

SoftMorph: Differentiable Probabilistic Morphological Operators for Image Analysis

Lisa Guzzi, Maria A. Zuluaga, *Senior Member IEEE*, Fabien Lareyre, Gilles Di Lorenzo, Sébastien Goffart, Andrea Chierici, Juliette Raffort, Hervé Delingette

Abstract—Morphological operators are crucial in image analysis. Their integration into deep learning pipelines could improve performances by extracting or enhancing important image features, either within network architectures or loss functions. However, the difficulties in rendering those operators differentiable hinder their integration. In this paper, we present SoftMorph, a novel framework designed to convert any binary morphological operator defined as a Boolean expression into its differentiable and probabilistic counterpart, compatible with gradient-based optimization. Specifically, we define probabilistic operators as the expectation of the binary operator with respect to the probability of generating each binary configuration. This expectation can be computed trivially from the truth table of the binary morphological filter, as a multi-linear polynomial function. Moreover, we approximate the probabilistic operators with quasi-probabilistic operators directly translated from the Boolean expressions leveraging Fuzzy logic. These quasi-probabilistic operators therefore maintain the computational complexity of the original binary operator. We demonstrate the efficiency and reliability of our method through validation experiments, and evaluate the backpropagation capability of the proposed operators. Finally, we showcase several applications of morphological operators integrated into neural networks for image segmentation tasks.

Index Terms—Morphological operations, Image analysis, Deep Learning, Fuzzy logic.

I. INTRODUCTION

MATHEMATICAL morphology is a fundamental theory used to process images by extracting or enhancing information based on geometrical shapes, structures, and spatial

Lisa Guzzi is with the Université Côte d’Azur, Inria, Epione Team, Sophia Antipolis, France, and also with the Data Science Department, EURECOM, Sophia Antipolis, France, and also with the University Hospital of Nice, Nice, France (e-mail: lisa.guzzi@inria.fr).

Maria A. Zuluaga is with the Data Science Department, EURECOM, Sophia Antipolis, France.

Fabien Lareyre is with the Department of Vascular Surgery, Hospital of Antibes Juan-les-Pins, Antibes, France, and also with the Université Côte d’Azur, CNRS, UMR7370, LP2M, Nice, France.

Gilles di Lorenzo is with the Department of Vascular Surgery, Hospital of Antibes Juan-les-Pins, Antibes, France.

Sébastien Goffart is with the Université Côte d’Azur, Inria, Epione Team, Sophia Antipolis, France, and also with the University Hospital of Nice, Nice, France.

Andrea Chierici is with the University Hospital of Nice, Nice, France.

Juliette Raffort is with the University Hospital of Nice, Nice, France, and also with the Université Côte d’Azur, CNRS, UMR7370, LP2M, Nice, France.

Hervé Delingette is with the Université Côte d’Azur, Inria, Epione Team, Sophia Antipolis, France.

(Corresponding author: Lisa Guzzi.)

The code is available at <https://github.com/lisaGUZZI/SoftMorph2>.

This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P31A-0002. The authors are grateful to the OPAL infrastructure from Université Côte d’Azur.

relationships [1]. Its methodology revolves around two fundamental operations: dilation expands the foreground objects in an image, increasing their size and connectivity, while erosion shrinks the foreground, eliminating small objects. Dilation and erosion can be combined or iterated to achieve more complex operations. These operations can enhance the essential shape and characteristics of an image, and extract shape features such as edges, holes, cracks, and corners [2]. Morphological operations operate on a Structuring Element (SE) that defines the pattern of the image feature to be processed, encompassing a large variety of shapes and sizes. Originally defined on binary images, morphological operations have subsequently been extended to gray-scale images. An important extension involves *fuzzy morphological operators* [3], [4], which adapt the operations to continuous values, mainly by replacing erosion and dilation with minimum and maximum operations.

Therefore, morphological operations are particularly useful for many image analysis and processing tasks such as image segmentation [5]–[8], feature extraction [9]–[12], noise reduction [13]–[16], or image enhancement [17]–[20]. For instance, morphological filters were applied in radar imaging for image restoration by effectively removing noise while retaining critical information [21]. Similarly, new morphological operations were designed for edge detection in noisy medical imaging [22]. In medical applications, skeletonization is the operation that extracts the centerline of a segmented structure such as vessels or the heart. It is often used to assess the organ’s topology, connectivity and trajectory [23].

Morphological operations can enhance the performance of convolutional neural networks (CNNs) by extracting meaningful structural information to analyze images. Researchers have started integrating these operations into loss functions to improve specific characteristics of segmentation tasks [24], [25]. Moreover, some operations such as dilation and erosion could be integrated as pre-processing, intermediate, or post-processing layers within CNN architectures. This allows to handle different patterns and tasks than traditional convolutional layers, and hereby refine complex image analysis [26], [27]. However, traditional binary morphological operations are based on Boolean expressions, and some are implemented using minimum and maximum operations, ill-suited for the gradient-based optimization methods used in deep learning models. Indeed, that optimization process requires differentiable functions that are smooth and continuous to compute gradients effectively during the training phase [28], [29]. Binary decisions and min/max operations introduce discontinuities, making them inherently non-differentiable.

To address this, fuzzy morphological neurons were designed to align operators with gradient-based optimization, by approximating the *soft minimum and maximum functions* [30], [31]. Standard min/max operations can also be seen as piece-wise differentiable, directly replacing erosion and dilation with *min and max-pooling layers*, and combining them to achieve more complex operations [24], [25], [32], [33]. Although these methods can successfully incorporate morphological filters in neural networks, they cannot handle general morphological operators. Additionally, the use of min and max-pooling only can lead to discontinuous skeletons and homotopy inaccuracies [34], [35]. Others have focused on *Learning Morphological Operations* by training CNNs to replicate specific morphological filters [36], [37]. These models can backpropagate gradients and be integrated into deep learning pipelines. However, they correspond to approximative morphological filters trained on a specific domain and may require fine-tuning. Lastly, one may consider *Convolution-like approaches* such as successive convolutional layers with pre-defined kernels to replicate a particular operation [34], but this method can be hard to apply to some complex operations.

To the best of our knowledge, there is no existing approach to seamlessly integrate any morphological operation with any SE into a CNN in a smooth and differentiable manner.

To address these limitations, we recently proposed a novel framework [38] to extend binary morphological operations on probability maps that can be seamlessly integrated into neural networks either as a loss function or as a final morphological layer. This framework translates any morphological operation based on Boolean expressions into a single multi-linear or proxy polynomial. The resulting soft morphological filters are differentiable, require no hyperparameter tuning, and can be derived from any binary morphological filter. We demonstrated their effectiveness on two medical segmentation applications, comparing them against state-of-the-art morphological filters.

