Journal of Mathematical Imaging and Vision (2023) 65:861–873
https://doi.org/10.1007/s10851-023-01146-x

Morphology on Categorical Distributions

Silas Nyboe Ørting[1,2] · **Hans Jacob Teglbjærg Stephensen**[1,2] · **Jon Sporring**[1,2]

Received: 9 January 2022 / Accepted: 10 March 2023 / Published online: 8 May 2023

Abstract
Mathematical morphology (MM) is an indispensable tool for post-processing. Several extensions of MM to categorical images, such as multi-class segmentations, have been proposed. However, none provide satisfactory definitions for morphology on probabilistic representations of categorical images. The categorical distribution is a natural choice for representing uncertainty about categorical images. Extending MM to categorical distributions is problematic because categories are inherently unordered. Without ranking categories, we cannot use the standard framework based on supremum and infimum. Ranking categories is impractical and problematic. Instead, we consider the probabilistic representation and operations that emphasize a single category. In this work, we review and compare previous approaches. We propose two approaches for morphology on categorical distributions: operating on Dirichlet distributions over the parameters of the distributions and operating directly on the distributions. We propose a "protected" variant of the latter and demonstrate the proposed approaches by fixing misclassifications and modeling annotator bias.

Keywords Mathematical morphology · Categorical distributions · Dirichlet distributions · Multi-class segmentation

1 Introduction

Multi-class segmentation problems are common in analysis of biomedical images. A typical solution is to train a neural network pixel classifier. Commonly, these networks predict a probability distribution over all classes in each pixel, which can be thresholded to obtain a final segmentation. These predictions often contains holes, partial misclassifications, shrinkage of small classes and rough borders between classes, resulting in errors in the final segmentation. To improve the segmentation, post-processing is often used to close holes, reclassify uncertain pixel labels based on proximity, grow objects and smoothen rough boundaries.

Mathematical morphology is a powerful framework for post-processing binary and grayscale images. Binary and grayscale morphology are special cases of morphology on complete lattices [1]. A complete lattice is a partially ordered set (poset), where each non-empty subset has an infimum and a supremum. For complete lattices, the core operators, dilation and erosion, can be defined using supremum and infimum: for binary morphology using set union and intersection; and for grayscale morphology using maximum and minimum under the standard total ordering of the reals; see [1] for an in depth treatment of the theoretical foundations of mathematical morphology.

For general multi-class images, there is no natural ordering of the classes, and hence, they do not form a complete lattice. For example, for a segmentation of microscope images of cells into cell membrane, mitochondria and background, any ordering of the classes is task-dependent and not given by the images themselves. A natural representation of this kind of data is the categorical distribution, which can represent both crisp segmentation masks and uncertainty as encountered in prediction images. In the remainder of this work, we will use the term "categorical" instead of "multi-class."

In this work, we provide a thorough review of previously proposed approaches to morphology on categorical images. We then propose two approaches for morphology on categorical distributions: an indirect approach where we operate on Dirichlet distributions that are then transformed to categorical distributions and a direct approach where we operate on the categorical distributions themselves. We then define protected variants of the direct operations that allow finer control over the processing. Finally, we illustrate the utility

✉ Silas Nyboe Ørting
silas@di.ku.dk

[1] Department of Computer Science, University of Copenhagen, Copenhagen, Denmark

[2] QIM - Center for Quantification of Imaging Data from MAX IV, Copenhagen, Denmark

of the proposed approach on two tasks: fixing misclassified mitochondria and modeling annotator bias.

2 Background and Related Work

In this section, we briefly restate morphology on complete lattices and on binary and grayscale images, before we review the most relevant literature [2–8]. What we refer to as categorical images have various names in the literature: color-coded images, label images and n-ary images. In the sections below, we will use the original names in the section titles, but otherwise we will refer to categorical images and categorical morphology.

In the literature, there are three main approaches for extending morphology to images with values that do not have a natural ordering: impose an order on the values, which is the common approach for color images; operate on all categories simultaneously [2, 5]; and operate on a single category at a time [6, 7]

Morphology on color images has received a lot of attention, with the primary focus on ordering colors by exploiting the relationship between dimensions of color spaces. See, for example, [9] for an overview of approaches for defining an ordering of colors. Our focus is on categorical images, where such approaches are less relevant.

2.1 Morphology on Complete Lattices

Let Γ be a set with the partial order $\leq$. The poset $(\Gamma, \leq)$ is a complete lattice if every subset of Γ has an infimum $\wedge$ and a supremum $\vee$. We define an image as a function f from pixel-coordinates $\mathbb{D} = \mathbb{Z}^d$ to Γ and a structuring element B as a subset of $\mathbb{D}$

$$f \in \mathcal{F} = \{g \mid g : \mathbb{D} \mapsto \Gamma\}, \tag{1}$$

$$B \subseteq \mathbb{D}. \tag{2}$$

The dilation (δ) and erosion (ϵ) of f by B are then defined as the supremum and infimum over the local neighborhoods in f given by B

$$\delta(f; B)(x) = \bigvee_{\{y \mid (y-x) \in B\}} f(y), \tag{3}$$

$$\epsilon(f; B)(x) = \bigwedge_{\{y \mid (y-x) \in B\}} f(y). \tag{4}$$

Opening (γ) and closing (ϕ) are the compositions of dilation and erosion

$$\gamma(f; B)(x) = \delta(\epsilon(f; B); B), \tag{5}$$

$$\phi(f; B)(x) = \epsilon(\delta(f; B); B). \tag{6}$$

2.2 Binary and Grayscale Morphology

We define a grayscale image as in (1) with $\Gamma = [0, 1]$. Let $\leq$ be the usual ordering of the reals, then the poset $([0, 1], \leq)$ is a complete lattice, where the min function gives the infimum and the max function the supremum. Let B be defined as in (2). Dilation and erosion can then be obtained from (3) and (4) as

$$\delta(f; B)(x) = \max_{\{y \mid (y-x) \in B\}} f(y), \tag{7}$$

$$\epsilon(f; B)(x) = \min_{\{y \mid (y-x) \in B\}} f(y). \tag{8}$$

If we restrict Γ to $\{0, 1\}$, we obtain binary morphology.

2.3 Morphology on Color-Coded Images

In [2], the authors propose a framework for categorical morphology where pixels have a set of categories. Let $C = \{c_1, c_2, \ldots, c_n\}$ be a set of n categories. The powerset of C, $\mathbb{P}_C$, is the set of all subsets of C, including the empty set. An image f is then defined as in (1) with $\Gamma = \mathbb{P}_C$. In this framework, the value of a pixel can be any element of $\mathbb{P}_C$, e.g $\{c_1\}$, $\{c_1, c_n\}$ or $\{\}$. Let $\subseteq$ be the usual subset relation, then the poset $(\mathbb{P}_C, \subseteq)$ is a complete lattice where set intersection is the infimum and set union is the supremum. In [2], the authors propose to use structuring elements that are images that is $B \in \mathcal{F}$. For the sake of comparison, we first consider the simpler case where B is defined as in (2). Dilation and erosion can then be obtained from (3) and (4) as

$$\delta(f; B)(x) = \bigcup_{\{y \mid (y-x) \in B\}} f(y), \tag{9}$$

$$\epsilon(f; B)(x) = \bigcap_{\{y \mid (y-x) \in B\}} f(y). \tag{10}$$

An example of these operations is shown Fig. 1a.

