



HAL
open science

ccDice: A topology-aware Dice score based on connected components

Pierre Rougé, Odyssée Merveille, Nicolas Passat

► To cite this version:

Pierre Rougé, Odyssée Merveille, Nicolas Passat. ccDice: A topology-aware Dice score based on connected components. Workshop on Topology- and Graph-Informed Imaging Informatics (TGI3@MICCAI), 2024, Marrakesh, Morocco. pp.11-21, 10.1007/978-3-031-73967-5_2. hal-04653406v3

HAL Id: hal-04653406

<https://hal.science/hal-04653406v3>

Submitted on 30 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ccDice: A Topology-Aware Dice Score Based on Connected Components

Pierre Rougé^{1,2}[0000-0003-2448-3935], Odysée Merveille²[0000-0002-9918-3761], and Nicolas Passat¹[0000-0002-0320-4581]

¹ Université de Reims Champagne Ardenne, CRESTIC, Reims, France

² Univ Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220, U1294, F-69100, Lyon, France
pierre.rouge@creatis.insa-lyon.fr

Abstract. Image segmentation is a complex task that aims to simultaneously satisfy various quality criteria. In this context, topology is being increasingly considered. Guaranteeing correct topological properties is indeed crucial for objects presenting challenging (e.g. small, elongated, numerous) shapes. This is especially true in medical imaging. Designing topology-aware metrics is then relevant, both for assessing the quality of segmentation results and for designing losses involved in learning procedures. In this article, we introduce ccDice (connected component Dice), a topological metric that generalises the popular Dice score. By contrast to Dice, that acts at the scale of pixels, ccDice acts at the scale of connected components of the compared objects, leading to a topological assessment of their relative structure and embedding. ccDice is a simple, explainable, normalized and low-computational topological metric. We provide a formal definition of ccDice, an algorithmic scheme for computing it, and we assess its behaviour by comparison to other usual topological metrics, thus emphasizing its relevance. Code is available on GitHub: <https://github.com/PierreRouge/ccDice>.

Keywords: Topology · Metric · Binary objects · Connectedness · Segmentation

1 Introduction

Segmentation is one of the most crucial tasks in medical image analysis. The result of a segmentation process is a set of pixels that represents a structure of interest (e.g. organ, tissue, lesion). This digital object should exhibit correct properties with respect to the real structure it describes, in terms of morphology, geometry and topology.

Early in the raise of computer science, various approaches were proposed to design topological models for numerical imaging [18, 12] with the purpose to be compliant with the underlying continuous topology of Euclidean spaces [13, 14]. This opened the way to the development of a rich panel of tools for digital objects, based on topological invariants (e.g. connectedness [17], homology [7], homotopy [5]). These topological concepts were progressively involved in the design of methods and tools dedicated to medical image analysis and processing [19]. For decades, topological notions were directly embedded in algorithms to guide them by modeling topological priors or providing topological regularization schemes. The rapid development of deep learning in

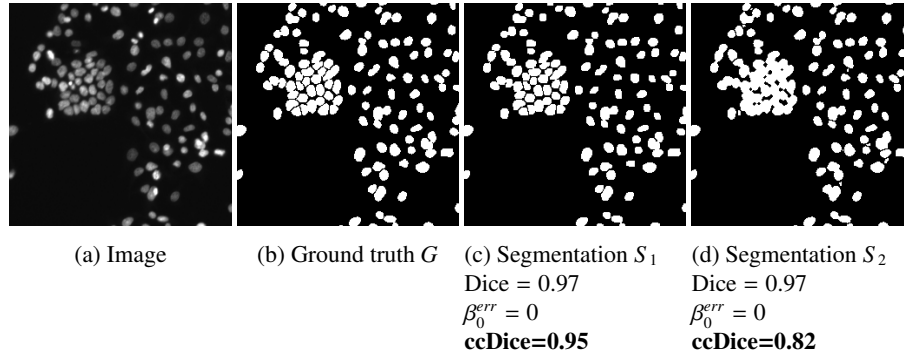


Fig. 1: An image of cells (a), its ground truth G (b) and two segmentation results S_1 (c) and S_2 (d). The cluster of cells is incorrectly connected in S_2 , which also contains false positive artifacts, resulting in an identical number of CCs as in (b–c), and thus a null β_0^{err} metric. Although the false positives and false negatives are not at the same locations, their total number is equal in both segmentations, resulting in equal Dice scores (0.97). Based solely on Dice and β_0^{err} , it is impossible to differentiate these segmentations, even though S_1 is better than S_2 . By contrast, ccDice rates S_1 higher than S_2 (0.95 vs. 0.82).