In this paper, we significantly extend the work by [38] to SoftMorph, a family of operators based on various fuzzy logics, providing a generalized approach for translating any binary morphological operator with any SE into differentiable and probabilistic equivalents. The new contributions of this work are summarized as follows:

- *We provide a detailed explanation for the definition of probabilistic morphological filters* as the expectation of the binary filter over the probability of generating each possible binary configuration, expressed as a multi-linear polynomial deduced from the truth table.
- *We define quasi-probabilistic operators* for intractable truth tables of binary morphological filters. We use various fuzzy logic operators to convert the Boolean expression defining the morphological operation into a soft, differentiable expression. This quasi-probabilistic operator approximates the probabilistic operator while maintaining the complexity of the original binary filter. Unlike previous works on fuzzy morphological operators that applied to erosion and dilation only, our proposed expressions apply to any Boolean expression, hence easily adaptable to new custom operations.

- *We quantify the backpropagation capability* of the SoftMorph operators to ensure their smooth integration into deep learning pipelines.
- *We extend the applications of this framework* beyond the medical domain for various segmentation tasks, across 7 different 2D and 3D datasets. We integrate the SoftMorph operators as final layers of a neural network’s architecture and in loss function applications.

Overall, these contributions make the SoftMorph framework a powerful tool for integrating morphological filters into deep learning models.

In section II, we review related works to make soft, differentiable morphological operations and their integration into CNNs. Section III details the methodology of SoftMorph. We provide experiments in section IV-B to validate our morphological operations and the backpropagation capability of the proposed operator representations, and in section IV-C we demonstrate several applications of the proposed framework, specifically for segmentation tasks.

II. RELATED WORKS

A. Continuous Fuzzy Morphological Operators

Mathematical morphology has been extended from binary to gray-scale image processing partly through fuzzy set theory, as introduced by [3]. In this context, gray-scale images are treated as fuzzy sets, where each pixel value is interpreted as a degree of membership to a set, and morphological operators are redefined within that fuzzy space. This approach enables more nuanced image analysis and has led to various representations, employing fuzzy conjunction and disjunction operators. The most notable are the minimum and maximum operations to represent erosion and dilation respectively [4].

B. Differentiability of Morphological Filters

Subsequent works have addressed the differentiability of morphological filters through different techniques to integrate them into deep learning pipelines.

Counter Harmonic Mean: The counter harmonic mean (CHM) has been studied to approximate the erosion and dilation operations in [39]. This approach involves raising the pixel values of an image to a certain power, combining them with weights, and normalizing the result. Thus, morphological dilation and erosion correspond to the limit cases when the power tends to $+\infty$ (dilation) and $-\infty$ (erosion). The differentiable nature of the CHM formulation facilitates gradient-based optimization and its application in CNNs [40]–[42]. A morphological neuron, the so-called Pconv layer, implements the CHM filter to learn the appropriate power value and weight parameters to optimize the operations. However, this method is subject to exploding gradients due to the power function [32].

Log-Sum-Exp: The morphological filters have also been approximated using differential approximations of the minimum and maximum operations with the Log-Sum-Exp function. In [31] and [43], it has been tailored to optimize the SE of the morphological operations. In [44], it is furthermore improved to learn non-flat SEs by applying bias variables to correct the rounding errors caused by the soft approximation of min/max.

Non-smooth minimum and maximum operators: In theory, minimum and maximum operations are not fully differentiable. While they are differentiable almost everywhere, they lack a well-defined gradient at points where multiple arguments of the operations are equal. In practice, these operations are treated as piece-wise differentiable. That means that the gradient is propagated only towards the maximum or minimum element during backpropagation. When multiple arguments share the same maximum or minimum value, different strategies are employed depending on the context. In typical minimum/maximum operations, the gradient is distributed equally between those arguments. In commonly used max-pooling layers, the gradient is assigned exclusively to the first occurrence of the maximum value within each pooling window. Therefore, some researchers argue that these non-smooth operators can be used in deep learning pipelines to replicate the erosion and dilation operators, while still ensuring the optimization of models. For example, in [24], the combination of min and max-pooling layers allows to obtain the borders of the foreground. Similarly, the pooling layers are iterated to extract the skeleton in [25]. It is obtained by iteratively getting the difference between the erosion of the image and the opening of that erosion. However, this method leads to inconsistencies creating disconnected skeletons and topological errors. As the standard pooling layers are fixed with a square SE, [32] and [33] propose to replace them with learned morphological pooling layers that can optimize the SE. In [32] non-flat SEs are optimized to learn the exact morphological operators.

Learning-Based Emulation of Morphological Operations: Morphological operations trained from data by neural networks inherently support gradient-based optimization. It can provide differentiable solutions for complex operations such as skeletonization, trained in CNNs for example by [36] and [37]. Other works try to optimize the morphological operation and the SE learned from some training data, by alternating dilation and erosion-like convolution layers [45]. However, training from data is always prone to domain shift and can produce topological errors [34].

Convolution-like approaches: A specific operation can be emulated by applying several convolution layers with pre-defined kernels to detect expected configurations and patterns in an image. For example, [34] detects simple points, corresponding to pixels that can be removed from the foreground to obtain a skeleton, only using convolutional and matrix operations followed by non-linear functions. They defined specific kernels to recognize 2D and 3D configurations and their rotational equivalent to check the presence of simple points. That process requires 57 convolutions to match the Boolean rules for simple point detection. The reparametrization trick is employed as the detection is based on binary criteria. Depending on the operation to replicate, this method can be computationally expensive due to the high number of convolutions required. Designing specific kernels to detect various configurations and their rotational equivalents can be complex and time-consuming, potentially limiting its adaptability.

C. Morphological Filters in Deep Learning applications

Through the aforementioned methods to approximate differentiable morphological operations defined on continuous values, several applications of these operators have been explored in neural networks [26], [46], [47].

Deep morphological neural networks: There has been a growing interest in replacing traditional convolutional layers of CNNs with morphological operations. Deep Morphological Neural Networks (DMNNs) leverage the inherent non-linearity of morphological operations to substitute the linear convolution and non-linear activation functions typically used in CNNs. These networks define morphological filters or layers that approximate the erosion or dilation operations and optimize the SE based on target data. In some works, the appropriate sequence of erosion and dilation within the network architecture is also learned [40], [44], [48]. DMNNs can alternate morphological layers with standard convolutional layers, and can also incorporate fully connected layers for classification tasks [27], [32], [41], [42], [44]. Other main applications of these networks include image restoration [30], [32], [33], [40], [48] and edge detection [32]. Besides, [42] consider these networks more interpretable than usual CNNs because the learned sequence of morphological operations can be explicitly recovered and analyzed.

Loss functions with morphological operators: Lastly, some morphological operations can extract specific image features to be used in the loss function of a CNN to improve particular characteristics of the model. The cDice loss function in [25] requires the extraction of the skeleton to compute its intersection with the foreground volume, to maximize the topological preservation in the segmentation of tubular structures. In [24], they extract the boundaries of the segmentation to minimize the perimeter difference of the prediction with the ground truth volume. The segmentation of small and thin structures is enhanced in [49] by integrating the white top-hat operation to detect small structures. The final loss corresponds to the weighted linear combination of the usual cross-entropy function and its application on small objects only.