Let $B \in \mathcal{F}$. Under this scheme, an operation is only performed when one or more categories in the structuring element match a category in the image, and the result depends on the categories in both image and structuring element. Several variations of dilation and erosion are proposed in [2]; here, we only consider the "transparent" operations. Let $\mathbb{D}_f$ be the domain of f and $\mathbb{D}_B$ the domain of B. A specified reference point, $y_0 \in \mathbb{D}_B$, is used to determine whether B matches f and could, for example, be the center of a ball-shaped $\mathbb{D}_B$. Dilation and erosion are then defined as

$$\delta(f; B)(x) = f(x) \cup \bigcup_{\{y \in \mathbb{D}_B \mid f(x+y) \cap B(y_0) \neq \emptyset\}} B(y) \tag{11}$$

$$\epsilon(f; B)(x)$$

$$= \begin{cases} f(x), & \text{if } f(x) \cap B(y_0) = \emptyset \\ f(x) \setminus B(y_0), & \text{if } [\exists y \in \mathbb{D}_B](f(x+y) \cap B(y_0) = \emptyset), \\ f(x) \cup B(y_0), & \text{otherwise} \end{cases} \tag{12}$$

An example of these operations is shown Fig. 1b using a cross-shaped structuring element with y_0 in the center.

2.4 Morphology on Label Images

In [5], the authors propose a framework for categorical morphology where pixels have no category ($\bot$), a unique category or conflicting categories ($\top$). Let $C = \{c_1, c_2, \ldots, c_n\}$ be a set of n categories and let $C_* = C \cup \{\bot, \top\}$. An image f is then defined as in (1) with $\Gamma = C_*$. The poset $(C_*, \leq)$ where $\leq$ satisfies $[\forall c \in C](\bot \leq c \leq \top)$ is a complete lattice. Let B be defined as in (2) and let $V(x) = \{f(x-y) \mid y \in B\}$. Dilation and erosion are then defined as

$$\delta(f; B)(x) = \begin{cases} \top, & \text{if } \top \in V(x) \\ \top, & \text{if } |V(x) \cap C| > 1 \\ V(x) \cap C, & \text{if } |V(x) \cap C| = 1 \\ \bot, & \text{otherwise} \end{cases} \tag{13}$$

$$\epsilon(f; B)(x) = \begin{cases} \bot, & \text{if } \bot \in V(x) \\ \bot, & \text{if } |V(x) \cap C| > 1 \\ V(x) \cap C, & \text{if } |V(x) \cap C| = 1 \\ \top, & \text{otherwise} \end{cases} \tag{14}$$

An example of these operations is shown Fig. 1c. In the context of categorical distributions, where we have detailed information about label uncertainty, this approach is unsuitable due to the loss of information.

2.5 N-ary Morphology

In [6], the authors propose a framework for categorical morphology where pixels have a unique category. Let $C = \{c_1, c_2, \ldots, c_n\}$ be a set of n categories. An image f is then defined as in (1) with $\Gamma = C$. Instead of operating on all categories simultaneously, the authors propose to operate on a single category at a time. Let B be defined as in (2) and let i be the category we operate on. We use subscripts to distinguish single category operations from standard operations. Dilation and erosion are then defined as

$$\delta_i(f; B)(x) = \begin{cases} f(x), & \text{if } [\forall y \in B](f(x+y) \neq i) \\ i, & \text{otherwise} \end{cases} \tag{15}$$

$$\epsilon_i(f; B)(x) = \begin{cases} f(x), & \text{if } f(x) \neq i \\ i, & \text{if } [\forall y \in B](f(x+y) = i) \\ \theta(x, f), & \text{otherwise} \end{cases} \tag{16}$$

where θ is a function that assigns a value in the case where there are different categories in the neighborhood of x. A natural choice for θ, which is also suggested in [6], is to pick the value of the closest pixels. However, this does not help when the closest pixels have different values, which is a fundamental problem when pixel values cannot represent uncertainty. This is solved by ranking the categories and using the ranking to break ties. In general, there is no obvious way of ranking categories based on the image alone, and as the number of multi-category interfaces increases, it becomes more difficult to understand how one particular ranking influence the outcome.

Without ranking categories a priori, the above definition implies an ordering $\leq_i$, which is not a partial order, and thus, $(C, \leq_i)$ is not a complete lattice. In [7], the authors show that $\leq_i$ is a preorder, and formalize constraints for choosing θ such that dilation and erosion form an adjunction and their compositions are an opening and a closing. However, this does not help decide which category to choose when multiple categories are closest, as the constraints on θ do not yield a unique rule for breaking ties. An example of these operations is shown in Fig. 1d, where the question marks highlight two pixels that cannot be assigned a value without a method for breaking ties.

2.6 Fuzzy n-Ary Morphology

In [6], the authors also propose an extension of n-ary morphology to images of categorical distributions. Let $C = \{c_1, c_2, \ldots, c_{n+1}\}$ be a set of $n+1$ categories. The categorical distribution of $n+1$ categories is completely determined by a point in the n-simplex $\Delta^n = \{\pi \in \mathbb{R}^{n+1} \mid \pi_k \geq 0, \sum \pi_k = 1\}$, where π_k is the probability of c_k. An image f is then defined as in (1) with $\Gamma = \Delta^n$. Operations are again defined on a single category at a time. Let B_r be a closed ball of radius r centered at the origin and let i be the category we operate on. Let $f_k(x) = f(x)_k$ be the probability of observing category c_k in pixel x and let $\omega_k(x) = 1 - f_k(x)$. Dilation is then defined as

$$\delta_i(f; B_r)(x)_k = \begin{cases} \delta(f_k; B_r)(x), & \text{if } k = i, \\ [1 - \delta(f_i; B_r)(x)]\frac{f_k(x)}{\omega_i(x)}, & \text{if } k \neq i, \end{cases} \tag{17}$$

where $\delta(f_i; B_r)(x) = 1 \implies [1 - \delta(f_i; B_r)(x)]\frac{f_k(x)}{\omega_i(x)} = 0.$

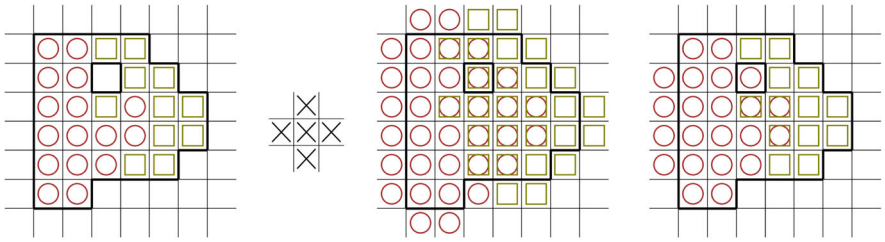

(a) Morphology on color-coded images using (2) as structuring element. The bold boundary highlights the pixels with one or more categories and is only for illustration.

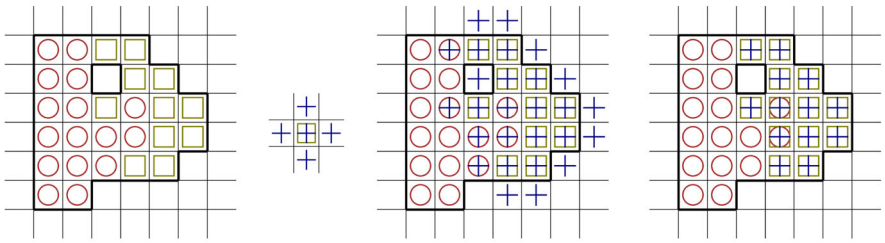

(b) Morphology on color coded images using structuring element from [2] . This figure extends Figure 2. in [2] with the closing operation. Notice that $B(y_0) = \{plus, square\}$, meaning that B will match any pixel with either *plus* or *square*.