medical image segmentation, together with the growing trend to consider topology as a quality feature, led to a new paradigm. Topological concepts now tend to be embedded in metrics, for designing both quality scores dedicated to the evaluation of segmentation results, and losses dedicated to train topology-aware machine/deep learning models.

Therefore, recent works have introduced novel metrics specifically designed to take into account topological information. In the context of tubular structure segmentation, Shit et al. [21] proposed the centerline Dice (clDice), which considers the medial axes of both the predicted and ground truth segmentation to avoid bias effects induced by larger vs. smaller structures. Some important works leverage discrete Morse theory (DMT) to identify important topological structures from the likelihood map outputted by a neural network. For example, Hu et al. [11] designed a loss focusing on the correct segmentation of these critical structures. Also, Gupta et al. [9] introduced a probabilistic DMT framework which allows computing a structure-wise uncertainty estimation which is more interpretable than pixel-wise uncertainty, especially for curvilinear structures. Algebraic topology, including topological descriptors, such as Betti numbers or persistent barcodes [1], were also investigated. For instance, Clough et al. [4] proposed to include topological priors in the segmentation. Their loss function enforces the persistent barcodes of the predicted segmentation to be in line with the theoretical Betti numbers of the target object. Byrne et al. [2] extended this approach to the case of multiclass segmentation. Alternatively, Perret et al. [16] proposed to rely on morphological trees in order to model topological features related to the first and last homology groups, leading to a differentiable, low-cost topological loss likely to constrain the structure of a grey-level image. Hu et al. [10] introduced a loss enforcing the segmentation to have the same Betti numbers as the ground truth. However, these topological descriptors lack spatial awareness. Indeed, two images with identical topological descriptors can

represent significantly different structures. To address this issue, Stucki et al. [22] leveraged persistence barcodes and designed a loss function that spatially matches objects with the same topological features between the predicted and ground truth segmentations. This approach marked a significant progress as it proposed a topological and spatially coherent matching of observed features. A drawback of this approach is its high computational cost, which makes it hardly tractable in many cases, in particular in 3 dimensions.

In this article, we propose a novel metric called ccDice, for *connected component Dice*, to assess both the topological and spatial accuracy of a segmentation. We designed a metric being (1) explainable, (2) fast to compute, (3) related to both the topology and the spatial embedding of the objects. The notion of connected components (CCs), that provides an intermediate level of topological representation of an object, is a good trade-off regarding objectives (1–3). ccDice is a generalization of the usual Dice score [6] and thus shares its relevant properties: an explainable and normalized metric.

The core idea is to design a mapping between the CCs of the prediction and those of the ground-truth. This mapping is driven by the spatial matching between the CCs, thus involving spatial embedding in the topological analysis. The behaviour of ccDice is illustrated in Fig. 1.

We provide a formal definition of ccDice in Sect. 2. In Sect. 3, we describe an algorithm to compute efficiently ccDice. We experimentally evaluate the relevance of ccDice in Sect. 4, by comparison with other classic overlap and topological metrics.

2 From Dice to ccDice

We aim to compare binary images based on their CCs. This requires to embed images in a topological space (e.g. [18, 12]). Connected components will be handled as elements of a partition of an image. Thus ccDice is valid for digital images, but more generally for any set subdivided as a partition. Given a (nonempty) binary image X , the set of its CCs is noted $C[X]$ and is a partition of X .

2.1 Dice from (mis)matching

Let $S, G \subseteq \Omega$ be two subsets of a set Ω (the support of the image), that may represent a segmentation result (S) and ground truth (G). The Dice score [6] of S with respect to G is defined as:

$$Dice(S, G) = \frac{2|S \cap G|}{|S| + |G|}. \quad (1)$$

By noting $tp(S, G) = |S \cap G|$, $fp(S, G) = |S \setminus G|$ and $fn(S, G) = |G \setminus S|$, Eq. (1) can be written as follows:

$$Dice(S, G) = \frac{2 \cdot tp(S, G)}{2 \cdot tp(S, G) + fp(S, G) + fn(S, G)}. \quad (2)$$

This formulation of Eq. (2) is generally the one used in the context of segmentation evaluation.