Compared to existing methods, our approach does not involve optimizing the SE or learning the morphological operations. Instead, we focus on adapting any known binary operator into a differentiable and probabilistic form, that one might want to use in a CNN to optimize an image analysis problem. It provides a solution to precisely replicate the binary operator, that can be integrated either in the loss function of a CNN or as a post-processing layer within the network. One key innovation of our approach is defining the differentiable operators on probability maps, enhancing their applicability. Additionally, we ensure to define quasi-probabilistic operators that match the computational complexity of the binary operators. Previous methods either approximate those operations, resulting in a mismatch with the exact binary operator, or are challenging to apply to complex operations. Overall, our method offers a more precise and flexible solution to incorporate morphological operations into deep learning frameworks.

III. METHODOLOGY

A. Definition of binary morphological filters

We define an image X consisting of N voxels $\{X_n\}$, $n = 1 \dots N$ and its corresponding binary segmentation image Y with $Y_n \in \{0, 1\}$.

A binary morphological operator $F()$ is applied to the binary image Y resulting in a final binary segmentation Z such that $Z = F(Y) \in \{0, 1\}^N$.

This operator $F()$ takes as input k binary variables selected from the neighborhood function $\mathcal{N}(i, n)$, providing the index of the i th neighbor of the voxel n in Y as $Y_{\mathcal{N}(1,n)}, \dots, Y_{\mathcal{N}(k,n)}$. It outputs a binary variable $Z_n = F(Y_{\mathcal{N}(1,n)}, \dots, Y_{\mathcal{N}(k,n)}) \in \{0, 1\}$. The neighborhood function corresponds to the SE of F . In mathematical morphology, it defines the domain of the geometrical features processed in the morphological operation. For instance, in a 2D image, typical neighborhoods are defined as $k = 4 + 1$ (Fig 1a) or $k = 8 + 1$ (Fig 1b). Similarly, in 3D they are defined as $k = 6 + 1, 18 + 1$ or $26 + 1$.

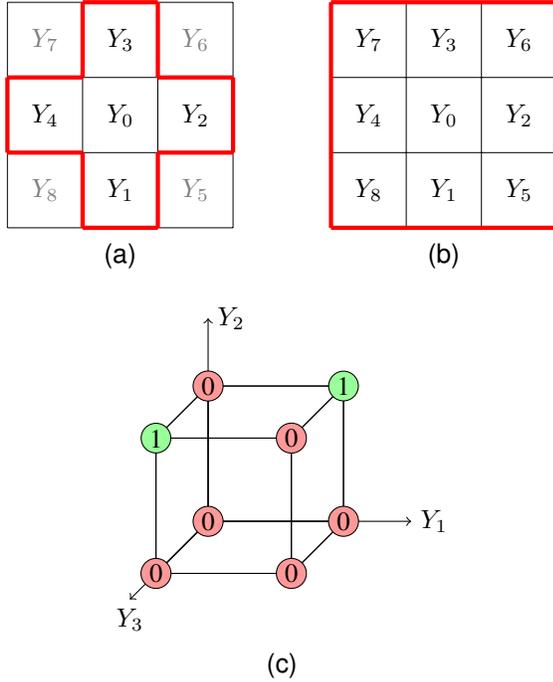


Fig. 1. Definition of the neighborhood variables in a 2D image. (a) $4 + 1$ neighborhood. (b) $8 + 1$ neighborhood. (c) Representation of the truth table of F_{Ex} on a hypercube. Each dimension corresponds to a binary input variable of the morphological operator.

The binary operator $F()$ is a Boolean function that has 2^k different possible input values and outputs a binary variable. The exhaustive list of those values $F(a) \in \{0, 1\}$, for $a \in \{0, 1\}^k$ is called the *truth table* of $F()$ and can be provided exhaustively for small values of k . Besides, it can be shown that any Boolean function can be written as a propositional formula involving the k binary variables with the logical operator AND (\wedge), OR (\vee) and NOT (\neg). In fact, two operators are sufficient in Boolean logic, since the AND and OR operators can be expressed with the remaining two

according to De Morgan's theorem [50]. Similarly, the XOR (\oplus) operator can be rewritten as: $(A \oplus B) = (A \vee B) \wedge \neg(A \wedge B)$.

Based on the notation defined in Fig.1, the dilation operator acting on a $4 + 1$ -neighborhood of a 2D image can be written as $f_{\text{Dil}} = Y_0 \vee Y_1 \vee Y_2 \vee Y_3 \vee Y_4$ whereas the erosion operator is $f_{\text{Ero}} = Y_0 \wedge Y_1 \wedge Y_2 \wedge Y_3 \wedge Y_4$.

As an example, the operator f_{Ex} will be reused throughout this section to illustrate key concepts. The operator f_{Ex} is defined on a 1×3 grid such that $Z_n = f_{\text{Ex}}(Y_{n-1}, Y_n, Y_{n+1}) = F_{\text{Ex}}(Y_1, Y_2, Y_3)$. It is computed with the Boolean expression:

$$F_{\text{Ex}} = Y_2 \wedge ((Y_1 \wedge \neg Y_3) \vee (\neg Y_1 \wedge Y_3))$$

Consequently, with $k = 2 + 1$, there are 8 possible input configurations represented on the 3-hypercube of Fig 1c.

B. Definition of probabilistic morphological filters

We want to extend these morphological filters to the output of segmentation algorithms. Typically, the output of a neural network consists of probabilities y_n , which correspond to the posterior probability $y_n = p(Y_n = 1|X) \in [0, 1]$ of the binary variables $Y_n \in \{0, 1\}$.

We seek to formalize the definition of a probabilistic morphological filter \mathcal{F}^* applied on the probabilistic segmentation image $\mathcal{Y} = \{y_n\} \in [0, 1]^N$ and generate a new probabilistic image $\mathcal{Z} = \{z_n\} \in [0, 1]^N = \mathcal{F}^*(\mathcal{Y})$. This filter should generalize the given binary filter $Z = F(Y)$ such that both give the same result when the input probabilistic image is binary $Z = F(Y) = \mathcal{F}^*(Y)$. More precisely, we aim to apply the deterministic morphological operator $F()$ on a binary image Y only known through its posterior probability $\mathcal{Y} = p(Y|X)$. Therefore we estimate the posterior $z_n = p(Z_n = 1|X)$ of the final segmentation Z knowing that it results from the application of the morphological operation $Z = F(Y)$.

Lemma 1 *The posterior probability z_n can be obtained through the law of total probability as the expectation of filtered binary segmentation $F(a)$:*

$$\begin{aligned} z_n = p(Z_n = 1|X) &= \sum_{Y_1=0}^1 \dots \sum_{Y_N=0}^1 p(Z_n = 1|Y) p(Y|X) \\ &= \sum_{a \in \{0,1\}^N} F(a) p(a|X) = \mathbb{E}_{a \sim p(Y|X)} F(a) \end{aligned}$$

Lemma 1, as represented in Fig.2, defines implicitly the relationship $z_n = \mathcal{F}^*(\mathcal{Y})$ of the soft morphological filter.