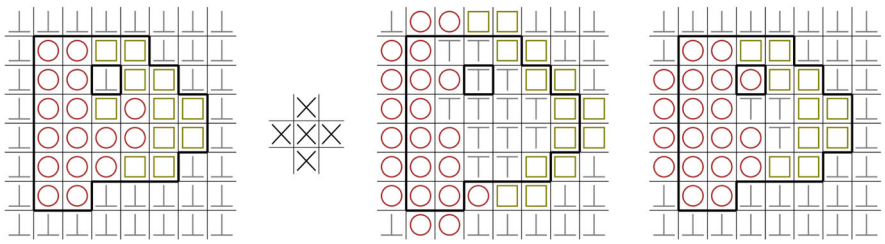

(c) Morphology on label images.

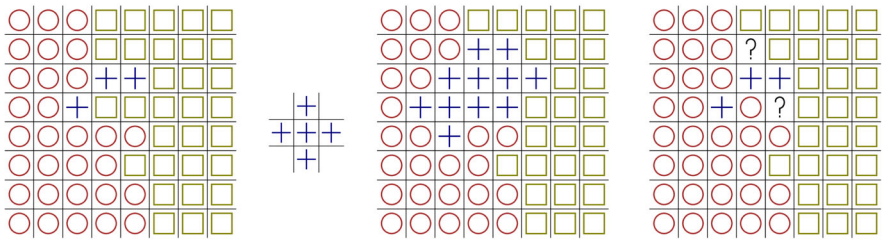

(d) N-ary morphology. The coloring of the structuring element with *plus* indicates that operations are on the *plus* category. The question mark ? indicate pixels that cannot be assigned a valued without additional ordering of the categories. Notice that it is not possible to assign pixels "no-category" as in (a), (b) or (c).

Fig. 1 Comparison of categorical morphologies from the literature. From left to right: image, structuring element, dilation, closing

Two variations on erosion are proposed in [6], neither of which we find satisfactory. The first requires that we pick a ranking of all categories and does not yield idempotent opening and closing

$$\epsilon_i(f; B_r)(x)_k = \begin{cases} \epsilon(f_k; B_r)(x) & \text{if } k = i \\ f_k(x) + f_i(x) - \epsilon(f_i; B)(x) & \text{if } k = \min(\arg\min_{j \neq i}(\delta(f_j; B))) \\ f_k(x) & \text{otherwise} \end{cases} \tag{18}$$

The second assumes that the image is restricted to the edges of the simplex (at most two categories are nonzero in any pixel) and opening and closing are again not idempotent

$$\epsilon_i(f;B_r)(x)_k = \begin{cases} \epsilon(f_k;B_r)(x) & \text{if } k=i \\ \frac{1-\epsilon(f_i;B)(x)}{1-f_i(x)} f_k(x) & \text{if } f_i(x) \le 0.5 \vee \max_{j\neq i}(\delta(f_j;B)(x) < 0.5) \\ 1-\epsilon(f_i;B)(x) & \text{if } k = \min(\arg\max_{j\neq i} \delta(f_j;B)(x)) \\ 0 & \text{otherwise} \end{cases} \tag{19}$$

We refer the reader to [6] for the motivation for these formulations and their properties.

2.7 Fuzzy Pareto Morphology

In [3], the authors propose fuzzy Pareto morphology for color images. An RGB color image can be seen as a three-dimensional fuzzy set, where the membership function for each set corresponds to the value of each color channel. This can equivalently be seen as point in the half-open unit cube. An image f is then defined as in (1) with $\Gamma = (0,1]^d$. For each $a \in \Gamma$, we can associate a hyperrectangle defined by the vector from the origin to a. Fuzzy Pareto morphology is based on the idea of dominance. For $a, b \in \Gamma$, let $a \cap b = \{\min(a_i, b_i)\}_{i=1\ldots d}$ be the intersection of a and b. Let $A(a) = \prod_i a_i$ be the area function, yielding the area of the hyperrectangle of a. The degree to which a dominates b is then

$$\mu_D(a,b) = \frac{A(a\cap b)}{A(b)}, \tag{20}$$

which measures how much of the hyperrectangle of b is contained in the hyperrectangle of a.
Let $B(x) = \{x + y \mid y \in B\}$, dilation and erosion are then defined as

$$\delta(f;B)(x) = f\left(\arg\min_{y\in B(x)} \left\{ \max_{z\in B(x)\wedge z\neq y} \mu_D(f(z), f(y)) \right\}\right), \tag{21}$$

$$\epsilon(f;B)(x) = f\left(\arg\max_{y\in B(x)} \left\{ \min_{z\in B(x)\wedge z\neq y} \mu_D(f(z), f(y)) \right\}\right). \tag{22}$$

Although not directly applicable to categorical distributions, it could easily be extended by either restricting Γ to $\{v \in (0,1]^d \mid \sum_i v_i = 1\}$ or by considering it in the context of the Dirichlet distribution. However, (21) and (22) are not guaranteed to yield a unique solution, requiring us to come up with an arbitration rule.

2.8 Morphology on the Unit Circle

In [10], the authors propose morphology on the unit circle for processing the hue space of color images. The idea is to use structuring elements from the hue space and define an ordering based on the shortest distance along the unit circle between values in the image and values in the structuring element. Although not directly applicable to categorical images, it could be relevant to consider structuring elements that are themselves categorical distributions and base morphology on distance between distributions.

Morphology on the unit circle is also considered in [4] where the authors propose three approaches: using difference operators (e.g., gradient), using grouped data and using "labeled openings." It is the labeled openings that are most relevant in our context. Let f be an image as defined in (1) with $\Gamma = [0, 2\pi]$. In a labeled opening, the unit circle is partitioned into segments $S(\omega) = \{[0,\omega), [\omega, 2\omega), \ldots, [2\pi - \omega, 2\pi)\}$ and each segment $s \in S(\omega)$ gives rise to a binary image $f(x;s) = f(x) \in s$. A labeled opening is then the union of the binary openings of all segments, $\gamma_\omega(f) = \cup_{s\in S(\omega)} f(x;s)$, indicating for each pixel if it is present in at least one of the opened segments. The resulting image highlights areas of uniform direction and the inverse of that image highlights areas with change in direction. Categorical images have a natural partitioning based on the categories. A labeled opening of a categorical image would then be a binary image indicating those pixels where at least one category was preserved after opening each category. The resulting image highlights areas where at least one category is uniformly present, and the inverse of that image highlights areas without a category after the opening, similarly to $\perp$ in Sect. 2.4.

2.9 Morphology on Component Graphs

In [8], the authors propose a framework for morphology on multi-valued images based on component graphs. Let an image be defined as in (1) with $\Gamma = \mathbb{R}^d$. The component graph is constructed from the connected components of the threshold sets of an image. For example, for $d = 2$ and $f(x) \in \{0,1\}^2$ the levels of f are $\{(0,0), (0,1), (1,0), (1,1)\}$ and the level set of $(0,1)$ is $\{x \mid f(x) = (0,1)\}$. Let $\le$ be a partial order on $\{0,1\}^2$. The threshold set of $(0,1)$ is then $T_{(0,1)} = \{x \mid (0,1) \le f(x)\}$. We can represent the threshold set as a binary image and find the connected components in this image. For two levels l_i and l_j, we have $l_i \le l_j \implies T_i \supseteq T_j$, so any connected component in T_j must be contained in a connected component in T_i. The component graph is then constructed by adding a node for each connected component and an edge from node u to node v if the connected component of v is contained in the connected component of u. In order to construct the com-

ponent graph, it is required that Γ allows a minimum, e.g., $\{0\}^d$, such that the graph will be connected. For categorical images, this would require that we have a special background category as in Sect. 2.3 and Sect. 2.4. Further, it requires that each pixel can have multiple categories; otherwise, no component will be nested inside another and the graph will be the root with all connected components as children.

