

Distance Transformations Based on Ordered Weighted Averaging Operators

C. Lopez-Molina, L. De Miguel, S. Iglesias-Rey, H. Bustince
 Universidad Publica de Navarra, Spain
carlos.lopez@unavarra.es

B. De Baets
 KERMIT, Universiteit Gent, Belgium
bernard.debaets@ugent.be

Abstract

Binary image comparison has been a study subject for a long time, often rendering in context-specific solutions that depend upon the type of visual contents in the binary images. Distance transformations have been a recurrent tool in many of such solutions. The literature contains works on the generation and definition of distance transformations, but also on how to make a sensible use of their results. In this work, we attempt to solve one of the most critical problems in the application of distance transformations to real problems: their oversensitivity to certain spurious pixels which, even if having a minimal visual impact in the binary images to be compared, may have a severe impact on their distance transforms. With this aim, we combine distance transformations with Ordered Weighted Averaging (OWA) operators, a well-known information fusion tool from Fuzzy Set Theory.

1. Introduction

Quantitative comparison has been a key topic in different scientific fields. For example, it is crucial to mathematical psychology [1, 2], and has proven critical in very different data science procedures, e.g. Support Vector Machines (SVMs), optimization and clustering [3]. In fact, almost all data science procedures rely on data comparison in some manner.

The comparison of individual elements within a given universe, often referred to as 1-to-1 comparison, is deeply studied in the literature. Often, this study has led to operators that fit within the axiomatic definition of metric, but similar classes of operators (as quasi- or pseudo-metrics) have also had a relevant role. In fact, the fitness of metrics is often revisited, despite the mathematical and computational convenience of their axiomatic constraints. For example, from the perspective of mathematical psychology [4, 5], the triangle inequality has been often disregarded [1]. Nevertheless, the use of 1-to-1 comparison operators

is both recurrent and well founded in different mathematical theories [6].

The literature is significantly more scarce when comparisons are made in configurations other than 1-to-1. Such is the case of 1-to- N or N -to- M comparisons, that often lead to problem-specific solutions. Usually, 1-to- N (e.g. point-to-set) comparison is reduced to the minimum (resp. maximum) distance (similarity) between the point and any element in the set. While some exceptions hold [7, 8], most 1-to- N have mathematical foundations significantly shallower than those of 1-to-1 comparison operators. N -to- M comparison operators are, as well, far less frequently studied than 1-to-1 ones. Exceptions to this are the Earth Mover's metric [9] or the Hausdorff metric [10]. Still, it is very common to produce context-specific operators for given problems.

Binary image processing is one of the fields in which N -to- M comparison operators are recurrently needed. Binary images, while hardly useful to represent visual information, are critical to represent intermediate or final results in very different image processing tasks. Certain image segmentation tasks, as well as pattern recognition or boundary detection tasks, typically represent their final results as binary images. Hence, the comparison with ground truth, with the final goal of quantitative quality evaluation, is often based on binary image comparison operators [11]. As long as binary images can (and often must) be seen as subsets of positions, the comparison of binary images becomes a matter of N -to- M comparison in the universe of positions of the image [12].

A significant part of the effort devoted to the comparison of binary images has relied on Distance Transformations (DTs) to overcome the problem of set comparison with variable cardinality [13]. This is so because matching-based solutions (as the classification-based approaches in [14]) rely on pixel matching and counting, which can hardly succeed when the cardinality of the sets to be compared is different [14]. Hence, DTs became a recurrent

tool in the design of binary image comparison operators. Nevertheless, several studies illustrate how DTs have endemic problems in over-representing certain visual artefacts in the images despite their little relevance. This has sparked interest from researchers to solve the problem in both theoretical and practical approaches [15]. In this work, we combine state-of-the-art developments on DTs with Ordered Weighted Averaging (OWA) operators [16, 17], a recurrent tool for information fusion in the Fuzzy Set Theory [18]. Our goal is to produce DTs which are able to cope with the oversensitivity to visually irrelevant binary features in the images.

The remainder of this work is organized as follows. Section 2 recovers some concepts of interest for the rest of the manuscript. Section 3 covers the formulations and uses of DTs, while Section 4 illustrates our proposal for DTs using OWA operators. Our proposal is put to the test in the context of boundary image comparison in Section 5. Finally, Section 6 includes discussion and future lines.

2. Preliminaries

This section recalls some basic definitions of the concepts used hereafter.

Definition 1 Let $\mathbf{x} = (x_1, \dots, x_n)$, with $x_i \in \mathbb{R}$ for all $i \in \{1, \dots, n\}$. The power mean of \mathbf{x} is given by

$$G_q(\mathbf{x}) = \left(\frac{1}{n} \sum_{i=1}^n x_i^q \right)^{1/q},$$

with $q \in \mathbb{R}$.

The power mean generalizes well-known operators. For example, if $q = -\infty$ (resp., $q = \infty$), the power means becomes the minimum (resp., maximum) operator.

Among the different families of aggregation operators [18], OWA operators have become popular in recent years.

Definition 2 [16] A function $w : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is called an OWA operator of dimension n if there exists a weighing vector $\mathbf{h} = (h_1, \dots, h_n) \in [0, 1]^n$ with $\sum_i h_i = 1$, and such that

$$\Theta^{\mathbf{h}}(a_1, \dots, a_n) = \sum_{j=1}^n h_j b_j, \quad (1)$$

with b_j the j -th largest of the a_i , for any $(a_1, \dots, a_n) \in [0, 1]^n$.




Figure 1. Visual representation of OWA operators on $[0, 1]^2$ with different weighing vectors \mathbf{h} .