C. Soft operators using Multi-linear polynomials

It is furthermore required to make the filter differentiable, i.e. to estimate the derivatives $\frac{\partial \mathcal{F}^*(\mathcal{Y})}{\partial y_m}$. To provide a closed-form expression of a soft filter defined in Lemma 1, we propose to adopt a polynomial representation of the Boolean function $F(a)$, $a \in \{0, 1\}^k$. Indeed we can associate with any Boolean function $F(a)$ a multilinear polynomial $\mathcal{F}^*(x)$, $x = (x_1, \dots, x_k)^T \in \mathbb{R}^k$ defined as:

$$\mathcal{F}^*(x) = \sum_{a \in \{0,1\}^k} F(a) \prod_{i=1}^k x_i^{a_i} (1 - x_i)^{1-a_i} \quad (1)$$

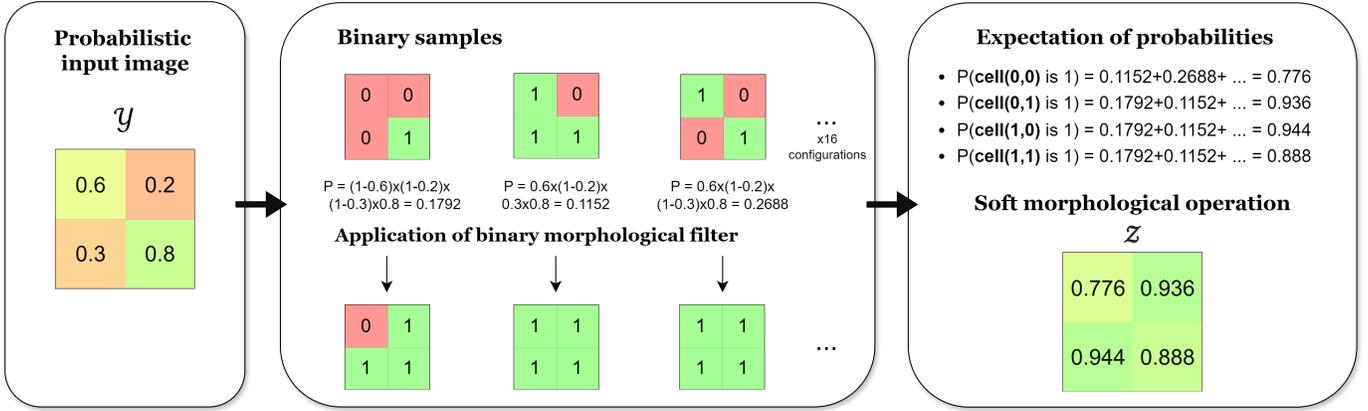


Fig. 2. Given a 2x2 input grid with probabilistic values, the probability of each possible binary configuration is calculated. By applying a specific morphological operator to each binary configuration, we compute the expectation over all configurations to estimate the result of the soft operator on the probabilistic input.

It is easy to see that $F(a) = \mathcal{F}^*(a)$, i.e. that the polynomial $\mathcal{F}^*(x)$ coincides by construction with the Boolean function on the hypercube $\{0, 1\}^k$. Each monomial $F(a) \prod_{i:a_i=1} x_i \prod_{i:a_i=0} (1-x_i)$ is equal to 0 if $x \neq a$ and equal to $F(a)$ otherwise. Besides, it is of degree k and linear with each variable x_j , making the polynomial multilinear. This property leads to the following result:

Theorem *The expectation of a Boolean function $F(a) \in \{0, 1\}$, $a \in \{0, 1\}^k$ over a set of k independent variables with $a \sim \text{Bernoulli}(p)$, $p \in [0, 1]^k$ is $\mathcal{F}^*(p)$*

Proof *It is easy to show that $\mathbb{E}_{a_i \sim \text{Bernoulli}(p_i)}(\alpha + \beta a_i) = \alpha + \beta p_i$ using the linearity of expectation. Thus, we have :*

$$\begin{aligned} \mathbb{E}_{a_1}(F(a)) & \prod_{j=1}^k a_j^{a_j} (1-a_j)^{1-a_j} \\ & = F(a) p_1^{a_1} (1-p_1)^{1-a_1} \prod_{i=2}^k a_i^{a_i} (1-a_i)^{1-a_i}. \end{aligned}$$

By taking the expectation over each variable a_i , we get :

$$\mathbb{E}_a(F(a)) = \sum_{a \in \{0,1\}^k} F(a) \prod_{i=1}^k p_i^{a_i} (1-p_i)^{1-a_i} = \mathcal{F}^*(p). \quad \square$$

Therefore, assuming that the marginal posteriors $y_n = p(Y_n|X)$ are independently distributed (which is the case when dealing with the output of segmentation neural networks or mean field approximations), we define the soft morphological filter associated with the binary filter $F()$ as the polynomial value $\mathcal{F}^*(y_{\mathcal{N}(1,n)}, \dots, y_{\mathcal{N}(k,n)}) \in [0, 1]$.

With our example binary operator F_{Ex} , we get the multilinear polynomial representation of the soft operator :

$$\begin{aligned} \mathcal{F}_{\text{Ex}}^*(y_1, y_2, y_3) & = \sum_{a \in \{0,1\}^3, F(a) \neq 0} \prod_{i=1}^3 y_i^{a_i} (1-y_i)^{1-a_i} \\ & = y_1 y_2 (1-y_3) + (1-y_1) y_2 y_3 \\ & = y_1 y_2 + y_2 y_3 - 2 y_1 y_2 y_3 \end{aligned}$$

This function defined over the 3D cube can be seen as the trilinear interpolation of the binary truth table as illustrated in Fig. 3a. More generally, the probabilistic morphological

operator defined in Theorem 1, can be interpreted as the multilinear interpolation of the binary operator F over the hypercube. Multilinear interpolation is among the most basic interpolation methods and is symmetric with respect to all variables, which makes it a natural extension of the binary operator F to the probabilistic context.

D. Limitations of Multi-linear polynomial representation

The construction of the multilinear polynomial $\mathcal{F}^*(\cdot)$ (Eq.(1)) requires the summation over non-zero elements $F(a)$ of the truth table of size 2^k . For non-trivial truth tables, writing such polynomials requires the use of symbolic computation software such as SymPy or Maple. But the complexity of such polynomials grows exponentially with the number k of variables. For example, if $k = 26+1$ in a 3D images, there are $2^{27} = 134\,217\,728$ possible input configurations. In practice, when $k > 10$ for non-trivial filters, the number of monomials often becomes prohibitively large.