Because the component graph directly exposes the spatial relationship between differently valued regions, it is possible to apply morphological filters, e.g., noise reduction, by pruning some nodes and reconstructing the image from the pruned component graph. Directly pruning the component graph can lead to ambiguity in the reconstruction when a node with two non-comparable parents is removed. The authors propose to solve this by building a component tree of the component graph, prune the tree and then reconstruct the graph from the tree and the image from the graph. In order to construct the component tree, it is necessary to impose a total order on the nodes of the component graph, for example, by using a shape measure on the connected components in the component graph.

Because the component graph only captures spatial relationships when connected components overlap for different level sets, some common post-processing operations, such as closing holes in segmentations, are challenging to perform.

3 Morphology on Categorical Distributions

In this section, we propose two approaches for morphology on categorical distributions. In Sect. 3.1, we show how to operate on all categories simultaneously by operating on Dirichlet distributions. The limitations of this approach will then motivate single category operations that work directly on categorical distributions, which we will define in Sect. 3.2.

3.1 Morphology on Dirichlet Distributions

Let $\mathbb{R}_+$ be the positive real line. We consider the Dirichlet distribution of order $n \geq 2$ with parameters $\alpha \in \mathbb{R}_+^n$, written as $\mathrm{Dir}(\alpha)$, as a distribution over the $(n-1)$-simplex $\Delta^{n-1} = \{\pi \in \mathbb{R}^n \mid \pi_k \geq 0, \sum \pi_k = 1\}$ with density function

$$\mathrm{pdf}(\pi) = \frac{1}{\mathrm{Beta}(\alpha)} \prod_{k=1}^{n} \pi_k^{\alpha_k - 1} \tag{23}$$

where $\mathrm{Beta}(\cdot)$ is the multivariate Beta function defined with the Gamma function as

$$\mathrm{Beta}(\alpha) = \frac{\prod_{k=1}^{n} \mathrm{Gamma}(\alpha_k)}{\mathrm{Gamma}(\sum_{k=1}^{n} \alpha_k)}. \tag{24}$$

Let $X^\alpha \sim \mathrm{Dir}(\alpha)$, with $\alpha \in \mathbb{R}_+^n$. A realization of X^α is a point in the $(n-1)$-simplex, which can be taken as parameters of the categorical distribution with n categories. The expectation of X^α is

$$\mathbb{E}[X_k^\alpha] = \frac{\alpha_k}{\sum \alpha}, \tag{25}$$

which maps each Dirichlet distribution to a specific categorical distribution. Note that $0 < \alpha_k < \infty$ implies that we can only represent categorical distributions in the open simplex. In practice, this is not a problem as we can get arbitrarily close to the boundary of the simplex.

Let f_k be the kth category in f. An image f is defined as in (1) with $\Gamma = \mathbb{R}_+^n$. If we equip f with the ordering $f \leq g \iff [\forall k](f_k(x) \leq g_k(x))$, we obtain a complete lattice. Dilation and erosion are then defined as their grayscale counterparts applied to each category independently

$$\delta(f; B)(x)_k = \delta(f_k; B), \tag{26}$$

$$\epsilon(f; B)(x)_k = \epsilon(f_k; B). \tag{27}$$

An example of these operations is provided in Fig. 2. It is interesting to compare the images of entropy (uncertainty) and α parameter magnitude for dilation and erosion. We can see that entropy and magnitude are positively correlated for dilation, and negatively correlated for erosion. Try to think of dilation as "increasing the probability of everything" and erosion as "decreasing the probability of everything." It is, of course, impossible to change the "probability of everything," all we can do is shuffle probability around between categories. Nevertheless, the idea captures our intent and the magnitude image reflects this. High entropy and low magnitude can be interpreted as uncertainty due to a lack of confidence, whereas high entropy and high magnitude can be interpreted as uncertainty due to over confidence. Opening and closing appear to be more straightforward as they, respectively, decrease and increase uncertainty at the boundaries between overlapping categories.

We can easily extend these operators to operate on a subset of categories S by simply only updating those categories

$$\delta(f; B|S)(x)_k = \begin{cases} \delta(f_k; B), & \text{if } k \in S \\ f_k, & \text{otherwise} \end{cases} \tag{28}$$

$$\epsilon(f; B|S)(x)_k = \begin{cases} \epsilon(f_k; B), & \text{if } k \in S \\ f_k, & \text{otherwise} \end{cases} \tag{29}$$

An example of these operations is provided in Fig. 3 where we operate on the green category. Consider the gray/blue region surrounded by green that is indicated with a white ellipse in the left image of the second row. When we dilate the green

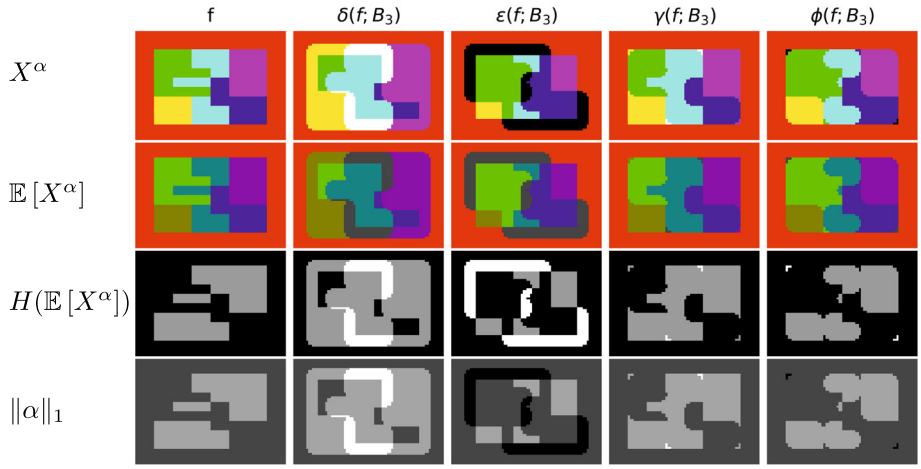

Fig. 2 Morphology on Dirichlet distributions. The top left image is an RGB representation of an image f with three categories, where the colors red, green, and blue correspond to points very close to the vertices of Δ^2 and the remaining colors are mixtures of these three colors. The first row is the Dirichlet distribution. The second row is the probability vectors obtained from (25). The third row is entropy of the probability vectors, and the fourth row is the magnitude (l_1 norm) of the parameter vectors. We can see that dilation increases both entropy and magnitude, whereas erosion decreases magnitude and increases or decreases entropy depending on the local distribution (Color figure online)

category, we would expect this region to become green in the probability image, but in the Dirichlet space these pixels already have the same green value as the green region, so they are unaffected by the dilation. We could partly solve this by carefully setting the α values, e.g., setting the pixels with only green to have very large green values. However, if our goal is to work on categorical distributions, this becomes too large a burden to be practical and we now turn our attention to morphological operators that work directly on categorical distributions.