Although OWA operators were originally designed as mappings $\mathbb{R}^n \rightarrow \mathbb{R}^n$, they have been recurrently reformulated as mappings $[0, 1]^n \rightarrow [0, 1]^n$, so as to fit the definition of aggregation operators. In fact, OWA operators are general cases for very significant aggregation operators. For example, the OWA operator based on weighing vector $(1, 0, \dots, 0)$ (resp. $(0, \dots, 0, 1)$) is the minimum (resp. maximum) operator. Also, the vector $(\frac{1}{n}, \dots, \frac{1}{n})$ yields the arithmetic mean operator. Figure 1 displays some paradigmatic OWA operators on $[0, 1]^2$.

Definition 3 A function $g : U \times U \rightarrow \mathbb{R}^+$ is called a metric (or distance function) on a universe U if and only if it satisfies the following properties, for any $a, b, c \in U$:

- (i) Identity of indiscernibles: $g(a, b) = 0$ iff $a = b$;
- (ii) Symmetry: $g(a, b) = g(b, a)$;
- (iii) Triangle inequality: $g(a, c) \leq g(a, b) + g(b, c)$.

Definition 4 A function $g : U \times U \rightarrow \mathbb{R}^+$ is called a pseudo-metric on a universe U if and only if it satisfies the following properties, for any $a, b, c \in U$:

- (i) $g(a, a) = 0$ for all $a \in U$;
- (ii) Symmetry: $g(a, b) = g(b, a)$;
- (iii) Triangle inequality: $g(a, c) \leq g(a, b) + g(b, c)$.




Figure 2. Distance transforms of different binary images. The three leftmost images are synthetic examples, while the two rightmost images are taken from the BSDS500 [19]. The upper row displays the binary images, with 1-valued pixels represented in black for better visualization. The lower rows display the results by different distance transformations. The notation for the transformations is that in Section 3.1. Note that each distance transformation uses its own numerical scale.

3. Distance transformations and binary image comparison

In the remainder of this work, we represent images as functions $f : \Omega \rightarrow T$, with $\Omega = \{1, \dots, R\} \times \{1, \dots, C\}$ the set of positions and T the set of tones. The set of images with tonal palette T is denoted \mathbb{I}_T .

3.1. Distance transformations

Definition 5 The Distance Transform (DT) of a binary image $I \in \mathbb{I}_{\{0,1\}}$, for some metric m on Ω , is the image $\mathcal{T}_m[I] \in \mathbb{I}_{\mathbb{R}^+}$ given by

$$\mathcal{T}_m[I](p) = \min_{p' \in I} m(p, p'), \quad (2)$$

for all $p \in \Omega$.

In the remainder of this work, generic metrics will be referred to as m , while the Euclidean metric will be represented as d .

Distance transformations provide a powerful tool to ease the comparison of binary images by representing them into \mathbb{R}^+ -valued images. Regardless of the number of featured pixels in any two binary images, their transforms are always in $\mathbb{I}_{\mathbb{R}^+}$. Figure 2 contains the visualization of distance transforms for a set of images. In Fig. 2, the top row contains the original images, and the middle row displays the transforms using $\mathcal{T}_d[I]$, with d standing for the Euclidean metric. We observe how DTs allow for a representation based on the distribution of the 1-valued pixels in the original images. Also, it is noticeable how DTs are very sensitive to small responses. The second and third rightmost columns in Fig. 2 offer a visual example of this. Both columns

feature almost identical images, the only exception being visually insignificant path of 2×2 pixels in the upper right area. Still, as it can be observed in the middle row of that same figure, such pixels have a deep impact on the distance transform of the image, as they become the closest 1-valued pixels for a large area in the image.

A recurrent solution to oversensitivity in DTs is the adjustment of the metric m . As DTs can be based on any metric on Ω , a smart selection of such metric may reduce the impact of spurious responses in the transform.

Let m be a metric on U . The bounded version of m , on a universe U , is given by

$$m_t(a, b) = \min(t, m(a, b)), \quad (3)$$

for any $a, b \in U$, with $t \in \mathbb{R}^+$, $t > 0$. A distance transformation based on a bounded metric is referred to as *bounded distance transformation*. Note that, since $m = m_\infty$, the subindex is ignored for $t = \infty$.

The lowest row in Fig. 2 contains the results for the distance transformations of the images in the top row using $\mathcal{T}_{d_{20}}[I]$. The impact of the bounded metric is severe: spurious responses have a significant reduced impact on the transformation; also, this is associated to a far less smooth representation of the binary features in the image.

While the bounded metric has been the primary tool to prevent oversensitivity to spurious responses [14, 20], it has not been the only one. In [15], Brunet and Sills proposed to generalize the distance transformation by replacing the minimum operator by the power mean; that is, computing the value at each pixel of the distance transform as the power mean of the distances to all 1-valued pixels in the image, instead of yielding the minimum of such distances.

Definition 6 [15] *The Generalized Distance Transform of a binary image $I \in \mathbb{I}_{\{0,1\}}$, for some metric m on Ω , is the image $\mathcal{P}_m^q[I] \in \mathbb{I}_{\mathbb{R}^+}$ so that*

$$\begin{aligned} \mathcal{P}_m^q[I](p) &= G_q(w(m(p, p'))) \\ &= \left(\frac{1}{|I|} \sum_{p' \in I} w(m(p, p'))^q \right)^{1/q}, \end{aligned} \quad (4)$$

for all $p \in \Omega$, with $q \in [-\infty, \infty]$ and w a concave, increasing function such that $w(0) = 0$.