E. Quasi-probabilistic operators using fuzzy logic

The computation of \mathcal{F}^* is based on the exhaustive list of positive binary configurations which can become intractable to produce. We are looking for alternative computation methods that have the same complexity as the binary filter $F(a)$, $a \in \{0, 1\}^k$. To this end, we notice that the probabilistic version of the AND operator is the product of the probabilities $\text{AND}(X_1, X_2) \rightarrow x_1 x_2$ while the OR operator is transformed into $\text{OR}(X_1, X_2) \rightarrow x_1 + x_2 - x_1 x_2 = 1 - (1-x_1)(1-x_2)$ and the NOT operator as $\text{NOT}(X_1) \rightarrow 1 - x_1$. For any Boolean operator $F()$ represented by a proposition formula, involving the AND, OR and NOT operators, we propose to create a soft quasi-probabilistic operator $\mathcal{F}^\bullet()$ by substituting the logical operators AND, OR, and NOT by their probabilistic versions. This soft operator is a polynomial expression which can be of a degree greater than k , and is not necessarily multilinear. However, its computation complexity is the same as the one of the propositional form of the operator since we have substituted a simple logical expression with another simple

polynomial one. The *quasi-probabilistic* polynomial corresponds to a compact and factorized form whereas the multilinear polynomial corresponds to its polynomial expansion combined with the application of the idempotence rule $x^i = x$, $\forall i > 0$ on all probability variables. The quasi-probabilistic polynomial can be interpreted as an alternative non-linear interpolation method of the values over the hypercube whereas the probabilistic polynomial is a (multi) linear one.

While AND, OR and NOT operators were substituted by their probabilistic versions, one could think of other ways to replace the 3 logical operators with algebraic expressions of continuous values in the range $[0, 1]$. This has been the focus of the Fuzzy logic [51] and Fuzzy set theories. They introduce Triangular Norms (T-norms), and Triangular Conorms (S-norms) as substitutes for the AND and OR operators with the complementation operator $1 - x$ as the substitute for the NOT operator. T and S-norms are defined on the unit square $[0, 1] \times [0, 1]$ with values on the unique segment $[0, 1]$, and follow the commutativity, monotonicity, associativity and element identity properties. Besides, an S-norm $S()$ is dual of a T-norm $T()$ under the action on the complementation operator, $S(x_1, x_2) = 1 - T(1 - x_1, 1 - x_2)$, which can be seen as the generalization of the De Morgan's rules.

There exist many different T/S-norms proposed in the literature, and they can be ordered pointwise as follows:

$$T_1 \leq T_2 \quad \text{if} \quad T_1(a, b) \leq T_2(a, b) \quad \text{for all} \quad a, b \in [0, 1].$$

The expressions of the main T and S-norms are listed in increasing order in Table I and their graphs on the unit square are displayed in Fig. 4. The *product* T-norm corresponds to the probabilistic AND operators introduced previously and it is easy to see on Fig. 4 that the Drastic T-norm is the smallest whereas the Minmax logic is the largest. A number of those T-norms have additive generators [52], which means that there exists a function $f : [0, 1] \rightarrow \mathbb{R}^+$ such that $T(x, y) = f^{-1}(f(x) + f(y))$. This is the case for the product logic (with $f(x) = -\log(x)$) and this allows to easily factorize multiple applications of the T and S-norms. For the product rule, we have for instance : $\text{AND}(Y_0, Y_1, \dots, Y_l) \equiv y_0 y_1 \dots y_l$, $\text{OR}(y_0, y_1, \dots, y_l) \equiv 1 - (1 - y_0)(1 - y_1) \dots (1 - y_l)$

With each T-norm, we can substitute the AND, OR, and NOT operations of the binary function $F()$ to obtain an approximation $\mathcal{F}^\bullet()$ of the associated probabilistic operators \mathcal{F}^* . Both functions coincide on the vertices of the unit hypercube, but they correspond to different interpolation functions inside the hypercube. Thus, we call *SoftMorph* the family of soft operators derived from a binary morphological operator $F(Y_{N(1,n)}, \dots, Y_{N(k,n)})$ which can be either *probabilistic operators* as multilinear polynomials $\mathcal{F}^*(y_{N(1,n)}, \dots, y_{N(k,n)})$ or *quasi-probabilistic operators* $\mathcal{F}^\bullet(y_{N(1,n)}, \dots, y_{N(k,n)})$ as derived from a T / S-norm. The Einstein and product logic generate smooth and differentiable expressions whereas the other T-norms are only piecewise differentiable, in particular due to the min and max functions. As an example, the *SoftMorph* operator associated with the binary function F_{Ex} using product logic writes as $\mathcal{F}^\bullet_{\text{Ex}}(y_1, y_2, y_3) = y_2(1 - (1 - y_1(1 - y_3))(1 - y_3(1 - y_1)))$. This is a factorized polynomial which differs from its probabilistic version and its expansion involves

TABLE I
T-NORMS AND S-NORMS FORMULA FOR DIFFERENT FUZZY LOGICS

Logic	Operation	Formula
Boolean	AND	$a \wedge b$
	OR	$a \vee b$
Drastic	T-norm	$\min(a, b)$ if $\max(a, b) = 1$, else = 0
	S-norm	$\max(a, b)$ if $\min(a, b) = 0$, else = 1
Bounded	T-norm	$\max(0, a + b - 1)$
	S-norm	$\min(1, a + b)$
Einstein	T-norm	$(ab)/(2 - (a + b - ab))$
	S-norm	$(a + b)/(1 + ab)$
Product	T-norm	$a * b$
	S-norm	$a + b - 2ab$
Hamacher	T-norm	0 if $a = b = 0$, else $(ab)/(a + b - ab)$
	S-norm	1 if $a = b = 1$, else $(a + b - 2ab)/(1 - ab)$
Min-Max	T-norm	$\min(a, b)$
	S-norm	$\max(a, b)$

monomials of degree 5 (instead of 3 for the multilinear case) $y_1 y_2 + y_2 y_3 - 3y_1 y_2 y_3 + y_1^2 y_2 y_3 + y_1 y_2 y_3^2 - y_1^2 y_2 y_3^2$. Both polynomials have the same values on the unit hypercube, but differ elsewhere as seen in Fig. 3. The largest difference between the probabilistic and quasi-probabilistic functions on the unit cube is only around 0.06, showing that $\mathcal{F}^\bullet_{\text{Ex}}$ is a good approximation of $\mathcal{F}^*_{\text{Ex}}$. In this simple case, there is no benefit to use $\mathcal{F}^\bullet_{\text{Ex}}$ instead of $\mathcal{F}^*_{\text{Ex}}$, but for more complex functions, one must resort to quasi-probabilistic *SoftMorph* functions.

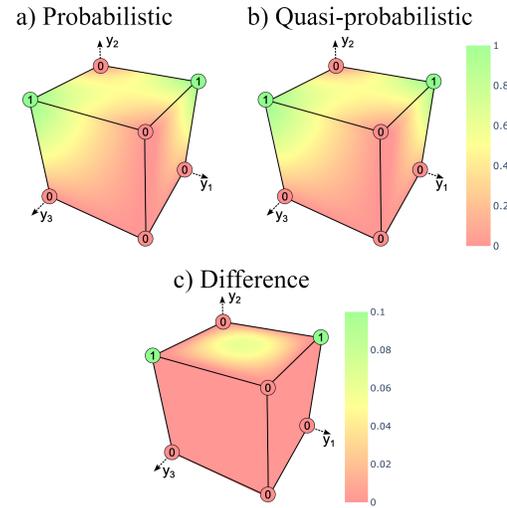


Fig. 3. Field plots of the probabilistic (a) and product-based quasi-probabilistic (b) morphological filters of F_{Ex} over the hypercube. (c) Field plot of the difference between the two filters. The maximum difference is reached at the center of the faces or the center of the volume.