3.2 Morphology on Categorical Distributions

Recall from Sect. 2.6 that for a set of $n+1$ categories, $C = \{c_1, c_2, \ldots, c_{n+1}\}$, the categorical distribution over these categories is completely determined by a point in the n-simplex $\Delta^n = \{\pi \in \mathbb{R}^{n+1} \mid \pi_k \geq 0, \sum \pi_k = 1\}$, where π_k is the probability of c_k. An image f is then defined as in (1) with $\Gamma = \Delta^n$. Operations are again defined on a single category at a time. Let B_r be a closed ball of radius r centered at the origin and let i be the category we operate on. Let $f_k(x) = f(x)_k$ be the probability of observing category c_k in pixel x and let $\omega_k(x) = 1 - f_k(x)$.

3.2.1 Dilation

For the dilated category i, the operation is the same as standard grayscale dilation. For the remaining set of categories, the operation is a rescaling to ensure that the probabilities sum to one, while the conditional probabilities $\Pr(k = j | x, k \neq i)$, for $j \neq i$ are unchanged

$$\delta_i(f; B_r)(x)_k = \begin{cases} \delta(f_k; B_r)(x), & \text{if } k = i, \\ [1 - \delta(f_i; B_r)(x)]\frac{f_k(x)}{\omega_i(x)}, & \text{if } k \neq i. \end{cases} \tag{30}$$

If $\delta(f_i; B_r) = 1$, then the conditional probabilities are not defined and we simply set the probabilities to $1 - \delta(f_i; B_r) = 0$. This definition is the same as (17) and equivalent to the definition from [6].

3.2.2 Erosion

Erosion is defined similarly to dilation, with the exception of the case when $f_i(x) = 1$ where we cannot rescale the remaining categories because $\omega_i(x) = 0$

$$\epsilon_i(f; B_r)(x)_k = \begin{cases} \epsilon(f_k; B_r)(x) & \text{if } k = i \\ [1 - \epsilon(f_i; B_r)(x)]\frac{f_k(x)}{\omega_i(x)} & \text{if } k \neq i \wedge f_i(x) < 1 \\ [1 - \epsilon(f_i; B_r)(x)]\frac{\theta(f_k, B_r)(x)}{\sum_{j \neq i} \theta(f_j, B_r)(x)} & \text{if } k \neq i \wedge f_i(x) = 1 \end{cases} \tag{31}$$

The function θ must only depend on the neighborhood defined by B_r and defined such that $\epsilon(f_i; B_r)(x) < 1 \implies [\exists k \neq i](\theta(f_k, B_r)(x) > 0)$. In addition, we require that, when disregarding discretization issues, eroding with $B_{r+\rho}$ is equivalent to first eroding with B_r and then eroding with B_ρ

$$\epsilon_i(\epsilon_i(f, B_r), B_\rho)(x) = \epsilon_i(f, B_{r+\rho})(x). \tag{32}$$

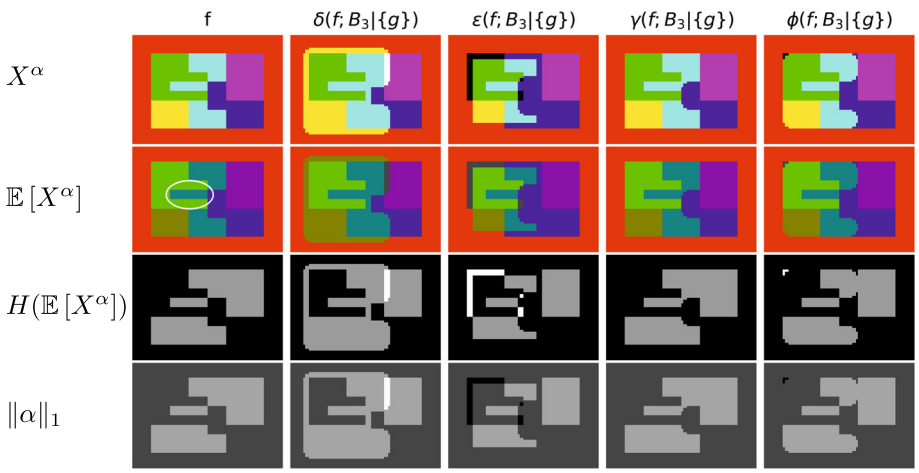

Fig. 3 Morphology on Dirichlet distributions using a subset of categories, in this case the green category $\{g\}$. See also Fig. 2 (Color figure online)

Since θ is only used in the case where $f_i(x) = 1$, we must have that

$$\epsilon(f_i; B_r)(x) < 1 \implies \frac{\theta(f_k, B_{r+\rho})(x)}{\sum_{j\neq i}\theta(f_j; B_{r+\rho})(x)} = \frac{\theta(f_k; B_r)(x)}{\sum_{j\neq i}\theta(f_j; B_r)(x)} \tag{33}$$

So θ must only depend on the smallest possible neighborhood B_{r^*} where $\epsilon(f_i; B_{r^*}) < 1$, leading to

$$\theta(f_k, B_r)(x) = \delta(f_k; B_{r^*})(x) \tag{34}$$
$$r^* = \arg\min_{r'>0} r', \text{ s.t. } \epsilon(f_i; B_{r'})(x) < 1.$$

This amounts to picking the closest category as suggested for crisp categorical images in [6, 7], although without the need for breaking ties since multiple closest categories are now handled by rescaling. In Appendix 1, we show that these definitions have the same properties as the definitions in [7] for operating on n-ary images.

An example of the proposed operations is provided in Fig. 4, where we operate on the green category. Compared to morphology on Dirichlet distributions using subsets in Fig. 3, the operations now work directly on the probabilities, making it much easier to understand and control.

4 Protected Morphological Operations

In [2], the authors introduce the concept of protected morphological operations, where a subset of categories are protected from being updated. Here we adapt the idea of protected morphological operations to categorical distributions and define protected dilation and erosion.

Let L be a set of categories, we then write $\epsilon_i(f; B_r|L)$ for an erosion of i that protects L. Let $J = C\backslash(\{i\} \cup L)$ be the set of categories that are not protected nor operated on. Let $f_K(x) = \sum_{k\in K} f_k(x)$ be the sum over a set of categories $K \subset C$. If L is empty, or $[\forall x](f_L(x) = 0)$, protected operations reduce to their non-protected counterparts. Because L can change the topology of the domain, we cannot just define operations based on Euclidean distance. Instead we introduce a distance function $d_\Omega(x, y)$, which computes the distance from x to y on the domain Ω. If $\Omega = \mathbb{Z}^d$, then $d_\Omega(x, y)$ is the Euclidean distance. Computing exact Euclidean distance on a Euclidean domain with holes is non-trivial. Here we use the simplified fast marching method (FMM) from [11] with the update rule defined in [12], which results in a small approximation error. For brevity, when possible we leave out function application and write f instead of $f(x)$ in the following.

