max_s f(\pi_1(s))$.

Next, we study the probability $P(B)$ that p and q belong to the same component of $T_{I,t}$, and that $f(p)$ and $f(q)$ are both greater than or equal to t . This is true if there exists a path $\pi \in \Pi_{p,q}^D$ such that $f(\pi(s)) \geq t$ for all $s \in [0, 1]$. Such a path exists iff $t < \sup_{\pi_2 \in \Pi_{p,q}^D} \min_s f(\pi_2(s))$. The probability of the event B occurring is $P(B) = \sup_{\pi_2 \in \Pi_{p,q}^D} \min_s f(\pi_2(s))$.

From eq. (6), it thus follows that

$$BMD(p, q) = \inf_{\pi_1 \in \Pi_{p,q}^D} \max_s f(\pi_1(s)) - \sup_{\pi_2 \in \Pi_{p,q}^D} \min_s f(\pi_2(s)) = \varphi(p, q) . \quad (11)$$

□

3 The discrete Boolean Map Distance

In this section, we introduce a discrete formulation of the BMD.

We define a discrete gray-scale digital image \hat{I} as a pair $\hat{I} = (\hat{D}, f)$ consisting of a finite set \hat{D} of image elements and a mapping $f : \hat{D} \rightarrow [0, 1]$. We will refer to elements of \hat{D} as *pixels*, regardless of the dimensionality of the image. Additionally, we define a mapping $\mathcal{N} : \hat{D} \rightarrow \mathcal{P}(\hat{D})$ specifying an *adjacency* relation over the set of pixels \hat{D} . For any $p, q \in \hat{D}$, we refer to $\mathcal{N}(p)$ as the *neighborhood* of p and say that q is *adjacent* to p if $q \in \mathcal{N}(p)$. We require the adjacency relation to be symmetric, so that $q \in \mathcal{N}(p) \Leftrightarrow p \in \mathcal{N}(q)$ for all $p, q \in \hat{D}$.

A discrete path $\hat{\pi} = \langle \hat{\pi}(0), \hat{\pi}(1), \dots, \hat{\pi}(k) \rangle$ of length $|\hat{\pi}| = k + 1$ from $\hat{\pi}(0)$ to $\hat{\pi}(k)$ is an ordered sequence of pixels in \hat{D} where each consecutive pair of pixels are *adjacent*. We use the symbol $\hat{\Pi}_{p,q}^{\hat{D}}$ to denote the set of all discrete paths from p to q . For a set of pixels $S \subseteq \hat{D}$, the symbol $\hat{\Pi}_{p,S}^{\hat{D}}$ denotes the set of all discrete paths from p to any $q \in S$. The *reverse* $\hat{\pi}^{-1}$ of a discrete path $\hat{\pi}$ is defined as $\hat{\pi}^{-1}(i) = \hat{\pi}(k - i)$ for all $i \in \{0, 1, \dots, k\}$. Given two discrete paths $\hat{\pi}_1$ and $\hat{\pi}_2$ such that the endpoint of $\hat{\pi}_1$ equals the starting point of $\hat{\pi}_2$, we denote by $\hat{\pi}_1 \cdot \hat{\pi}_2$ the concatenation of the two paths.

Throughout, we assume that the combination of the set \hat{D} and the adjacency relation \mathcal{N} defines a connected graph, so that for every pair of pixels $p, q \in \hat{D}$ there exists a path between them.