There is a list of remarks to be pointed out in this definition. Firstly, since the concave mapping of a metric m is still a metric, removing the function w saves a parameter while preserving the flexibility in the operator. In the remainder of the work, w will be unused. Secondly, as stated in [21] the name *generalized distance transform* was used, prior to [15], by Felzenszwalb and Huttenlocher [22]. In this work, we adhere to referring as *Power Distance Transformation (PDT)* to the transformation introduced in [15]. Thirdly, the original proposal by Brunet and Sills [15] imposes $q < 0$. The reason is that, if $q < 0$, then $\mathcal{P}_m^q[I](p) = 0$ iff $p \in I$. This ensures \mathcal{P}_m^q to be invertible, impacting the properties of subsequent operators.

The behaviour of PDTs, as compared to that of DTs, can be seen in Fig. 3. In this figure, the four lowest rows contain the results of PDTs using different configurations with $q \in \{-1, 1\}$ and $m \in \{d, d_{20}\}$. It is clear how the PDTs yield 0-valued pixels if and only if $q < 0$, making them invertible. Also, we observe that, similarly to what happened for the original DT, the Euclidean metric d yields larger values and more intense smoothing. This is specially significant in the case of $\mathcal{P}_d^1[I]$, which produces almost unrecognizable representations of the binary images.

3.2. Image comparison using Baddeley's Delta Metric

Binary images can be dually interpreted as mappings $\Omega \rightarrow \{0, 1\}$ or as subsets of positions (in Ω). Elaborating on this dual interpretation, binary image comparison is often carried out using the Hausdorff metric [23]. Let \mathcal{T}_m be a distance transformation on binary images based on some metric m . The Hausdorff

distance between two images, $A, B \in \mathbb{I}_{\{0,1\}}$ is given by

$$\text{HM}(A, B) = \max \left(\max_{a \in A} \mathcal{T}_m[B](a), \max_{b \in B} \mathcal{T}_m[A](b) \right). \quad (5)$$

The Hausdorff metric, while simple and sound, also features some problems. Most of such problems stem from the fact that the value yielded for any two images depends upon just two pixels. While this is positive for certain tasks, e.g. those in which noise and spurious responses are non-existing, it also makes the comparison unstable for most applications.

The Hausdorff metric has been further tuned for specific goals. For example, Dubuisson and Jain [24], proposed a variation of this metric to create asymmetric similarity measures, in combination with bidirectional aggregation operators. In total, Dubuisson and Jain designed 24 comparison operators evolved from the Hausdorff metric, and tested them in the context of delineated object recognition. This work was continued by Takacs and Wechsler [25, 26], who proposed a different tuning to adapt the results in [24] to the context of boundary-based face recognition. Also Baudrier *et al.* [10] presented a variation of the Hausdorff metric by combining the ideas in [24] with a local, sliding window-based, analysis of subregions of the image.

The most relevant proposal evolving the Hausdorff metric is Baddeley's Delta Metric (BDM) [13, 27]. In these works, Baddeley introduced a metric for binary images that aims at increasing the flexibility of the Hausdorff metric, while preserving the metric axioms.

Baddeley's reformulates the Hausdorff metric as

$$\text{HM}(A, B) = \max_{p \in \Omega} |\mathcal{T}_m[A](p) - \mathcal{T}_m[B](p)|, \quad (6)$$

which is equivalent to Eq. (5). Baddeley intends to exploit the comparison of $\mathcal{T}_m[A]$ and $\mathcal{T}_m[B]$, as representatives of A and B , and to produce an interpretation richer than that by max.

Let $A, B \in \mathbb{I}_{\{0,1\}}$ be two binary images on Ω , and let m be some metric on Ω . The distance between them, in terms of BDM, is given by

$$\Delta^k(A, B) = \left[\frac{1}{|\Omega|} \sum_{p \in \Omega} |w(\mathcal{T}_m[A](p)) - w(\mathcal{T}_m[B](p))|^k \right]^{\frac{1}{k}}, \quad (7)$$

where $w : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a concave function with $w(x) = 0$ iff $x = 0$, \mathcal{T}_m is a distance transformation and $k \in \mathbb{R}^+$.

The role of the function w is, again, potentially redundant. As explained for the PDT, any function $w \circ m$