F. Relation to prior work

The dilation and erosion operators have been defined on any grayscale images using the notion of supremum and infimum [2]. On finite sets, dilation and erosion are obtained by considering the maximum, and minimum values within a SE. They have been further generalized as fuzzy morphological operators [4] with fuzzy (i.e. within the $[0, 1]$ range) SE and based on fuzzy logic (T / S-norms). However, when restricted

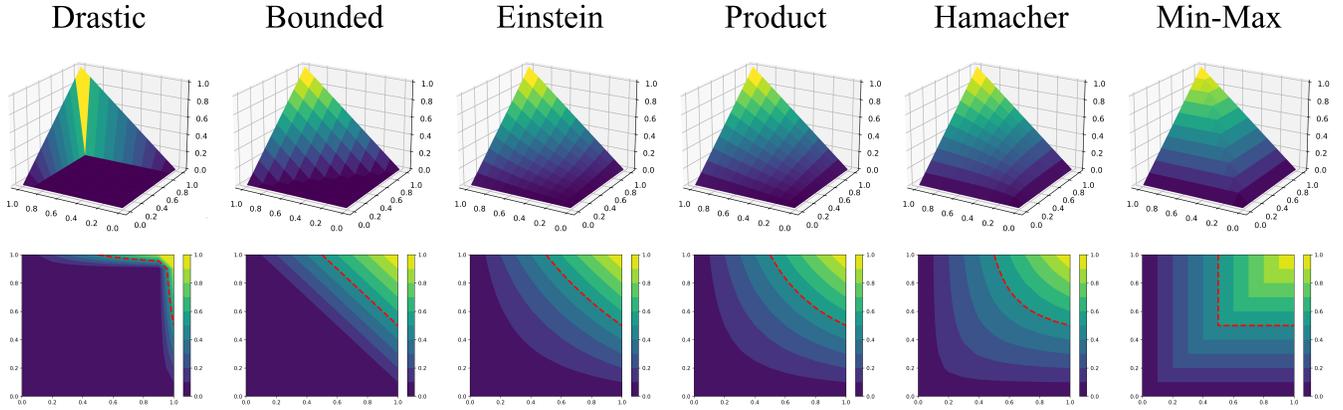


Fig. 4. Graphs of the drastic, bounded, Einstein, product, Hamacher and minmax T-norms on the $[0, 1] \times [0, 1]$ unit square. The red dashed contours correspond to the isocontour at 0.5.

to binary SEs, the fuzzy erosion and dilation operators are also equivalent to the ones defined in mathematical morphology (independently of the selected fuzzy logic), i.e. taking the maximum/minimum values in the neighborhood of a pixel.

It is important to note that those dilation and erosion operators are specific cases of the quasi-probabilistic morphological operators when adopting the Min-Max logic. Therefore, the *SoftMorph* operators are novel *soft* formulations of binary morphological operators that supersede the existing erosion and dilation operators. Those multi-linear and product-based operators are smooth and fully differentiable whereas most other quasi-probabilistic ones require to adopt smooth min-max approximations to be differentiable. Furthermore, probabilistic morphological operators and their approximations are "averaged morphological operators" as the expectation of a binary operator. This implies that the *SoftMorph* closing and opening operations are not idempotent (producing the same output irrespective of the number of times it is applied) unlike the classical closing and opening using the Min-Max logic.

G. Morphological Operators Of Interest

In this paper, we generate the *SoftMorph* operators on 5 main morphological operations: erosion, dilation, closing, opening, and skeletonization. Depending on the SE and the image dimension, the erosion and dilation F_{Dil} Boolean expressions are written as $F_{Ero} = \bigwedge_{i=1}^k Y_k$ and $F_{Dil} = \bigvee_{i=1}^k Y_k$. The closing operation is a dilation followed by an erosion whereas the opening is an erosion followed by a dilation.

Skeletonization is a more complex morphological operation for which many algorithms have been proposed [53]. This operation involves repeated thinning operations, applied iteratively until the final skeleton is obtained. The morphological thinning on 2D images proposed by Wagner *et al.* [54] is divided into 4 sub-iterations, one for each North, East, South, and West directions, consists of the following Boolean function for the North direction defined on the $k = 8 + 1$

neighborhood as :

$$F1 = \neg Y_6 \wedge \neg Y_7 \wedge (Y_2 \oplus Y_5 \oplus \neg Y_1) \wedge (Y_5 \oplus Y_8) \wedge (Y_5 \oplus Y_4)$$

$$F2 = (Y_6 \oplus Y_8) \wedge (Y_2 \oplus Y_8) \wedge (Y_1 \oplus \neg Y_8) \wedge \neg Y_4 \wedge \neg Y_7$$

$$F3 = Y_2 \wedge Y_1 \wedge \neg Y_7$$

$$F4 = \neg Y_6 \wedge Y_1 \wedge Y_4$$

$$F5 = Y_2 \wedge Y_1 \wedge Y_4$$

$$F6 = \neg Y_6 \wedge \neg Y_2 \wedge \neg Y_5 \wedge Y_4 \wedge Y_7$$

$$F = \neg Y_3 \wedge (F1 \vee F2 \vee F3 \vee F4 \vee F5 \vee F6).$$

The Boolean formula must be rotated for the 3 other directions.

Similarly, we have translated the 3D thinning algorithm proposed by Palàgy [55] as a Boolean function defined on the $k = 26 + 1$ neighborhood which is provided in the supplementary material. In binary skeletonization, the algorithm ends when the binary structure is stable upon the application of each thinning sub-iteration. However, the *SoftMorph* are in general not idempotent (except with the Min-Max logic) and therefore a stopping criterion must be defined. We propose to stop the soft thinning process when the change between two thinning operations is less than 2% for all pixels in the initial foreground object.

IV. EXPERIMENTS

A. Experimental setup

Datasets: Experiments are conducted on six 2D and two 3D datasets. In 2D, The DRIVE dataset [56] corresponds to retinal blood vessels. The Massachusetts Road [57] dataset comprises satellite images of road networks. Labeled images are extracted from the Open Images Dataset V7 [58] for classes Sea turtle, Starfish and Croissant. The Butterfly dataset [59] is composed of butterfly images (masks have been corrected and available in the repository). In 3D, We use the Vessap dataset featuring synthetic brain vessels [60] and the Liver task from the Medical Segmentation Decathlon [61].