Alternatively to this definition in terms of true positives (tp), false positives (fp) and false negatives (fn), we can view the Dice score as a combination of the number of matching pixels (m) and mismatching pixels (\bar{m}) between S and G :

$$Dice(S, G) = \frac{m(S, G) + m(G, S)}{m(S, G) + \bar{m}(S, G) + m(G, S) + \bar{m}(G, S)} = \frac{m(S, G) + m(G, S)}{|S| + |G|}, \quad (3)$$

where:

$$m(S, G) = m(G, S) = |S \cap G| = tp(S, G) = tp(G, S) \quad (4)$$

$$\bar{m}(S, G) = |S \setminus G| = fp(S, G) = fn(G, S) = |S| - m(S, G). \quad (5)$$

2.2 Matching connected components

From Eq. (3), the Dice score can be built based on the notions of (mis)matching. In Eqs. (4–5), the notions of (mis)matching are defined in a point-wise paradigm. Our purpose is to extend the notion of Dice from point comparison to CC comparison. Following Eq. (3), this requires to generalize the notions of (mis)matching (Eqs. (4–5)) from sets of points to sets of CCs.

Intuitively, a CC X of the segmentation S matches a CC Y of the ground truth G if their intersection is significant with respect to the size of X .

Let $\varphi_{S,G}^\lambda : C[S] \rightarrow C[G]$ be the matching function of CCs from S to a set G , where $\lambda \in (0, 1]$ is a parameter controlling the required degree of the overlap.

We define the *embedding score*, $\mathcal{E}(X, Y)$, as a function quantifying the degree of overlap of X with respect to Y :

$$\mathcal{E}(X, Y) = \frac{|X \cap Y|}{|X|} \in [0, 1]. \quad (6)$$

Then, the matching $\varphi_{S,G}^\lambda$ is defined such that for any $X \in C[S]$, we have:

$$\mathcal{E}(X, \varphi_{S,G}^\lambda(X)) \geq \lambda. \quad (7)$$

Remark 1 A CC $X \in C[S]$ has an image by $\varphi_{S,G}^\lambda$ in $C[G]$ iff there exists a CC $Y \in C[G]$ such that $\mathcal{E}(X, Y) \geq \lambda$.

Remark 2 If $\lambda > 0.5$ then $\varphi_{S,G}^\lambda$ is unique.

With this definition, each CC $X \in C[S]$ is associated to *at most one* CC $Y \in C[G]$. However, each CC $Y \in C[G]$ could be associated to *many* CCs $X \in C[S]$. This property is undesirable as we need each CC $Y \in C[G]$ to be associated with *at most one* CC $X \in C[S]$. To enforce this behavior, we require $\varphi_{S,G}^\lambda$ to be injective (see Sect. 3).

Based on these notions and hypotheses, We can now define the number of matching (μ) and of mismatching ($\bar{\mu}$) on CCs as follows:

$$\mu(S, G) = |\{X \in C[S] \mid \exists Y \in C[G], Y = \varphi_{S,G}^\lambda(X)\}| \quad (8)$$

$$\bar{\mu}(S, G) = |\{X \in C[S] \mid \forall Y \in C[G], Y \neq \varphi_{S,G}^\lambda(X)\}| = |S| - \mu(S, G). \quad (9)$$

Algorithm 1: Compute ccDice

Input: $S, G \subseteq \Omega$
Input: $\lambda \in (0, 1]$
Output: ccDice $\in [0, 1]$

- 1 **Build** $C[S]$
- 2 **Build** $C[G]$
- 3 $\mu(S, G) := \text{Compute Matching}(C[S], C[G], \lambda)$
- 4 $\mu(G, S) := \text{Compute Matching}(C[G], C[S], \lambda)$
- 5 ccDice $:= (\mu(S, G) + \mu(G, S)) / (|C[S]| + |C[G]|)$