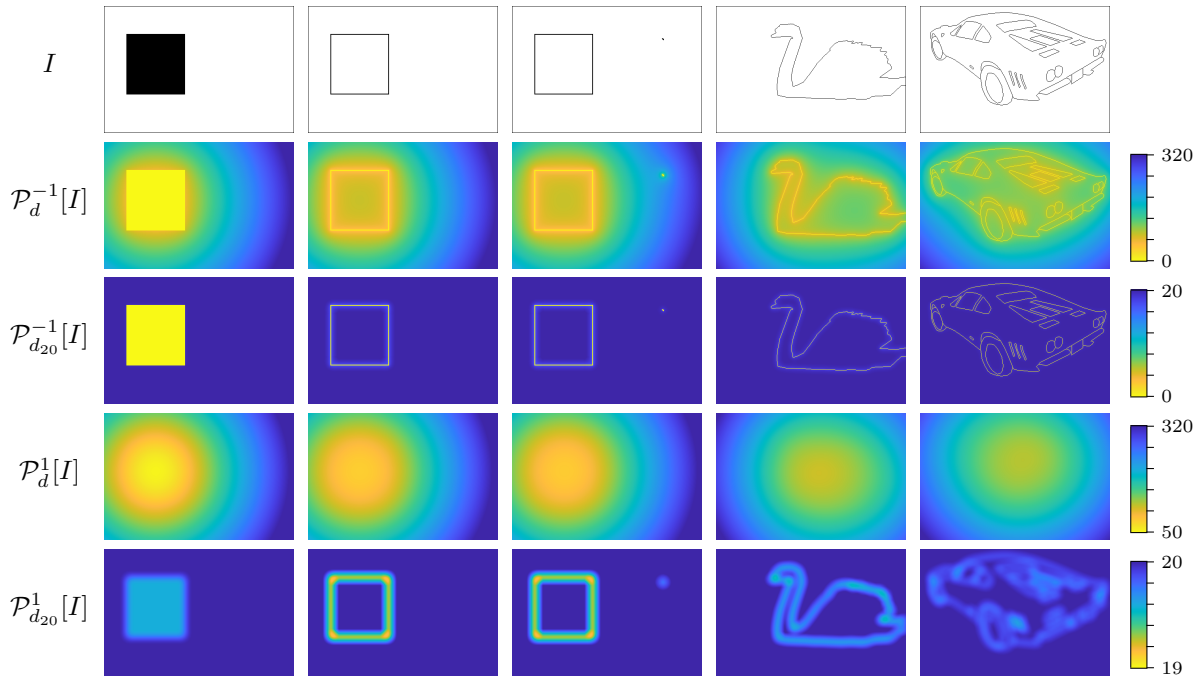


Figure 3. Distance transforms of different binary images. The three leftmost images are synthetic examples, while the two rightmost images are taken from the BSDS500 [19]. The upper row displays the binary images, with 1-valued pixels represented in black for better visualization. The lower rows display the results by different distance transformations. The notation for the transformations is that in Section 3.1. Note that each distance transformation uses its own numerical scale.

is in fact a metric. Hence, any effect obtained by setting w , can be equivalently inferred by modifying the metric m . In this work, we set $w(x) = x$, giving w no role in the metric.

Note that, in Eq. (7), DTs can be replaced by PDTs. BDM will remain a metric as long as $q < 0$, as proven in [15]; for $q > 0$, the BDM becomes a pseudometric.

Baddeley’s Delta Metric has been used for its original purpose in different contexts and tasks [28, 29], but has also been evolved to compare items other than binary images, e.g. grayscale images [30, 31], or hyperspectral signatures [32].

4. OWA-based distance transformations

One of the foreseeable problems in the PDT, as presented both in [15] and [21], is the fact that each value in a transform is dependent upon the distance to all 1-valued pixels in the original image. While depending on 1 pixel (as in the original DT) made the representation prone to oversensitivity, we believe that considering the distance to all pixels might produce oversmoothed representations of images, unable to characterize small details or regions. An intermediate option (considering the n closest 1-valued pixels) might

be able to solve both problems.

We propose in this work a reformulation of distance transformations so as to solve the problems in both the original DT and the PDT. Our proposal is based on three principles:

- *Interpretability*- The distance transform of an image should be an understandable representation of the information in the original image. In this sense, the 1-valued regions in the original image must have a clear impact in the transformation when relevant; Also, they should have a light or non-existing impact when spurious or irrelevant.
- *Robustness*- The transformation must be as robust as possible to changes in its parameterization, as well as to changes in the binary image. In this sense, the addition or removal of a small set of 1-valued pixels in the binary image should not modify the transform dramatically.
- *Intuitive parameterization*- The idea of spurious information in a binary image is often associated to small, noise-related 1-valued (or 0-valued) pixels. However, small segments might be interesting for some tasks. In fact, different tasks

are likely to induce different definitions of how short (resp. small) a segment (resp. region) needs to be in order to be labelled as spurious. Hence, the reformulation of the distance transformation should allow a straightforward configuration of the boundary between a spurious and a relevant object in the image.

The PDT does not totally fit the aforementioned principles. While robust, the parameter q in the PDT fails to achieve the goals of interpretability and intuitive parameterization. The parameter q modulates the value of the distance transform at each pixel, from the minimum ($q = -\infty$) to the maximum distance to a 1-valued pixel ($q = \infty$), passing by, e.g., the arithmetic mean ($q = 1$). Hence, it is a non-linear modulator of the orness-andness balance, a well-known property of information fusion operators [33, 34]. The parameter q does not, however, relate to the images themselves, or to the visual characteristics of the objects in them, and hence the configuration is far from trivial. When computing the PDT at a pixel, all other pixels have the same influence in the final result, what seems counter-intuitive. Also, the instabilities when $q \approx 0$ inspire a sense of delicacy in the process of parameter setting. Nevertheless, the PDT has been proven useful for both region- and boundary-oriented binary image comparison [21]. We believe that the same ideas could be embodied in a different manner in order to ease the study and application of distance transformations.

We propose to use OWA operators, instead of the power mean, for the aggregation of the pixel-to-pixel distances at each pixel. The OWA operators can recover both the original distance transformation (based on the minimum operator), as well as both extreme cases in the PDT (minimum and maximum operators). However, it also enables a configuration which is more evidently related to the visual facts in the images. Since OWA operators apply weights on the ordered arguments, we can easily state the number of (nearest) pixels which will be taken into account for each pixel. For example, the distance transform at each pixel could be the arithmetic mean to the n nearest 1-valued pixel, or to the n furthest ones. Also, it could be configured according to some percentage of the closest 1-valued pixels. Moreover, the fusion of the distances to such n pixels may not be based on the arithmetic mean. OWA operators allow for an easy configuration of the weights each argument receives.