4.1 Protected Dilation

Let $\Omega_p = \{x \in \mathbb{D} \mid f_L(x) \leq 1 - p\}$, this is the part of f where it is possible to set $f_i = p$. Protected dilation is then defined as

$$\delta_i(f; B_r|L)(x)_k = \begin{cases} f_k & \text{if } k \in L \\ \min\left(1 - f_L, \max_{p\in(0,1]} \max\{f_i(y) \mid d_{\Omega_p}(x, y) \leq r\}\right) & \text{if } k = i \\ [1 - f_L - \delta_i(f; B_r|L)_i]\frac{f_k}{f_J} & \text{otherwise} \end{cases} \tag{35}$$

4.2 Protected Erosion

Protected erosion is defined similarly to protected dilation, with the added complication of normalization

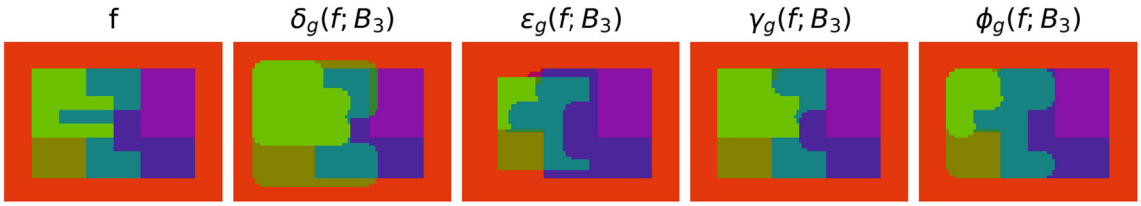

Fig. 4 Morphology on categorical distributions. Here we operate on the green category g (Color figure online)

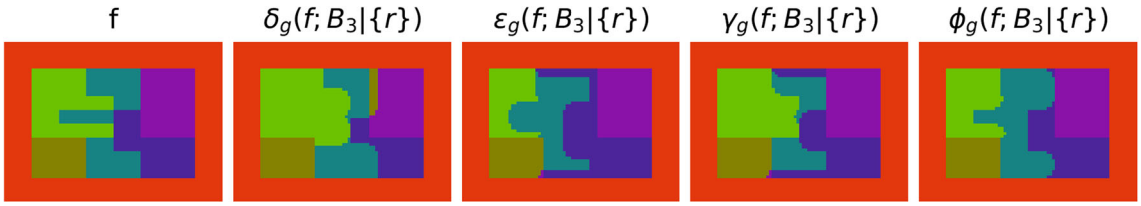

Fig. 5 Protected morphology on categorical distributions. The red category $\{r\}$ is protected while we operate on the green category g (Color figure online)

$$\epsilon_i(f; B_r|L)(x)_k = \begin{cases} f_k & \text{if } k \in L \\ f_k & \text{if } \max_{p\in(0,1]} \max\{f_J(y) \mid d_{\Omega_p}(x,y) \le r\} = 0 \\ \min_{p\in(0,1]} \min\{f_i(y) \mid d_{\Omega_p}(x,y) \le r\} & \text{if } k = i \\ [1 - f_L - \epsilon_i(f; B_r|L)_i]\frac{f_k}{f_J} & \text{if } k \in J \wedge f_J > 0 \\ [1 - f_L - \epsilon_i(f; B_r|L)_i]\frac{\theta(f_k)}{\sum_{j\in J}\theta(f_j)} & \text{if } k \in J \wedge f_J = 0 \end{cases} \tag{36}$$

The first case ensures that all protected categories are unchanged. The second case ensures that a pixel x is not updated, unless there is a path, not blocked by f_L, to a pixel y with $f_J(y) > 0$. The importance of this is easily seen by considering the case where f_i varies in an region, but $f_i + f_L = 1$ in the region. The third case states that if there is such a path, then it can be eroded. The fourth and fifth cases handle normalization. The θ function is defined in a similar manner as for non-protected erosion in (34),

$$\theta(f_k)(x) = \max_{p\in(0,1]} \max\{f_k(y) \mid d_{\Omega_p}(x,y) \le r^*\} \tag{37}$$
$$r^* = \arg\min_{r'>0} r' \ , \ \text{s.t. } [1 - f_L - \epsilon_i(f; r'|L)_i(x)] > 0.$$

An example of these operations is provided in Fig. 5, where the red category is protected while we operate on the green category. Compared to the non-protected operations in Fig. 4, we can see that changes are restricted to the green and blue categories.

5 Examples

The first example illustrates how morphology on categorical distributions (Sect. 3.2) can be used to remove noisy predictions. The second example illustrates how protected morphology on categorical distributions (Sect. 4) can be used to model annotator bias.

5.1 Removing Noisy Predictions

Despite the impressive performance of neural networks for segmentation, the results are rarely perfect. Figure 6 shows part of an electron microscopy image of the hippocampus, along with multi-class predictions and segmentations obtained from [13]. Notice the noisy mitochondria predictions resulting in misclassifications highlighted in Fig. 6a. We can remove these misclassification by opening the mitochondria class before the final classification. Figure 7 shows the opened predictions along with the final classifications. Notice in particular how the errors in circle 2 in Fig. 7a are fixed, such that the vesicle (teal) and the endoplasmic reticulum (yellow) are separated by cytosol. This would have been very difficult to achieve by working directly on the final segmentations. That the vesicle and endoplasmic reticulum are probably misclassified just illustrates that not all things should be fixed in post-processing.

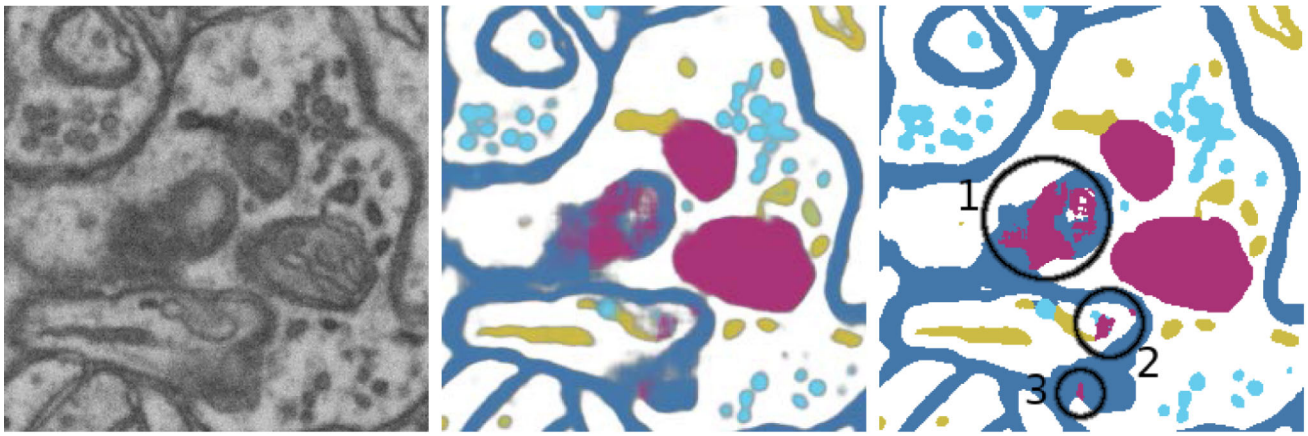

Fig. 6 Electron microscopy image of the hippocampus with predictions of five classes: cytosol (white), membrane (blue), mitochondria (purple), endoplasmic reticulum (yellow) and vesicle (teal). By examining neighboring slices, the areas 1–3 have been confirmed to wrongly contain mitochondria predictions (Color figure online)

Fig. 7 Fixing mitochondria misclassifications by opening the mitochondria predictions with B_{12}

5.2 Modeling Annotator Bias

Expert annotation is the gold standard in most clinical practice as well as for evaluating computer methods. However, annotation tasks are inherently subjective and prone to substantial inter-rater variation [14, 15]. When investigating the influence of this variation on statistics and decisions, it can be interesting to consider specific hypotheses regarding the variation. Consider the brain tumor annotation in Fig. 8. The annotation is derived from the QUBIQ[1] challenge brain tumor dataset, where three annotators each annotated whole tumor, tumor core and active tumor. From this, we obtain an image with four categories: background, edema, active core and inactive core. Although the annotators have a high level of agreement, there is still substantial variation in the extent of edema and in how much of the tumor core is active.

Using protected dilation, we can, for example, hypothesize how the merged annotation would appear under the assumption that the tumor core is oversegmented, but the active part is undersegmented. Figure 9 shows the results where we first dilate the active core while protecting edema and background, then dilate edema while protecting background. This would allow us to easily investigate if statistical differences in a case–control study could be explained by biased annotations.