Evaluation metrics: To evaluate the experimental performances on the final segmentations, we use the Dice similarity coefficient to measure the overlap with the ground truth. For datasets containing tubular structures (Drive, Massachusetts

road, and Croissant datasets), we additionally use the cIDice metric [25] to assess the topological preservation. Given the importance of topological accuracy in morphological operations, we also compute the mean absolute error of topological invariants: the Betti numbers β_0 (the number of connected components), β_1 (the number of holes), and Euler’s number (the difference between the betti numbers).

Implementation: Experiments are implemented with Python 3.11.4, Pytorch 2.0.1 and 3 Nvidia A40 PCIe GPUs.

B. Validation experiments

1) *Validation of SoftMorph on binary images:* We evaluate the reliability of our designed probabilistic morphological operators on binary images, as summarized in Table II. The primary objective is to ensure that our probabilistic operators replicate accurately the binary filters. The reference corresponds to the non-differentiable morphological operations for erosion, dilation and skeletonization, from the widely used scikit-image package [62]. We do not assess the opening and closing operations as they correspond to iterative applications of erosion and dilation. Additionally, we compare the performance of our operators against other existing differentiable morphological operators on binary images. For the erosion and dilation operations, we test the max and min-pooling layers with a kernel of 3x3. For skeletonization, we compare our method with the *soft-skeleton* approach from the cIDice paper [25] (corresponding to a combination of min and max-pooling layers) and the method by Menten et al. [34] (corresponding to the convolutional layers with specific kernels to detect simple points). In 2D, we also assess a neural network model trained for skeletonization from [36]. Each method is evaluated on 15 randomly selected images in 2D from the DRIVE dataset and in 3D with the VesSap dataset.

Our method precisely replicates the erosion and dilation operations of the reference, whereas the min and max-pooling layers demonstrate significant topological errors when the SE of the operations is defined on the $k = 4 + 1$ neighborhood. Especially, the soft min-pooling approach exhibits an average β_0 absolute error of 214.10. These discrepancies do not arise when the reference SE is set to $k = 8 + 1$. This is because these pooling layers are defined with a 3x3 kernel that corresponds to the $k = 8 + 1$ SE. Because they are limited to square SEs, they can not replicate the $k = 4 + 1$ SE as defined in 1a.

The skeletons produced by our method and from Menten et al. are both topologically accurate. However, the soft-skeleton from cIDice and the trained model show high topological errors. It is important to note that the Dice score is not an optimal metric to assess skeletonization performances compared with the reference. Multiple valid skeletons can be derived from the same initial volume, making the definition of the centerline ambiguous, for example when the width of the object is set on an even number of pixels. Therefore, the Dice scores from our method and Menten et al. indicate similarities with the reference with values of 0.65 in 2D and 0.71 in 3D.

Although our probabilistic operators are considerably slower than other methods, this trade-off ensures topological correctness. These results demonstrate that our probabilistic morphological operators are accurately designed and can replicate the

exact traditional binary morphological filters, whereas most other differentiable methods lack topological reliability or the adaptability to scale to various morphological operations.

2) *Backpropagation capability of SoftMorph:* We quantify the backpropagation capabilities of the probabilistic and quasi-probabilistic representations to evaluate their impact on gradient computation for optimization within neural networks. This experiment is similar to the one described in [34]. We initialize a tensor with random values that we pass through a morphological operator. The soft-Dice is used to compute the loss between the operation’s output and a ground truth image that is also passed through the same operator. The propagation of gradients enables the adjustment of the tensor’s values until the operation’s output converges towards that of the ground truth image. The experiment is performed for dilation, erosion and skeletonization operations, converging with a learning rate set to 1 over 20 epochs. The operators compared are the probabilistic multi-linear polynomial operator, and the family of proposed quasi-probabilistic operators based on fuzzy logic. We record the loss values at each epoch, as shown on Fig 5 to compare the convergence speed and performances of the different operators.

For dilation, the drastic and bounded logic-based operators fail to facilitate gradient backpropagation across all epochs. The other operators converge around the 15th epoch, with the product and Einstein operators reaching the best final loss values of 0.01 and 0 respectively. For erosion, we observe the same trend for the drastic and bounded logic operators, whereas all other operators converge to the exact ground truth operation output values after around 8 epochs only. In the skeletonization operation, the drastic and bounded logic operators achieve the best convergence after 10 epochs, fully learning the operation’s output. The Hamacher and product operators follow, reaching a loss of approximately 0.10, with the Einstein and multi-linear operators trailing at 0.20, 0.24. The Min-Max operator obtains a Dice loss of 0.47 after 20 epochs, although the Einstein, multi-linear and Min-max do not appear to have fully converged within this time frame.

Overall, the product-based operator emerges as the most stable operator representation, allowing an efficient convergence across all operations. The Hamacher and Einstein operators also perform effectively but are less consistent than the product operator. Finally, the drastic and bounded logic exhibit unique behavior: their very restricted formulations seem to hinder the backpropagation in simpler operations like erosion and dilation, yet allowing effective gradient backpropagation in the more complex skeletonization operation, compensating for their inherent sparsity and rigidity. Therefore, most of these operators effectively support gradient backpropagation for the optimization of CNNs. The product-based logic is recommended for such use due to its consistent performance.

C. Applications

1) *Final morphological layer for segmentation:* We test the integration of the probabilistic morphological operators into a CNN’s architecture, specifically in the U-net model [63] for semantic segmentation across seven datasets. In our implementation, we add a probabilistic morphological operator as

TABLE II
 VALIDATION OF DIFFERENTIABLE MORPHOLOGICAL OPERATORS ON 2D AND 3D BINARY IMAGES COMPARED TO NON-DIFFERENTIABLE REFERENCE OPERATORS. TIME RATIO CORRESPONDS TO THE MEAN RATIO OF RUNNING TIME OF EACH OPERATOR COMPARED TO THE REFERENCE.

Dataset	Operation	Method	$\beta_0 \downarrow$	$\beta_1 \downarrow$	Euler \downarrow	Dice \uparrow	Time ratio \downarrow
DRIVE 2D dataset	Dilation ($k = 4 + 1$)	Ours	0	0	0	1	13.33 \pm 17.02
		Soft maxpooling	0.10 \pm 0.31	6.65 \pm 6.64	6.75 \pm 6.60	0.95 \pm 0.002	0.33 \pm 0.91
	Erosion ($k = 4 + 1$)	Ours	0	0	0	1	9.58 \pm 7.02
		Soft minpooling	214.10 \pm 98.27	7.60 \pm 3.69	211.80 \pm 92.64	0.85 \pm 0.02	1.09 \pm 3.33
	Skeleton	Ours	0	0	0	0.65 \pm 0.02	115.31 \pm 28.09
		Neural Network	206.15 \pm 19.66	22.30 \pm 14.27	226.85 \pm 30.40	0.77 \pm 0.02	35.84 \pm 48.48
Soft-skeleton Menten et al.		1414.20 \pm 191.19	66.50 \pm 15.63	1480.70 \pm 203.26	0.65 \pm 0.01	1.86 \pm 6.98	
VesSap 3D dataset	Skeleton	Ours	0	0	0	0.72 \pm 0.002	121.26 \pm 4.93
		Soft-skeleton	8362.20 \pm 360.33	8.80 \pm 4.30	8371.00 \pm 358.99	0.64 \pm 0.003	0.90 \pm 0.048
		Menten et al.	0	0	0	0.65 \pm 0.01	89.94 \pm 75.27
		Ours	0	0	0	0.71 \pm 0.001	12.37 \pm 0.61