Let $t \in [0, 1]$. Given a discrete image $\hat{I} = (\hat{D}, f)$, we define the thresholding $T_{\hat{I}, t} : \hat{D} \rightarrow \{0, 1\}$ of \hat{I} by t as

$$T_{\hat{I}, t}(p) = \begin{cases} 0, & f(p) < t \\ 1, & \text{otherwise} \end{cases}, \quad (12)$$

For any $p, q \in \hat{D}$, we say that p and q belong to the same component of $T_{\hat{I}, t}$ if there exists a path $\hat{\pi} \in \hat{\Pi}_{p, q}^{\hat{D}}$ such that either $T_{\hat{I}, t}(\hat{\pi}(i)) = 0$ for all $i \in \{0, 1, \dots, k\}$ or $T_{\hat{I}, t}(\hat{\pi}(i)) = 1$ for all $i \in \{0, 1, \dots, k\}$. A path satisfying either of these criteria is called a *connecting path*. As before, we use the notation $p \underset{t}{\sim} q$ to indicate that p and q belong to the same component of $T_{\hat{I}, t}$, while the notation $p \not\underset{t}{\sim} q$ indicates that they belong to different components. Additionally, for a set of pixels S , the notation $p \underset{t}{\sim} S$ indicates that $p \underset{t}{\sim} q$ for at least one $q \in S$ while the notation $p \not\underset{t}{\sim} S$ indicates that $p \not\underset{t}{\sim} q$ for all $q \in S$.

Definition 4. Let t be a random value sampled from a uniform probability distribution over $[0, 1]$. For any set of pixels $S \subseteq \hat{D}$ and pixel $p \in \hat{D}$, the discrete boolean map distance $B\hat{M}D : \hat{D} \times \mathcal{P}(\hat{D}) \rightarrow [0, 1]$ is defined as

$$B\hat{M}D(p, S) = P(p \not\underset{t}{\sim} S) \quad (13)$$

In the above definition, $\mathcal{P}(\hat{D})$ denotes the power set of \hat{D} . If the set S consists of a single element q , we can consider $B\hat{M}D$ to be a mapping from $\hat{D} \times \hat{D}$ to $[0, 1]$, and the definition can in this case be reduced to $B\hat{M}D(p, q) = P(p \not\underset{t}{\sim} q)$.

Theorem 3. Let $S \subseteq \hat{D}$ consist of a single element q and let $p \in \hat{D}$. Then the discrete $B\hat{M}D$, viewed as a mapping from $\hat{D} \times \hat{D}$ to $[0, 1]$, is a pseudo-metric.

The proof of Theorem 3 is identical to that of Theorem 1, provided that the relevant continuous notions defined in Section 2 are swapped out for their discrete counterparts defined in this section.

3.1 Equivalence between the discrete $B\hat{M}D$ and $\hat{\varphi}$

In this section, we prove the equivalence between the discrete $B\hat{M}D$ mapping and the $\hat{\varphi}$ mapping defined by Strand et al. (eq. (4) in [4]) as a discrete counterpart of the φ mapping. We provide a slightly extended definition of $\hat{\varphi}$:

Definition 5. The mapping $\hat{\varphi} : \hat{D} \times \mathcal{P}(\hat{D}) \rightarrow [0, 1]$ is defined as

$$\hat{\varphi}(p, S) = \min_{\hat{\pi}_1 \in \hat{\Pi}_{p, S}^{\hat{D}}} \left(\max_{i \in \{0, 1, \dots, k\}} I(\hat{\pi}_1(i)) \right) - \max_{\hat{\pi}_2 \in \hat{\Pi}_{p, S}^{\hat{D}}} \left(\min_{i \in \{0, 1, \dots, k\}} I(\hat{\pi}_2(i)) \right) \quad (14)$$

for all $p \in \hat{D}$ and $S \subseteq \hat{D}$.

Note that if S consists of a single element q , the above definition of $\hat{\varphi}$ reduces to the definition given by Strand et al. [4].