Definition 7 *The Ordered Weighted Distance Transform of a binary image $I \in \mathbb{I}_{\{0,1\}}$, for some metric m on Ω , is the image $\mathcal{W}_m^{\mathbf{h}}[I] \in \mathbb{I}_{\mathbb{R}^+}$ so that*

$$\mathcal{W}_m^{\mathbf{h}}[I](p) = \Theta_{p' \in I}^{\mathbf{h}} m(p, p'), \quad (8)$$

for all $p \in \Omega$, with $\mathbf{h} \in [0, 1]^{|I|}$.

The Ordered Weighted Distance Transformation (OWDT) is a general form of the original DT, which can be recovered using the weighing vector $\mathbf{h} = (1, 0, \dots, 0)$. It also recovers both extreme cases of the PDT ($q \in \{-\infty, \infty\}$), as well as some other special cases (as $q = 1$, which is equivalent to the OWDT with $\mathbf{h} = (\frac{1}{n}, \dots, \frac{1}{n})$).

The OWDT is not an invertible operation, since it is not guaranteed that $\mathcal{W}_m^{\mathbf{h}}[I](p) == 0$ if and only if $p \in I$. In fact, $\mathcal{W}_m^{\mathbf{h}}[I](p) == 0$ only in extreme cases, e.g. if $\mathbf{h} = (1, 0, \dots, 0)$ or if p is the single 1-valued pixel in I . Hence, when the original DT is replaced by the OWDT in Baddeley's Delta Metric, it will remain a pseudometric (but not a metric).

Our proposal is mildly related to other works in literature. For example, it can be related to the k -distances [8], a relatively popular notion imported from geometry. This comparison measure yields the distance from a point to a set as the arithmetic mean of the individual distances to the k nearest elements in that set. Hence, it can be easily recovered by using weighing vectors within the definition of OWA operators. Also, our proposal is somehow similar, in both general goals and inspiration, to that by Öfverstedt et al. [35]. In this paper, the authors study the positive (and negative) noise in binary images, and audit the impact it has in the results by the original DT. However, the means and tools developed in [35] to overcome the problem are different from the ones proposed in this work.

Figure 4 displays the OWDT of the images in the upper row of Fig. 2. We consider exclusively triangular inferior weighing vectors; These vectors, which we denote \mathbf{t}_n , assign linearly increasing weights to the n smallest values to be fused. Hence, the j -th smallest value has a weight

$$t_j = \max \left(0, \frac{n - j + 1}{\sum_{j=1}^n j} \right), \quad (9)$$

satisfying that $\sum t_n = 1$ for any $n \in \mathbb{N}$.

By using triangular inferior weighing vectors, we establish the number of pixels to be taken into account (by setting n). Also, we also maintain the operator robust against variations of n , since the increase or decrease of n will render in small variations in the weight of each argument.




Figure 4. Ordered weighted distance transforms (OWDTs) of the binary images in Fig. 2. Each row displays the transforms by a different operators, featuring different metrics m and weighing vectors t_n . Note that each distance transformation uses its own numerical scale.

5. Experiments

This section is focused on illustrating the behaviour of OWDT, when combined with Baddeley’s Delta Metric, in the context of boundary detection comparison. Boundary images are binary images with the delineation of relevant objects/structures in the images. Boundary comparison has been recurrently tackled from the perspective of binary image comparison [11, 20].

Binary image comparison, either using metrics or any other class of functions, is a recurrent task in image processing. For example, it plays a role application in object tracking [36] or object recognition [37]. We focus this experiment in context of quality evaluation, specifically restricting to boundary detection and image segmentation. An automatic method for boundary detection is evaluated according to the similarity of its output to that by humans on the same task. The comparison operator plays, hence, a critical role in the process [11, 14]. In fact, one of the primary goals by Baddeley [13] was to produce an operator for quality evaluation.

The comparison of boundary images is an ongoing field of research, and connects with similar tasks as silhouette matching and graph matching [38]. Despite the diversity of comparison operators for the task, it is unclear how to rate or quantify their performance. That is, to evaluate how well a comparison operator performs, or to select which out of a pool of candidates is the best performing. Such questions belong to the field of metaquality: strategies to measure or enforce the quality

of the evaluation methods. Ideally, if humans could produce solid quantitative evaluations of the similarity of two images, the comparison operators shall be evaluated in terms of how correlated its results are to the quantitative evaluations by humans. However, humans rarely produce such inputs in a reliable manner. Hence, alternative strategies need to be drawn to evaluate comparison operators.

Metaquality often relies on psychological priors that could be applied to human behaviour in image comparison or recognition [20, 39]. In our experiments, we focus on the data distributed within the BSDS500 [19], a popular dataset for image segmentation and boundary detection. The BSDS500 dataset contains 500 images, each of them associated to 5 to 7 hand-made ground truth images. The ground truth images are produced by different humans, and hence feature a significant number of discrepancies among them. Nevertheless, any human could cluster all boundary images according to the image after which they were produced. Otherwise said, any human is able to discriminate when two boundary images are produced by humans either on the same scene or on different ones. This is illustrated in Fig. 5, which contains the ground truth images associated to two different original images in the BSDS500. In Fig. 5 it is clear how humans make different interpretations of what is important and unimportant in each scene. Still, humans would be able to cluster images in the same class, and discriminate images in different classes. This fact, apart from being an explicit example of