6 Discussion and Conclusion

We have provided a thorough review of morphology on categorically valued images. Based on this, we have defined morphology on Dirichlet distributions and morphology on categorical distributions. Inspired by [2], we have further defined protected morphology on categorical distributions. We have demonstrated the behavior of the proposed operations and shown how they can be used in real-world applications such as noise removal in multi-class predictions and modeling annotator bias.

[1] https://qubiq.grand-challenge.org/.

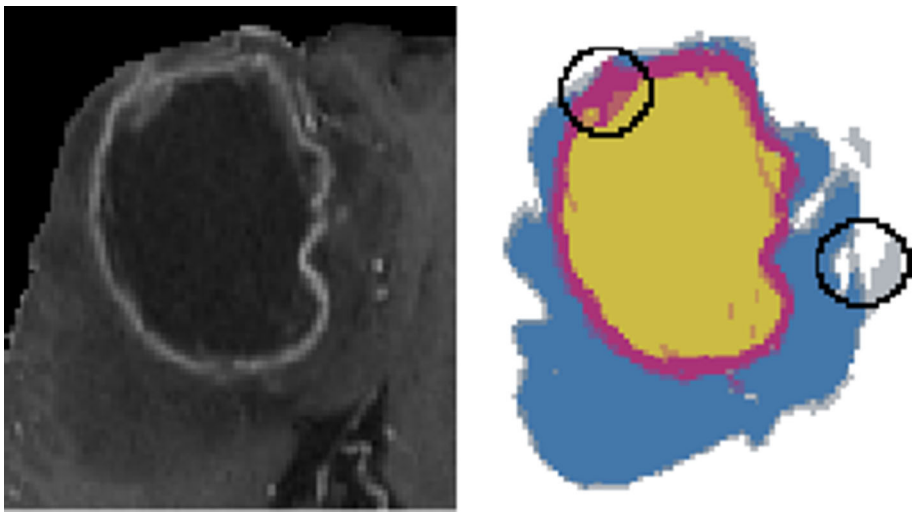

Fig. 8 Inter-rater variation in annotation of brain tumors. White is background, blue edema, yellow inactive core and purple active core. Variation is indicated by color mixing. The black circles highlights two regions with large variation (Color figure online)

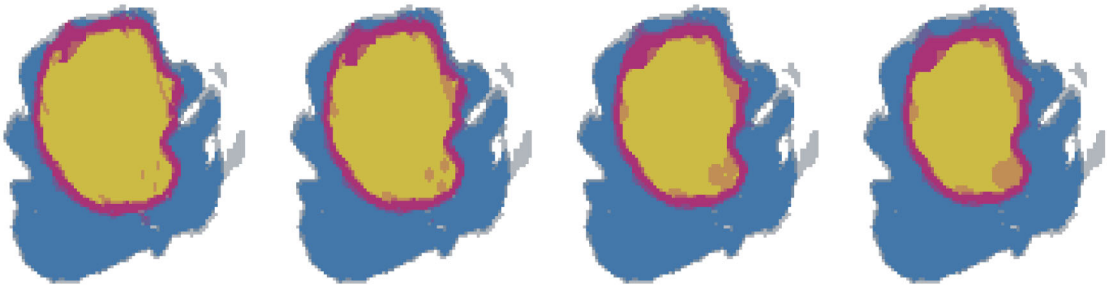

(a) Original (b) B_1 (c) B_2 (d) B_3

Fig. 9 What could the annotation look like if the core was oversegmented, but the active part undersegmented? Dilation of active core while protecting edema and background, followed by dilation of edema while protecting background using B_1, B_2, B_3

The definition of dilation is straightforward and no obvious alternatives present themselves. This is not so for erosion. In our definition, erosion corresponds to conditioning on a change in probability of the eroded category. An equally valid approach would be to also condition on where this change came from. Instead of simply rescaling the categories with nonzero mass, we could include information from the neighborhood. For example, when eroding i we would fill the difference $f_i(x) - \epsilon(f_i; B_r)(x)$ based on the pixels that contribute to the difference, that is, those with minimum mass for i. This would result in smoother boundaries, which could be a better representation of uncertainty. A downside is that categories can leak into each other, leading to undesirable results.

In this work, we have focused on the basic morphological operations, dilation and erosion, and their compositions, closing and opening. A logical next step is to investigate more complex morphological operations, such as the morphological gradient, which may be used to investigate spatial relationship between categories by measuring the change in one category as a function of change in another category.

We have defined protected versions of dilation and erosion. From these, we could define opening and closing in the standard way. Alternatively, by changing which categories are protected for dilation and erosion we get more control over how a category is opened or closed. In [2], the authors explore similar ideas for the so-called tunneling and bridging operations on their set-based morphology, which would be interesting to consider in the context of categorical distributions.

Our aim in this work was to bring morphological operations to probabilistic representations of categorical images. These representations can be considered as generative processes that can be sampled. Naive sampling will result in noisy and unrealistic samples. Combining the sampling process with the proposed morphological operations could be an easy approach to obtain smoother and more realistic samples.

In summary, morphology is an indispensable tool for post-processing segmentations. Extending morphology to categorical images and their probabilistic counterparts presents a particular problem since there is in no inherent ordering of categories. In this paper, we have proposed to view categorical images as images of categorical distributions and defined morphological operations that are consistent with this view.

Funding Open access funding provided by Royal Danish Library.

Data Availability Statement The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

Appendix A: Proofs for Sect. 3.2

We closely follow Section 5.2 in [7] by defining a preorder $\preceq_i$ on $\mathcal{F}$ and showing that our definitions of dilation and erosion form an adjunction in this preorder. We then show that their compositions are an opening ($\gamma_i = \delta_i\epsilon_i$) and a closing ($\phi_i = \epsilon_i\delta_i$), where we define an opening as an operator that is increasing, anti-extensive and idempotent, and a closing as an operator that is increasing, extensive and idempotent

$$\text{a } f \preceq_i g \implies \gamma_i f \preceq_i \gamma_i g, \quad \text{b } \gamma_i f \preceq_i f, \quad \text{c } \gamma_i\gamma_i f = \gamma_i f \tag{A1}$$

$$\text{a } f \preceq_i g \implies \phi_i f \preceq_i \phi_i g, \quad \text{b } f \preceq_i \phi_i f, \quad \text{c } \phi_i\phi_i f = \phi_i f \tag{A2}$$

For two images $f, g \in \mathcal{F}$, we define the preorder $\preceq_i$ as

$$f \preceq_i g \iff [\forall x](f_i(x) \leq g_i(x)) \tag{A3}$$

This preorder is not antisymmetric, as we can have $f \preceq_i g$ and $g \preceq_i f$, but not $f = g$.