2.3 ccDice from matching

We are now ready to define the notion of ccDice. This definition consists of substituting in Eq. (3) the number of (mis)matchings m and \bar{m} on pixels (Eqs. (4–5)) by the number of (mis)matchings μ and $\bar{\mu}$ on CCs (Eqs. (8–9)). We set:

$$ccDice(S, G) = \frac{\mu(S, G) + \mu(G, S)}{\mu(S, G) + \bar{\mu}(S, G) + \mu(G, S) + \bar{\mu}(G, S)} = \frac{\mu(S, G) + \mu(G, S)}{|C[S]| + |C[G]|}. \quad (10)$$

Remark 3 We have $ccDice(S, G) \in [0, 1]$.

Remark 4 We have $ccDice(S, G) = 1$ iff $|C[S]| = |C[G]|$ and both $\varphi_{S,G}^\lambda$ and $\varphi_{G,S}^\lambda$ are bijective.

The notion of ccDice generalizes the notion of Dice. This is established by the following proposition.

Proposition 5 *If we endow Ω with a totally disconnected topological space (i.e. a space that has only singletons as connected subsets), then for two nonempty sets $S, G \subseteq \Omega$, we have $Dice(S, G) = ccDice(S, G)$.*

3 Computing ccDice

Let $S, G \subseteq \Omega$ be two nonempty sets. We note $C[S] = \{X_i\}_{i=1}^t$ and $C[G] = \{Y_j\}_{j=1}^u$ with $t, u \geq 1$, the sets of CCs of S and G , respectively.

The computation of ccDice (Alg. 1) with respect to S and G mainly consists of building the two injective functions $\varphi_{S,G}^\lambda : C[S] \rightarrow C[G]$ and $\varphi_{G,S}^\lambda : C[G] \rightarrow C[S]$ in order to set the values $\mu(S, G)$ and $\mu(G, S)$. These two functions depend on the parameter $\lambda \in (0, 1]$ that determines the tolerance of the embedding between the compared CCs (Eqs. (6–7)).

The computation of $\varphi_{S,G}^\lambda$ and $\varphi_{G,S}^\lambda$ is described in Alg. 2. In practice, we process each candidate pair $(X_i, Y_j) \in C[S] \times C[G]$ for which we have precomputed the embedding score $\varepsilon_{i,j} = \mathcal{E}(X_i, Y_j)$.

We add a pair (X_i, Y_j) to $\varphi_{S,G}^\lambda$, i.e. we set $\varphi_{S,G}^\lambda(X_i) = Y_j$ and increment $\mu(S, G)$ if $\mathcal{E}(X_i, Y_j) \geq \lambda$. When doing so, the other pairs $\mathcal{E}(X_i, \star)$ can no longer be considered

Algorithm 2: Compute matching

Input: $C[S] = \{X_i\}_{i=1}^t$
Input: $C[G] = \{Y_j\}_{j=1}^u$
Input: $\lambda \in (0, 1]$
Output: $\mu(S, G)$

- 1 **Build** $\{\varepsilon_{i,j}\}_{(i,j) \in \llbracket 1,t \rrbracket \times \llbracket 1,u \rrbracket}$
- 2 **Sort** $E = \{(i, j) \mid \varepsilon_{i,j} \geq \lambda\}$ **by decreasing values of** $\varepsilon_{i,j}$
- 3 $\mu(S, G) := 0$
- 4 **foreach** $(i, j) \in E$ **(sorted)** **do**
- 5 **if** (i, \star) **and** (\star, j) **are not discarded** **then**
- 6 $\mu(S, G) := \mu(S, G) + 1$
- 7 **Discard** (i, \star)
- 8 **Discard** (\star, j)

(since $\varphi_{S,G}^\lambda$ is a function) and the same holds for the pairs $\mathcal{E}(\star, Y_j)$ (since $\varphi_{S,G}^\lambda$ is injective).

As stated in Rem. 2, when $\lambda > 0.5$, the definition of $\varphi_{S,G}^\lambda$ and $\varphi_{G,S}^\lambda$ is unique. By contrast, when $\lambda \leq 0.5$, many valid functions may be defined. In order to deal with this indeterminism, when building $\varphi_{S,G}^\lambda$ we process the candidate pairs $(X_i, Y_j) \in C[S] \times C[G]$ by decreasing value of $\mathcal{E}(X_i, Y_j)$ (the same holds for $\varphi_{G,S}^\lambda$). A complexity analysis (presented in supplementary materials) shown that the overall space and time costs of Alg. 1 are $\mathcal{O}(n)$ and $\mathcal{O}(n \log n)$, respectively.