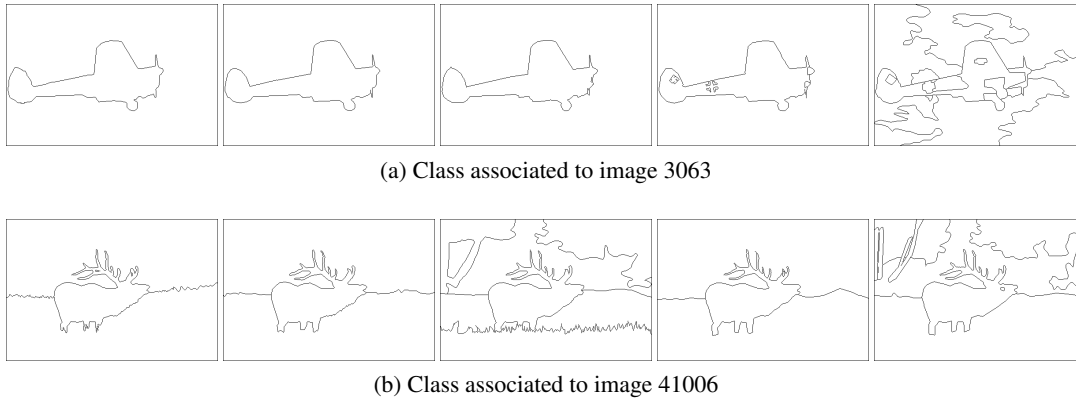


Figure 5. Hand-made (ground truth) boundary images extracted from the BSDS500. Each row contains images created from humans from the same original image. It can be observed how, despite the heterogeneity in the hand-made delineations, all images can be easily recognized as representing the same original scene.

human ability to combine bipolar information (in this case, coincidences and divergences between images), is also our accepted prior for the design of a metaquality strategy. Our meta-quality strategy is based on the idea that binary image comparison measures should replicate this human behaviour; that is, comparison measures should yield larger *dissimilarities* in the comparison of inter-class pairs of images than in the comparison of intra-class ones. This section intends to measure how the combination of OWDTs and BDM, under different parameter settings, performs in recovering the human class-discriminative behaviour.

For some given operator, we analyze the distributions of the inter- and intra-class comparisons in the BSDS500 Test Set. In this regard, we use BDM with both original DTs and different OWDTs. Recall that the BSDS500 Test Set contains 1063 images distributed in 200 classes, accounting for more than 10^6 comparisons.

Given the available space, we present a limited number of parameter settings. Specifically, we consider the original DT, together with two parameterizations of the OWDT ($\mathcal{W}_m^{t_{50}}$ and $\mathcal{W}_m^{t_{100}}$). Also, we consider two different distances on Ω , $m \in \{d, d_{50}\}$, so as to include both bounded and unbounded metrics.

Figure 6 contains the distribution of the values by different versions of BDM on inter- and intra-class pairs of images. Each row is devoted to a single metric on Ω (either d or d_{50}). Each plot also displays the accuracy (Acc) in the discrimination of both distributions for each threshold. The Acc is measured in the range $[0.5, 1]$, and is scaled on the rightmost axis of each plot. Also, note that $\text{Acc} = 1$ is reached if and only if the distributions are non-overlapping.

It is relevant to note that the all-vs-all analysis of the BSDS500 leads to around 4500 intra-class comparisons, and over 10^6 inter-class ones. Hence, the distributions

in Fig. 6 are scaled in percentual terms. The accuracy (Acc) is also computed w.r.t. normalized distributions for each class, to avoid problems due to class imbalance. Otherwise, the influence of the intra-class comparisons would be negligible.

The first noticeable fact in Fig. 6 is that the substitution of the original DT (leftmost column) by OWDTs (center and rightmost columns) results in a light improvement. This indicates that OWDTs are generally better in representing the boundary images than DTs; Also, that the improvement that can be gained in this setting is subtle (circa 0.02-0.03 in terms of Acc). There are, nonetheless, two relevant remarks to be made on this perception. Firstly, we have that the dataset used for the comparison is composed exclusively by hand-labelled images. Unlike computer-based algorithms, Humans do not normally include spurious responses or noise in the images they label. While being the most relevant dataset for metaquality evaluation [39], the BSDS500 does not suffer excessively from the kind of spurious responses to which original DTs are oversensitive. Secondly, larger, more ambitious experiments should be designed to fully understand how the OWA operators used to fusion information interact with the metrics in Ω . While it is clear why the original DT benefits from replacing d_{50} by d , it remains obscure how the OWDT, specifically the weighing vector \mathbf{h} , should adapt to such choices in metrics.

As a general conclusion, we can state that OWDTs produce sensible results, and seem to perform a better representation of binary images than the original DT. This holds for the qualitative point in view illustrated in Fig. 4, as well as from the quantitative point of view in Fig. 6. Such findings are, however, affected by limited space and scope of the experimental setup.