Theorem 1 *δ_i and ϵ_i form an adjunction in the preorder $\preceq_i$*

$$\delta_i(f; B_r) \preceq_i g \iff f \preceq_i \epsilon_i(g; B_r) \tag{A4}$$

Proof Since categories $j \neq i$ have no influence on $\preceq_i$, we only need to consider the case where $k = i$ in (30) and (31). These cases are standard grayscale dilation and erosion that form an adjunction. □

Lemma 1 *δ_i and ϵ_i are increasing in $\preceq_i$*

Proof Follows by the same argument as for Theorem 1. □

Theorem 2 *For a fixed structuring element B_r, δ_i and ϵ_i satisfy*

$$f \preceq_i \epsilon_i(\delta_i(f)) \tag{A5}$$

$$\delta_i(\epsilon_i(f)) \preceq_i f \tag{A6}$$

$$\delta_i(\epsilon_i(\delta_i(f))) = \delta_i(f) \tag{A7}$$

Proof (A5) and (A6) follow from substitution into (A4) with $g = \delta_i(f)$ and $f = \epsilon_i(g)$

$$\delta_i(f; B_r) \preceq_i \delta_i(f; B_r) \implies f \preceq_i \epsilon_i(\delta_i(g; B_r); B_r) \tag{A8}$$

$$\delta_i(\epsilon_i(g; B_r); B_r) \preceq_i g \impliedby \epsilon_i(g; B_r) \preceq_i \epsilon_i(g; B_r) \tag{A9}$$

To show (A7), we consider the three cases (I) $k = i$, (II) $k \neq i$ with $[\forall x](f_i(x) < 1)$, and (III) $k \neq i$ with $[\exists x](f_i(x) = 1)$. For brevity, we leave out the structuring element B_r, pixel index x and parentheses from operator and function application in the following.

For (I) $k = i$, we can directly substitute the definitions in (30) and (31) to get

$$(\delta_i\epsilon_i\delta_i f)_i = \delta\epsilon\delta f_i \tag{A10}$$

$$= \delta f_i \tag{A11}$$

$$= (\delta_i f)_i, \tag{A12}$$

where the second step follows from the properties of the standard grayscale operations.

For (II) $k \neq i$ with $[\forall x](f_i(x) < 1)$, we can ignore the third case in (31). After substitution and cancelation of terms, we get,

$$(\delta_i\epsilon_i\delta_i f)_k = [1 - \delta f_i]\frac{f_k}{f_J} = (\delta_i f)_k \tag{A13}$$

For (III) $k \neq i$ with $[\exists x](f_i(x) = 1)$, the third case in (31) is only relevant when $(\delta_i f)_i(x) = (\delta f_i)(x) = 1$, which leads to both sides of (A7) being zero

$$(\delta_i\epsilon_i\delta_i f)_k = [1 - \delta f_i]\frac{\theta(\epsilon_i\delta_i f)_k}{\sum_{j\in J}\theta(\epsilon_i\delta_i f)_j} = 0 \tag{A14}$$

$$(\delta_i f)_k = [1 - \delta f_i]\frac{f_k}{f_J} = 0 \tag{A15}$$

So $\delta_i\epsilon_i\delta_i f = \delta_i f$ □

Corollary 1 *$\gamma_i = \delta_i\epsilon_i$ is an opening and $\phi_i = \epsilon_i\delta_i$ is a closing*

Proof Corollary 1. in [7] □

References

1. Serra, J.: Morphological filtering: an overview. Signal Process. **38**(1), 3–11 (1994)

2. Busch, C., Eberle, M.A.: Morphological operations for color-coded images. Comput. Graph. Forum **14**, 193–204 (1995)
3. Köppen, M., Nowack, C., Rösel, G.: Pareto-morphology for color image processing: a comparative study of multivariate morphologies. Pattern Recognit Image Anal. **10**(4), 478–491 (2000)
4. Hanbury, A.G., Serra, J.: Morphological operators on the unit circle. IEEE Trans. Image Process. **10**(12), 1842–1850 (2001)
5. Ronse, C., Agnus, V.: Morphology on label images: flat-type operators and connections. J. Math. Imaging Vis. **22**, 283–307 (2005)
6. Chevallier, E., Chevallier, A., Angulo, J.: N-ary mathematical morphology. Math. Morphol. Theory Applicat. **1**, 42–59 (2016)
7. van de Gronde, J., Roerdink, J.B.T.M.: Chapter three—nonscalar mathematical morphology. In: Hawkes, P.W. (ed.) Advances in Imaging and Electron Physics. Advances in Imaging and Electron Physics, vol. 204, pp. 111–145. Elsevier, Amsterdam (2017)
8. Grossiord, E., Naegel, B., Talbot, H., Najman, L., Passat, N.: Shape-based analysis on component-graphs for multivalued image processing. Math. Morphol. Theory Appl. **3**(1), 45–70 (2019)
9. Aptoula, E., Lefèvre, S.: A comparative study on multivariate mathematical morphology. Pattern Recognit. **40**, 2914–2929 (2007)
10. Peters II, R.A.: Mathematical morphology for angle-valued images. In: Nonlinear Image Processing VIII, vol. 3026, pp. 84–94 International Society for Optics and Photonics (1997)
11. Jones, M.W., Baerentzen, J.A., Sramek, M.: 3d distance fields: a survey of techniques and applications. IEEE Trans. Visual Comput. Graphics **12**(4), 581–599 (2006)
12. Rickett, J., Fomel, S.: A second-order fast marching eikonal solver. Stanf. Explor. Proj. Rep. **100**, 287–293 (1999)
13. Stephensen, H.J.T., Svane, A.M., Villanueva, C.B., Goldman, S.A., Sporring, J.: Measuring shape relations using r-parallel sets. J. Math. Imaging. Vis. 63, 1069–1083 (2021). https://doi.org/10.1007/s10851-021-01041-3
14. Joskowicz, L., Cohen, D., Caplan, N., Sosna, J.: Inter-observer variability of manual contour delineation of structures in CT. Eur. Radiol. **29**(3), 1391–1399 (2019)
15. Becker, A.S., Chaitanya, K., Schawkat, K., Muehlematter, U.J., Hötker, A.M., Konukoglu, E., Donati, O.F.: Variability of manual segmentation of the prostate in axial t2-weighted MRI: a multi-reader study. Eur. J. Radiol. **121**, 108716 (2019)

Silas Nyboe Ørting received his Ph.D. degree from the Department of Computer Science, University of Copenhagen, Denmark (DIKU) in 2019. Since 2019 he has been a Post Doc at the Center for Quantification of Imaging Data at MAXIV (QIM) and at DIKU.

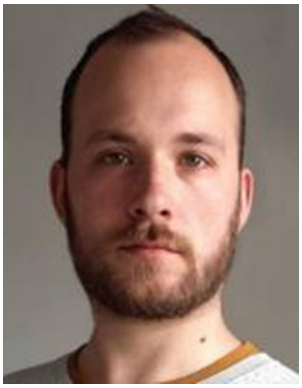

Hans Jacob Teglbjærg Stephensen received the Ph.D. degree in Computer Science in 2021 from the University of Copenhagen. His main focus is on method development, algorithm design and software development to solve geometrical, spatial and stochastic problems in bioimaging. He is now doing neural network design in Biometrics at Fingerprint Cards AB with a special focus on Attention Mechanisms and Reinforcement Learning.

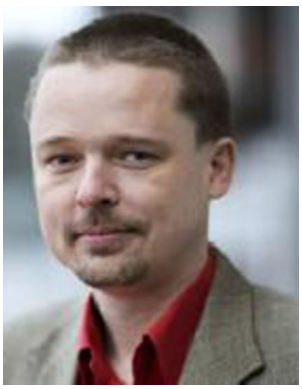

Jon Sporring received his Ph.D. degree from the Department of Computer Science, University of Copenhagen, Denmark (DIKU) in 1998. He was an assistant research professor at 3D-Lab, School of Dentistry, University of Copenhagen, an associate professor and since 2018 a full professor at DIKU. Jon has been a visiting researcher at IBM Research Center, Almaden, USA, at Computer Vision and Robotics Lab at Foundation for Research & Technology - Hellas, Greece, and at School of Computer Science, McGill University, Canada. He has been a Senior Researcher at Nordic Bioscience a/s and co-founded DigiCorpus Aps. For 10 years he served as Deputy head of the Department for Research at DIKU.