Time and space cost We set $n = |\mathcal{Q}|$, that corresponds to the number of points of the processed images. Alg. 1 relies on a pre-processing step (Lines 1–2) that consists of building the CCs of S and G . Standard binary images are composed of points (pixels/voxels) with a constant number of adjacent neighbours. In this context, building the CCs of S (resp. G) has time and space costs $\mathcal{O}(n)$ (in particular, we store the CC labels of S and G). The time and space costs of Lines 3–4 are those of Alg. 2 (see below). The time and space costs of Line 5 are $\mathcal{O}(1)$.

Alg. 2 first builds the embedding costs $\varepsilon_{i,j} = \mathcal{E}(X_i, Y_j)$ for the pairs of CCs of $C[S]$ and $C[G]$ (Line 1). This can be done by building a $t \times u$ (sparse) matrix to store these values and scanning the points of the CCs of $C[S]$ and $C[G]$. The time cost of Line 1 is then $\mathcal{O}(n)$. Since we are only interested by pairs (X_i, Y_j) such that $\varepsilon_{i,j} \neq 0$, by assuming that for each X_i , the number of Y_j such that $\varepsilon_{i,j} > 0$ is $\mathcal{O}(1)$ (which generally holds in imaging applications), the space cost of Line 1 for storing $\{\varepsilon_{i,j}\}_{(i,j)}$ is $\mathcal{O}(n)$. The time cost for sorting the useful subset of $\llbracket 1, t \rrbracket \times \llbracket 1, u \rrbracket$ with respect to $\varepsilon_{i,j}$ (i.e. the subset that correspond to non-null $\varepsilon_{i,j}$ values) is $\mathcal{O}(n \log n)$. The space cost for storing this subset is $\mathcal{O}(n)$. The space and time costs of Line 3 are $\mathcal{O}(1)$. Line 4 iterates $\mathcal{O}(n)$ times. At each iteration, checking whether (i, \star) and (\star, j) are not already discarded has a time cost $\mathcal{O}(1)$ by relying on the matrix built at Line 1. The space and time costs of Line 6 are $\mathcal{O}(1)$. Still by using the matrix built at Line 1, the time costs of Lines 7–8 are $\mathcal{O}(1)$. The overall space and time costs of Alg. 2 are then $\mathcal{O}(n)$ and $\mathcal{O}(n \log n)$, respectively.

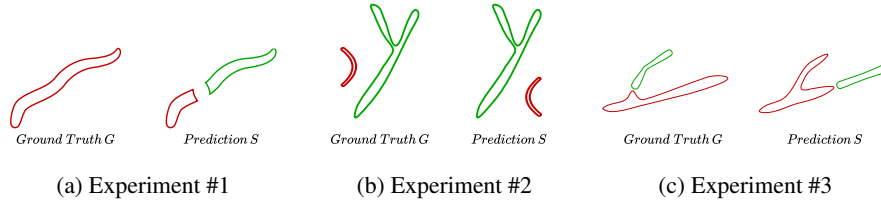


Fig. 2: Illustrations of the three proposed experiments. Green (resp. red) depicts a matched (resp. unmatched) CC. More formally, a green CC of S means that we have $\mathcal{E}(X, \varphi_{S,G}^\lambda(X)) \geq \lambda$ and a green CC of G means that we have $\mathcal{E}(X, \varphi_{G,S}^\lambda(X)) \geq \lambda$. Here λ is set to 0.5.

It follows that the overall space and time costs of Alg. 1 are also $O(n)$ and $O(n \log n)$, respectively. Note that when $\lambda > 0.5$, the construction of the two mapping functions is deterministic and does not require the sorting of the embedding values (i.e. we can skip Line 2 of Alg. 2). In that case, the overall time cost of Alg. 1 becomes $O(n)$.