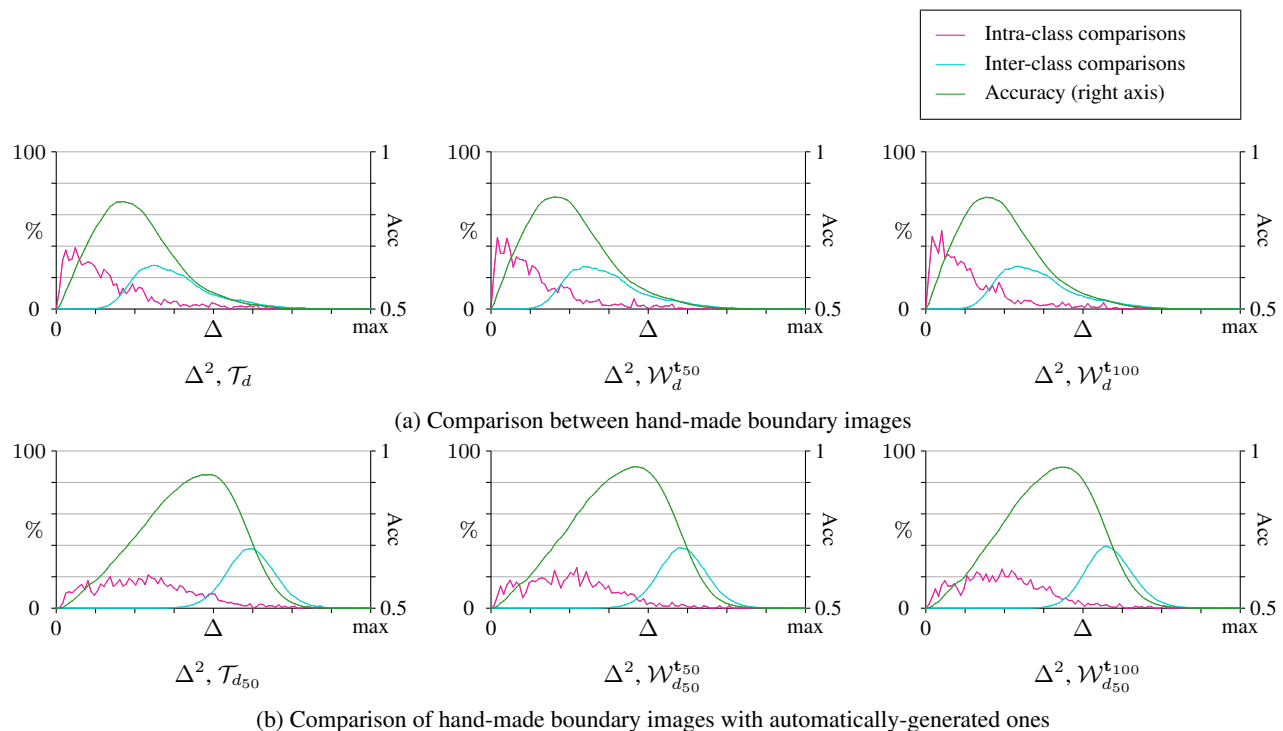


Figure 6. Distributions of distances for inter- and intra-class pairs of ground truth images in the BSDS500 Test set, using different configurations of BDM. The distributions have been configured with 100 bins.

6. Conclusions

In this work, we present a novel approach to distance transformations. Specifically, we propose to use OWA operators to avoid excessive sensitivity to unimportant 1-valued features in binary images. This sensitivity is a key problem for tasks in binary image comparison, specially in the context of image recognition/retrieval and quantitative quality evaluation. We present the mathematical and practical foundations of our approach, as well as an experimental setup in which the novel distance transformations have been put to the test. While promising, the results are not conclusive, as long as the large number of parameters make it advisable to perform comprehensive experiments that also, shall include different types of binary images.

As future lines of work, it is critical to design strategies that help researchers and practitioners in the adaptation of the different distance transformations to specific problems. One of the key benefits in the original distance transformation is the simplicity, with no parameter other than the metric on the set of image positions. Hence, the increasing interest in distance transformations will hardly render into applied developments unless the complexity in their configurations is reduced.

Acknowledgments

The authors gratefully acknowledge the financial support of the Spanish Research Agency, project PID2019-108392GB-I00 (AEI/10.13039/501100011033), as well as that by Navarra de Servicios y Tecnologías (Nasertic).

References

- [1] A. Tversky and I. Gati, "Similarity, separability, and the triangle inequality," *Psychological review*, vol. 89, no. 2, p. 123, 1982.
- [2] H. Blanton and J. Jaccard, "Arbitrary metrics in psychology," *American Psychologist*, vol. 61, no. 1, p. 27, 2006.
- [3] B. Nguyen and B. De Baets, "Kernel-based distance metric learning for supervised k -means clustering," *IEEE Trans. on Neural Networks and Learning Systems*, vol. 30, no. 10, pp. 3084–3095, 2019.
- [4] J. M. Yearsley, A. Barque-Duran, E. Scerrati, J. A. Hampton, and E. M. Pothos, "The triangle inequality constraint in similarity judgments," *Progress in Biophysics and Molecular Biology*, vol. 130, pp. 26–32, 2017.
- [5] F. Jäkel, B. Schölkopf, and F. A. Wichmann, "Similarity, kernels, and the triangle inequality," *Journal of Mathematical Psychology*, vol. 52, no. 5, pp. 297–303, 2008.
- [6] V. Gregori, S. Morillas, and A. Sapena, "Examples of