Theorem 4. *The mappings $B\hat{M}D$ and $\hat{\varphi}$ are equal, i.e., $\hat{\varphi}(p, S) = B\hat{M}D(p, S)$ for all $p \in \hat{D}$ and $S \subseteq \hat{D}$.*

Proof. We start by observing that

$$B\hat{M}D(p, S) = P(p \not\sim_t S) = 1 - P(p \sim_t S) = 1 - P(A \vee B) \quad (15)$$

where A is the event $(p \sim_t q_1 \wedge f(p) < t \wedge f(q_1) < t)$ for some $q_1 \in S$ and B is the event $(p \sim_t q_2 \wedge f(p) \geq t \wedge f(q_2) \geq t)$ for some $q_2 \in S$. Since the events A and B are mutually exclusive, it follows that

$$B\hat{M}D(p, S) = 1 - P(A) - P(B). \quad (16)$$

First, we study the probability $P(A)$ that p and q_1 belong to the same component of $T_{\hat{I}, t}$, and $f(p)$ and $f(q_1)$ are both less than t . This is true for some p_1 if there exists a path $\hat{\pi} \in \hat{\Pi}_{p, S}^D$ such that $f(\hat{\pi}(i)) < t$ for all $i \in \{0, 1, \dots, |\hat{\pi}| - 1\}$. Such a path exists iff $t \geq \min_{\hat{\pi}_1 \in \hat{\Pi}_{p, S}^D} \max_{i \in \{0, 1, \dots, k\}} f(\hat{\pi}_1(i))$. The probability of this event occurring is $P(A) = 1 - \min_{\hat{\pi}_1 \in \hat{\Pi}_{p, S}^D} \max_{i \in \{0, 1, \dots, k\}} f(\hat{\pi}_1(i))$.

Next, we study the probability $P(B)$ that p and q_2 belong to the same component of $T_{\hat{I}, t}$, and that $f(p)$ and $f(q_2)$ are both greater than or equal to t . This is true for some q_2 if there exists a path $\hat{\pi} \in \hat{\Pi}_{p, S}^D$ such that $f(\hat{\pi}(i)) \geq t$ for all $i \in \{0, 1, \dots, |\hat{\pi}| - 1\}$. Such a path exists iff $t < \max_{\hat{\pi}_2 \in \hat{\Pi}_{p, S}^D} \min_{i \in \{0, 1, \dots, k\}} f(\hat{\pi}_2(i))$. The probability of the event B occurring is $P(B) = \max_{\hat{\pi}_2 \in \hat{\Pi}_{p, S}^D} \min_{i \in \{0, 1, \dots, k\}} f(\hat{\pi}_2(i))$.

From eq. (16), it thus follows that

$$B\hat{M}D(p, S) = \min_{\hat{\pi}_1 \in \hat{\Pi}_{p, S}^D} \max_{i \in \{0, 1, \dots, k\}} f(\hat{\pi}_1(i)) - \max_{\hat{\pi}_2 \in \hat{\Pi}_{p, S}^D} \min_{i \in \{0, 1, \dots, k\}} f(\hat{\pi}_2(i)) = \hat{\varphi}(p, S). \quad (17)$$

□

A proof of Theorem 4 was previously provided by Zhang and Sclaroff [6].

4 Computing distance transforms for the discrete BMD

Given a discrete image $\hat{I} = (\hat{D}, f)$ and a set of *seed pixels* $S \subseteq \hat{D}$, the *distance transform* for the discrete BMD is a map assigning to each pixel $p \in \hat{D}$ the value $BMD(p, S)$, i.e., each pixel is assigned the discrete boolean map distance to the set S . In this section, we study various methods for computing distance transforms for the discrete BMD.

4.1 Monte Carlo approximation

From the definition of the BMD, it is straightforward to devise a Monte Carlo algorithm for approximating the BMD distance transform by iteratively selecting a random threshold, performing thresholding, and using a flood-fill operation to find the set of pixels connected to at least one seed-point. As the number of iterations increases, the relative frequency with which each pixel belongs to the complement of this set approaches the correct distance transform value.

4.2 The Zhang-Sclaroff algorithm

Assume that all intensities present in a given image can be written as i/k , for some fixed integer k and some i in the set $\{1, 2, \dots, k\}$. This situation occurs in practice if we, e.g., remap an image with integer intensity values to the range $[0, 1]$. If gray levels are stored as 8-bit integers, for example, we can take $k = 256$. Then the algorithm proposed by Zhang and Sclaroff for calculating Boolean Map Saliency [6] can be used directly for computing the exact BMD distance transform from any set of seed-points. Pseudo-code for this algorithm is listed in Algorithm 1.