4 Evaluation

4.1 Experiments

We compare ccDice with standard metrics dedicated to segmentation assessment: (1) the Dice score [6] which is the gold standard for pixel-wise segmentation; (2) β_0^{err} , the absolute difference between the number of CCs of the segmentation and the ground-truth; (3) cIDice [21] which is dedicated to topological assessment for curvilinear structures, and (4) β_{match}^{err} [22] which leverages persistence barcodes to obtain a spatial matching between topological objects.

For these comparisons, we use the CHASE dataset [8] that provides 28 retinal images with vascular annotations. To illustrate the behaviour of ccDice, we propose three experiments where we modified artificially the annotations to generate several pairs (G, S) . These experiments represent a type of error that could occur in segmentation, with a gradual escalation in the number of errors. These experiments are illustrated in Fig. 2.

Experiment #1: The goal is to study the behavior of metrics with an increasing number of disconnections in the segmentation. For each annotation, we progressively add random disconnections and compute the metrics accordingly.

Experiment #2: The goal is to study the behavior of the metrics when the number of CCs remains constant while the pixel-wise overlap of G and S progressively decreases. We start from annotations that we randomly disconnect to create G . S is then created from G by progressively removing true CCs and adding an equivalent number of false ones.

Experiment #3: The goal is to study the behavior of the metrics when the number of CCs and pixelwise overlap remain constant between S and G , but with distinct CCs. We start from an annotation, and we create S and G by randomly introducing the same number of disconnections in each, but at different locations.

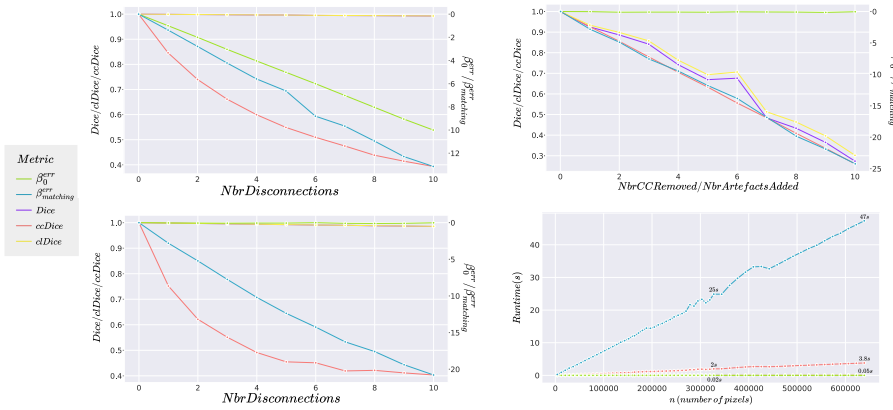


Fig. 3: Experiments #1 (upper-left), #2 (upper-right), #3 (bottom-left). Average runtimes per image with respect to image size (bottom-right). The sizes of images vary from 64×64 to 850×850 . Dice (purple), cIDice (yellow), ccDice (red), β_0^{err} (green) and β_{match}^{err} (blue)

4.2 Results

Fig. 3 depicts the evolution of each metric in the three experiments, where the addition of errors should be reflected by a decrease of their values.

In Exp. #2, the Dice and cIDice decrease, since the modifications to the CCs are carried out spatially. However, they remain close to 1 (best score) in Exps. #1 and #3 as the disconnections in S do not significantly alter the pixel-wise / medial axis coherence between S and G .

In Exp. #1, β_0^{err} increases, since the number of CCs between S and G diverges. However in Exps. #2 and #3, it remains close to 0 (best score) as the number of CCs remains unchanged.

This shows that neither the Dice/cIDice nor β_0^{err} are able to correctly assess the topological relevance of a segmentation across all scenarios. The combination of both metrics also fails in some cases, as evidenced by Exp. #3.

Both ccDice and β_{match}^{err} exhibit correct behaviors in the three experiments. Their respective values decrease when the number of CCs increases (Exp. #1) or when the spatial matching between the CCs of S and G diverges (Exps. #2 and #3).