- fuzzy metrics and applications,” *Fuzzy Sets and Systems*, vol. 170, no. 1, pp. 95–111, 2011.
- [7] O. Cuisenaire, “The euclidean k-distance transformation in arbitrary dimensions: a separable implementation,” in *IEEE International Conf. on Image Processing*, vol. 2, pp. II–658, 2005.
- [8] L. Guibas, D. Morozov, and Q. Mérigot, “Witnessed k-distance,” *Discrete & Computational Geometry*, vol. 49, no. 1, pp. 22–45, 2013.
- [9] Y. Rubner, C. Tomasi, and L. J. Guibas, “The earth mover’s distance as a metric for image retrieval,” *International Journal of Computer Vision*, vol. 40, no. 2, pp. 99–121, 2000.
- [10] E. Baudrier, G. Millon, F. Nicolier, and S. Ruan, “A new similarity measure using hausdorff distance map,” in *International Conf. on Image Processing*, vol. 1, pp. 669–672, 2004.
- [11] T. Peli and D. Malah, “A study of edge detection algorithms,” *Computer Graphics and Image Processing*, vol. 20, no. 1, pp. 1–21, 1982.
- [12] C. Lopez-Molina, C. Marco-Detchart, H. Bustince, and B. De Baets, “A survey on matching strategies for boundary image comparison and evaluation,” *Pattern Recognition*, vol. 115, p. 107883, 2021.
- [13] A. J. Baddeley, “An error metric for binary images,” in *Robust Computer Vision: Quality of Vision Algorithms* (W. Förstner and S. Ruwiedel, eds.), (Karlsruhe), pp. 59–78, Wichmann Verlag, 1992.
- [14] C. Lopez-Molina, B. De Baets, and H. Bustince, “Quantitative error measures for edge detection,” *Pattern Recognition*, vol. 46, no. 4, pp. 1125–1139, 2013.
- [15] D. Brunet and D. Sills, “A generalized distance transform: Theory and applications to weather analysis and forecasting,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 55, no. 3, pp. 1752–1764, 2016.
- [16] R. Yager, “On ordered weighted averaging aggregation operators in multicriteria decision making,” *IEEE Trans. on Systems, Man, and Cybernetics*, vol. 18, no. 1, pp. 183–190, 1988.
- [17] R. Yager, “Families of OWA Operators,” *Fuzzy Sets and Systems*, vol. 59, no. 2, pp. 125–148, 1993.
- [18] G. Beliakov, A. Pradera, and T. Calvo, *Aggregation Functions: A Guide for Practitioners*, vol. 221 of *Studies in Fuzziness and Soft Computing*. Springer, 2007.
- [19] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, “Contour detection and hierarchical image segmentation,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 898–916, 2011.
- [20] C. Lopez-Molina, B. De Baets, and H. Bustince, “Twofold consensus for boundary detection ground truth,” *Knowledge-Based Systems*, vol. 98, pp. 162–171, 2016.
- [21] C. Lopez-Molina, S. Iglesias-Rey, H. Bustince, and B. De Baets, “On the role of distance transformations in baddeley’s delta metric,” *Information Sciences*, vol. 569, pp. 479–495, 2021.
- [22] P. F. Felzenszwalb and D. P. Huttenlocher, “Distance transforms of sampled functions,” *Theory of computing*, vol. 8, no. 1, pp. 415–428, 2012.
- [23] D.-G. Sim, O.-K. Kwon, and R.-H. Park, “Object matching algorithms using robust Hausdorff distance measures,” *IEEE Transactions on Image Processing*, vol. 8, no. 3, pp. 425–429, 1999.
- [24] M.-P. Dubuisson and A. K. Jain, “A modified Hausdorff distance for object matching,” in *Proc. of the IEEE International Conference on Pattern Recognition*, vol. 1, pp. 566–568, 1994.
- [25] B. Takacs, “Comparing face images using the modified Hausdorff distance,” *Pattern Recognition*, vol. 31, no. 12, pp. 1873–1881, 1998.
- [26] B. Takacs and H. Wechsler, “Face recognition using binary image metrics,” in *Proc. of the IEEE International Conference on Automatic Face and Gesture Recognition*, pp. 294–299, 1998.
- [27] A. J. Baddeley, “Errors in binary images and an L^p version of the Hausdorff metric,” *Nieuw Archief voor Wiskunde*, vol. 10, pp. 157–183, 1992.
- [28] E. Gilleland, “Spatial forecast verification: Baddeley’s delta metric applied to the icp test cases,” *Weather and forecasting*, vol. 26, no. 3, pp. 409–415, 2011.
- [29] H. Rue and A. R. Syversveen, “Bayesian object recognition with baddeley’s delta loss,” *Advances in Applied Probability*, pp. 64–84, 1998.
- [30] D. L. Wilson, A. J. Baddeley, and R. A. Owens, “A new metric for grey-scale image comparison,” *International Journal of Computer Vision*, vol. 24, pp. 5–17, 1997.
- [31] D. Coquin and P. Bolon, “Application of Baddeley’s distance to dissimilarity measurement between gray scale images,” *Pattern Recognition Letters*, vol. 22, no. 14, pp. 1483–1502, 2001.
- [32] C. Lopez-Molina, D. Ayala-Martini, A. Lopez-Maestresalas, and H. Bustince, “Baddeley’s delta metric for local contrast computation in hyperspectral imagery,” *Progress in Artificial Intelligence*, vol. 6, no. 2, pp. 121–132, 2017.
- [33] J. Fernández Salido and S. Murakami, “Extending yager’s orness concept for the owa aggregators to other mean operators,” *Fuzzy Sets and Systems*, vol. 139, no. 3, pp. 515–542, 2003.
- [34] J. Dujmović, “Andness and orness as a mean of overall importance,” in *IEEE International Conf. on Fuzzy Systems*, pp. 1–6, 2012.
- [35] J. Öfverstedt, J. Lindblad, and N. Sladoje, “Stochastic distance transform: Theory, algorithms and applications,” *Journal of Mathematical Imaging and Vision*, vol. 62, no. 5, pp. 751–769, 2020.
- [36] A. Yilmaz, X. Li, and M. Shah, “Contour-based object tracking with occlusion handling in video acquired using mobile cameras,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 26, no. 11, pp. 1531–1536, 2004.
- [37] F. Mokhtarian and A. Mackworth, “Scale-based description and recognition of planar curves and two-dimensional shapes,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, no. 1, pp. 34–43, 1986.
- [38] R. M. Haralick and J. S. J. Lee, “Context dependent edge detection and evaluation,” *Pattern Recognition*, vol. 23, no. 1-2, pp. 1–19, 1990.
- [39] J. Pont-Tuset and F. Marques, “Measures and meta-measures for the supervised evaluation of image segmentation,” in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 2131–2138, 2013.