In this algorithm each iteration of the **foreach**-loop requires $\mathcal{O}(n)$ operations, where n is the number of image pixels. Thus, the entire algorithm terminates in $\mathcal{O}(nk)$ operations which, since k can be considered a constant, equals $\mathcal{O}(n)$ operations.

Algorithm 1: The Zhang-Sclaroff algorithm for computing the discrete BMD distance transform.

Input: An image I , a set of seed-points S , an integer k
Output: Distance transform D

- 1 Set $D(p) = 0$ for all pixels p in I ;
- 2 **foreach** $i \in \{1, 2, \dots, k\}$ **do**
- 3 Set $B \leftarrow T_{I,i/k}$;
- 4 Perform a flood-fill operation to identify the set of pixels belonging to the same component as at least one seed-point in B ;
- 5 Increase D by $1/(k)$ for all pixels not in this set;
- 6 **end**

4.3 Dijkstra's algorithm

As shown in Section 3.1, the discrete BMD can be written as the difference between two functions:

$$\min_{\hat{\pi} \in \hat{\Pi}_{p,S}^D} \left(\max_{i \in \{0,1,\dots,k\}} I(\hat{\pi}(i)) \right) \quad (18)$$




Fig. 1. Left: “Cameraman” image used in the experiments. The location of the single seed-point is indicated in green. Right: The corresponding BMD distance transform.




Fig. 2. Left: Approximation error of Monte Carlo estimation as a function of the number of samples. Right: Empirical comparison of running time for the Zhang-Sclaroff algorithm and Dijkstra’s algorithm.

and

$$\max_{\hat{\pi} \in \hat{I}_{p,S}^D} \left(\min_{i \in \{0,1,\dots,k\}} I(\hat{\pi}(i)) \right). \quad (19)$$

Both of these functions are *path based distance functions*, and are *smooth* in the sense defined by Falcão et al. [2]. Therefore, distance transforms for each term can be computed in $\mathcal{O}(n \log n)$ operations using the *image foresting transform* [2], a generalization of Dijkstra’s algorithm. In the case where the magnitude of the set of all intensities present in the image is bounded by a fixed integer, this can further be reduced to $\mathcal{O}(n)$ operations [2].

4.4 Empirical comparison of running time

In this section, we perform an empirical comparison of the algorithms described above.

First, we consider the Monte Carlo approximation method. This algorithm differs from the others in that it only produces approximate results. The error of the approximation decreases as the number of iterations is increased, but this also increases the computation time. We are thus interested in investigating the trade-off between number of iterations and approximation error. To this end, we use the Monte Carlo approximation method to compute an approximate BMD distance transform of the “Cameraman” image shown in Fig. 1, from the single seed-point indicated in the figure for a varying number of iterations. Each result was compared to the true distance transform, computed using Dijkstra’s algorithm. Fig. 2 (left) shows the average error per pixel as a function of the number of iterations. Due to the stochastic nature of the algorithm, the mean error is itself noisy. To increase clarity the figure therefore shows, for each number of samples, the average error obtained when repeating the experiment 20 times. As Fig. 2 (left) shows, a large number of samples is required to obtain an accurate approximation.

The Zhang-Sclaroff algorithm and the Dijkstra based method both have linear time complexity, but the constants involved differ substantially. To compare the algorithms empirically we calculated distance transforms for the “Cameraman” image shown in Fig. 1, scaled to various sizes using bi-cubic interpolation, using both algorithms. The gray-levels in this image are stored as 8-bit integers, so we take $k = 256$ for the Zhang-Sclaroff algorithm. In all cases, a single seed-point was placed at the top left corner of the image. The resulting running times are shown in Fig. 2 (right). Both algorithms show the expected linear dependence between image size and running time, but the approach based on Dijkstra’s algorithm is faster by about a factor 30-40.