However, ccDice presents various advantages. First it is a normalized metric, with values in $[0, 1]$, similarly to Dice. By contrast, β_{match}^{err} can be arbitrarily high which makes comparisons between images difficult. Second, ccDice can be computed with a low time complexity, which is never greater than quasi-linear (or even linear when $\lambda > 0.5$). Fig. 3 (bottom-right) presents the average running times of the metrics with respect to the image size on the CHASE dataset. The running time of ccDice is very close to that of Dice, cIDice and β_0^{err} , whereas the running time of β_{match}^{err} is significantly

greater. Third, we propose a 3D implementation of ccDice³, whereas the β_{match}^{err} is only in 2D.

5 Discussion

Analysing holes with ccDice ccDice allows one to compare two binary images $S, G \subseteq \Omega$ with respect to their respective “objects”, i.e. the CCs $X \in C[S]$ and $Y \in C[G]$, which correspond to the first homology groups of S and G . It is possible to use ccDice to investigate the “holes” of S, G , which correspond to their last homology group. We can apply ccDice on the complements $\bar{S} = \Omega \setminus S$ and $\bar{G} = \Omega \setminus G$, i.e. compute $\text{ccDice}(\bar{S}, \bar{G})$. It is possible to combine $\text{ccDice}(S, G)$ and $\text{ccDice}(\bar{S}, \bar{G})$ to get access to a normalized metric that evaluates a segmentation with respect to these two homology groups.

Towards ccDice loss One significant limitation of our method is its lack of differentiability, preventing its use as a loss function. In the following discussion, we introduce two extensions of the ccDice for fuzzy images and grey-level images, respectively. This sets the stage for future research aimed at developing a differentiable version of the ccDice.

Soft ccDice It is possible to extend ccDice to a soft version that evaluates fuzzy images $\mathcal{S}, \mathcal{G} : \Omega \rightarrow [0, 1]$, instead of binary images $S, G \subseteq \Omega$. A continuous soft version of ccDice can be defined as follows:

$$\text{ccDice}_{\text{soft}}(\mathcal{S}, \mathcal{G}) = \int_0^1 \text{ccDice}(\tau_v(\mathcal{S}), \tau_v(\mathcal{G})) dv \quad (11)$$

where τ_v is the upper thresholding at value v function defined by $\tau_v(\mathcal{S}) = \{x \in \Omega \mid \mathcal{S}(x) \geq v\}$. This soft version is compliant with ccDice. Indeed, if $\mathcal{S}, \mathcal{G} : \Omega \rightarrow [0, 1]$ are binary functions, i.e. if $\mathcal{S}(x)$ and $\mathcal{G}(x) = 0$ or 1 for any $x \in \Omega$, we have $\text{ccDice}_{\text{soft}}(\mathcal{S}, \mathcal{G}) = \text{ccDice}(\tau_1(\mathcal{S}), \tau_1(\mathcal{G}))$. This allows in particular to compare e.g. a fuzzy segmentation map $\mathcal{S} = \Omega \rightarrow [0, 1]$ with a binary ground truth $G \subseteq \Omega$.

In practice, Eq. (11) is computed in a discrete way, by sampling a finite number of values in $[0, 1]$. It is also possible to weight the ccDice components of this sum, leading to a computationally tractable formulation:

$$\text{ccDice}_{\text{soft}}(\mathcal{S}, \mathcal{G}) = \sum_{v \in V} \alpha_v \cdot \text{ccDice}(\tau_v(\mathcal{S}), \tau_v(\mathcal{G})) \quad (12)$$

where $V \subset [0, 1]$ is a finite set, $\alpha_v > 0$ and $\sum_{v \in V} \alpha_v = 1$. In particular, $\text{ccDice}_{\text{soft}}$ can be computed with a time cost $\mathcal{O}(|V| \cdot n \log n)$.

Grey-level ccDice More generally, one may want to define a generalization of ccDice to grey-level images without requiring to explicitly build upon persistent homology paradigm [22]. A convenient way may consist of relying on the notions of component-trees (min- and max-tree) [20] or the unifying complete tree of shapes [15], that allow to model the hierarchical structure of all the CCs (objects and holes) of the threshold sets of grey-level functions $\mathcal{S}, \mathcal{G} : \Omega \rightarrow \mathbb{R}$. In particular, these trees can be

³ <https://github.com/PierreRouge/ccDice>

built in quasi-linear time [3] and can be involved in the design of differentiable losses [16]. These properties open the way to future developments of a grey-level version of ccDice.