5 Extension to multi-channel images

To extend the BMD to multi-channel images, we consider the following procedure for creating a boolean map:

1. Randomly select one of the image channels according to some probability distribution over the set of image channels.
2. Randomly select a threshold from a uniform distribution over $[0, 1]$ and threshold the selected image channel at this value.

The multi-channel BMD between two pixels in an image with m channels is then defined as the probability that they belong to different components of the boolean map obtained by the above procedure. This probability is given by

$$BMD(p, q) = w_1 BMD_1(p, q) + w_2 BMD_2(p, q) + \dots + w_m BMD_m(p, q), \quad (20)$$



Fig. 3. Top left: “Flower” image (from Rhemann et al. [3]) with seedpoints overlaid in white. Top right: Gray-scale image. Bottom left: BMD distance transform of the color image, after transformation to the CIEL*a*b* color space. Bottom right: BMD distance transform of the gray-scale image. The values of both distance transforms have been scaled for display purposes. Compared to the single-channel distance transform, the multi-channel BMD distance transform better captures the contrast between the flower and the background.

where w_i denotes the probability of selecting channel i and BMD_i denotes the single channel BMD defined on channel i . For example, we may chose $w_i = 1/m$ for all $i \in \{1, 2, \dots, m\}$. Note that the above result applies to both the continuous and discrete BMD. In the discrete case, this means that we can compute BMD distance transforms for multi-channel images by computing the single channel BMD on each channel, and forming a weighted average of the results.

An illustration of computing a multi-channel BMD distance transform on a color image is shown in Fig. 3.

6 Conclusion

We have introduced the *Boolean map distance* (BMD), a pseudo-metric that measures the distance between elements in an image based on the probability that they belong to different components after thresholding the image by a randomly selected value. Formal definitions of the BMD have been given in both the continuous and discrete settings. The equivalence between the BMD and the φ mapping proposed by Strand et al. was previously shown in the discrete case [6]. We have extended this proof to also cover the continuous case, thereby further strengthening the connection between the BMD and the MBD.

We have summarized available algorithms for computing distance transforms for the discrete BMD. From the empirical comparison, we conclude that the

Monte Carlo approximation method is not suitable for practical applications, given the existence of efficient exact algorithms. In the comparison between exact algorithms, we found that the approach based on Dijkstra’s algorithm was faster than the Zhang-Scaroff algorithm by an order of magnitude for calculating distance transforms on gray-scale images stored using 8-bit pixels. With increased color depth, the difference in computation time will increase.

Various aspects of the ideas presented have been explored in previous publications [5, 6, 4, 1]. The BMD, however, has not to our knowledge previously been proposed as a distance function in its own right. By compiling and extending ideas previously scattered across multiple publications, we hope to highlight the BMD as a valuable distance function for image processing tasks.

References

1. Krzysztof Chris Ciesielski, Robin Strand, Filip Malmberg, and Punam K. Saha. Efficient algorithm for finding the exact minimum barrier distance. *Computer Vision and Image Understanding*, 123:53–64, 2014.
2. Alexandre X Falcão, Jorge Stolfi, and Roberto de Alencar Lotufo. The image foresting transform: Theory, algorithms, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(1):19–29, 2004.
3. Christoph Rhemann, Carsten Rother, Jue Wang, Margrit Gelautz, Pushmeet Kohli, and Pamela Rott. A perceptually motivated online benchmark for image matting. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1826–1833. IEEE, 2009.
4. Robin Strand, Krzysztof Chris Ciesielski, Filip Malmberg, and Punam K Saha. The minimum barrier distance. *Computer Vision and Image Understanding*, 117(4):429–437, 2013.
5. Jianming Zhang and Stan Sclaroff. Saliency detection: A boolean map approach. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 153–160, 2013.
6. Jianming Zhang and Stan Sclaroff. Exploiting surroundedness for saliency detection: a boolean map approach. *IEEE transactions on pattern analysis and machine intelligence*, 38(5):889–902, 2016.