6 Conclusion

In this paper, we introduced ccDice, which builds upon the notions of CCs and spatial matching to provide a metric evaluating both the topology and spatial matching of segmentations. Similarly to the Dice metric, it is a normalized, interpretable metric. It also has the notable property to generalize the Dice metric from pixels to CCs (Prop. 5), which emphasizes its theoretical soundness. Similarly to β_{match}^{err} , it takes into account topological and spatial embedding of the images in order to evaluate the relevance of a segmentation. In 2D, ccDice is as discriminant as its concurrent the β_{match}^{err} as it can assess the two first homology groups, but its computation time is largely lower. Finally, ccDice appears as a relevant tool for topological analysis of segmentation, which opens the way to further fuzzy / grey-level extensions, a step towards developing a differentiable version of the ccDice.

Acknowledgments. This work was supported by the *Agence Nationale de la Recherche* (Grants ANR-20-CE45-0011, ANR-22-CE45-0034, ANR-22-CE45-0018 and ANR-23-CE45-0015).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bauer, U., Lesnick, M.: Induced matchings and the algebraic stability of persistence barcodes. *J Comput Geom* **6**, 162–191 (2015)
2. Byrne, N., et al.: A persistent homology-based topological loss for CNN-based multiclass segmentation of CMR. *IEEE T Med Imaging* **42**, 3–14 (2022)
3. Carlinet, E., Géraud, T.: A comparative review of component tree computation algorithms. *IEEE T Image Process* **23**, 3885–3895 (2014)
4. Clough, J.R., et al.: A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE T Pattern Anal* **44**, 8766–8778 (2020)
5. Couprie, M., Bertrand, G.: New characterizations of simple points in 2D, 3D, and 4D discrete spaces. *IEEE T Pattern Anal* **31**, 637–648 (2009)
6. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**, 297–302 (1945)
7. Edelsbrunner, H., Harer, J.: Persistent homology – A survey. *Contemp Math* **453**, 257–282 (2008)
8. Fraz, M.M., et al.: An ensemble classification-based approach applied to retinal blood vessel segmentation. *IEEE T Bio-Med Eng* **59**, 2538–2548 (2012)
9. Gupta, S., et al.: Topology-aware uncertainty for image segmentation. In: *NeurIPS, Procs.* (2024)
10. Hu, X., et al.: Topology-preserving deep image segmentation. In: *NeurIPS, Procs.* (2019)
11. Hu, X., et al.: Topology-aware segmentation using discrete Morse theory. In: *ICLR, Procs.* (2021)

12. Kovalevsky, V.A.: Finite topology as applied to image analysis. *Comput Vision Graph* **46**, 141–161 (1989)
13. Mazo, L., et al.: Paths, homotopy and reduction in digital images. *Acta Appl Math* **113**, 167–193 (2011)
14. Mazo, L., et al.: Digital imaging: A unified topological framework. *J Math Imaging Vis* **44**, 19–37 (2012)
15. Passat, N., Mendes Forte, J., Kenmochi, Y.: Morphological hierarchies: A unifying framework with new trees. *J Math Imag Vis* **65**, 718–753 (2023)
16. Perret, B., Cousty, J.: Component tree loss function: Definition and optimization. In: *DGMM, Procs.* pp. 248–260 (2022)
17. Rosenfeld, A.: Adjacency in digital pictures. *Inform Control* **26**, 24–33 (1974)
18. Rosenfeld, A.: Digital topology. *Am Math Mon* **86**, 621–630 (1979)
19. Saha, P.K., Strand, R., Borgefors, G.: Digital topology and geometry in medical imaging: A survey. *IEEE T Med Imaging* **34**, 1940–1964 (2015)
20. Salembier, P., Oliveras, A., Garrido, L.: Anti-extensive connected operators for image and sequence processing. *IEEE T Image Process* **7**, 555–570 (1998)
21. Shit, S., et al.: clDice—A novel topology-preserving loss function for tubular structure segmentation. In: *CVPR, Procs.* pp. 16560–16569 (2021)
22. Stucki, N., et al.: Topologically faithful image segmentation via induced matching of persistence barcodes. In: *ICML, Procs.* pp. 32698–32727 (2023)