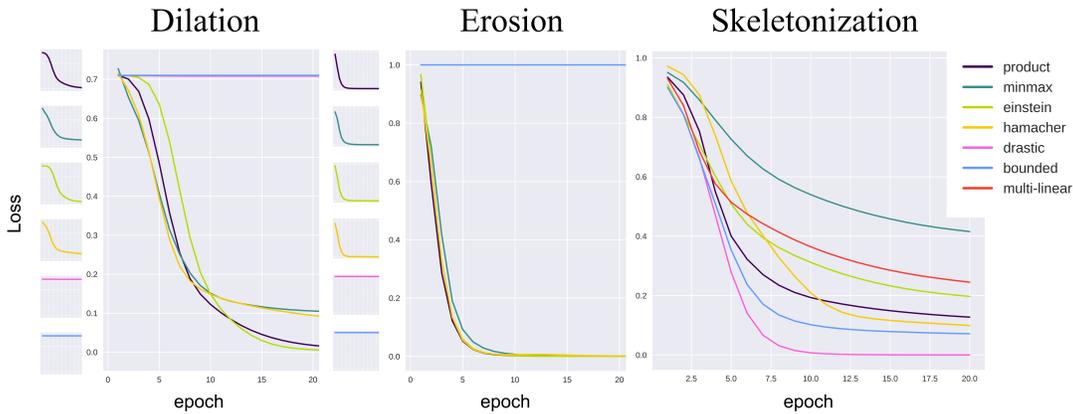


Fig. 5. Results of the backpropagation capability between the SoftMorph probabilistic and quasi-probabilistic operators for dilation, erosion and skeletonization.

the final layer of the network, following a sigmoid activation function. The operator is applied twice to enhance its effect on the network. For the 3D liver dataset, we use the nn-UNet model [64] that is specialized for medical segmentation. The training is conducted with the soft-dice loss function, a batch size of 16, and a learning rate of $1e-4$ over 1000 epochs. The datasets are split into 80% training and 20% testing sets. The operators tested include the erosion, dilation, opening and closing filters as the final network layers. We compare these configurations against the baseline U-Net model without any morphological operator. The averaged results from the testing set are presented in Table III. The cIDice metric is only measured for the datasets containing tubular structures.

Our results indicate a tendency for topological improvements across all datasets with the application of morphological filters. The β_0 absolute error is significantly lower with the inclusion of the morphological operators while maintaining a high Dice score on the DRIVE dataset. The cIDice is also improved in the tubular structures. However, no definitive pattern emerges to predict which operator can yield the best performances based on the data characteristics. We conclude that integrating morphological operators as final layers of a U-net architecture for segmentation tasks can considerably improve the topological accuracy. Currently, determining the most beneficial morphological operation for a specific dataset is based only on a trial-and-error approach.

2) *cIDice loss function*: We test the integration of the probabilistic morphological operators into a loss function. The

loss corresponds to the cIDice loss [25] designed to improve the topological preservation of tubular structure segmentation. This loss is combined with the SoftDice loss, weighted by a parameter α . We compare the segmentation performance on the 2D DRIVE dataset with the SoftDice alone, and the combination of cIDice and SoftDice with $\alpha = 0.5$ and $\alpha = 0.7$. In cIDice, it is required to extract the skeleton from both the prediction and the ground truth. To do so, we test the soft-skeleton from cIDice [25], the method from Menten et al., and our product-based quasi-probabilistic skeleton operator. We also evaluate the effects of dilation and erosion as final layers of the network in conjunction with the cIDice loss computed with our product operator. We use a U-Net model trained with a batch size of 16 and a learning rate of $1e-4$ over 500 epochs.

Our results in Table IV show that topological performances are improved with the cIDice loss function compared to BCE and SoftDice alone. The application of cIDice with our skeletonization method yields the best topological performances in terms of β_0 and Euler numbers, while also maintaining or even slightly improving the Dice and cIDice scores compared to other methods. The addition of the final morphological layer further enhances these results by significantly reducing the number of topological errors.

In conclusion, our skeletonization method demonstrates the best overall performance compared to the other skeletons. The soft-skeleton method is prone to topological errors as shown in the validation experiment in Table II, which likely contribute

TABLE III

DOUBLE MORPHOLOGICAL OPERATORS AS LAST LAYERS OF THE U-NET SEGMENTATION NETWORK. BOLD VALUES CORRESPOND TO IMPROVED PERFORMANCES COMPARED TO THE BASELINE AND (*) DENOTES STATISTICALLY SIGNIFICANT IMPROVEMENTS USING THE WILCOXON RANK TEST.

DATASET	Final layer	$\beta_0 \downarrow$	$\beta_1 \downarrow$	Euler \downarrow	Dice \uparrow	cDice \uparrow
	Normal	37.05 \pm 9.76	30.90 \pm 14.33	67.95 \pm 15.81	0.83 \pm 0.02	0.87 \pm 0.03
	Opening	27.40* \pm 8.44	33.30 \pm 13.45	60.70 \pm 14.05	0.81 \pm 0.01	0.88 \pm 0.03
	Closing	28.10* \pm 6.35	36.15 \pm 15.16	64.25 \pm 17.36	0.82 \pm 0.02	0.87 \pm 0.03
	Erosion	29.70* \pm 8.55	36.05 \pm 14.34	65.75 \pm 14.81	0.82 \pm 0.01	0.86 \pm 0.02
	Dilation	26.20* \pm 8.00	28.00 \pm 15.94	54.00* \pm 18.71	0.81 \pm 0.02	0.88 \pm 0.03
	Normal	3.30 \pm 2.52	3.55 \pm 5.42	4.55 \pm 5.56	0.57 \pm 0.32	0.56 \pm 0.32
	Opening	3.30 \pm 3.06	5.15 \pm 10.82	6.55 \pm 9.48	0.60 \pm 0.33	0.59 \pm 0.33
	Closing	1.80* \pm 2.19	3.40 \pm 6.06	3.90 \pm 7.30	0.60 \pm 0.34	0.58 \pm 0.34
	Erosion	7.15 \pm 4.82	4.45 \pm 6.16	5.60 \pm 5.59	0.57 \pm 0.29	0.53 \pm 0.27
	Dilation	2.85 \pm 2.54	4.25 \pm 5.38	4.80 \pm 5.45	0.58 \pm 0.32	0.57 \pm 0.32
	Normal	31.00 \pm 29.18	30.71 \pm 31.66	61.47 \pm 49.27	0.57 \pm 0.12	0.64 \pm 0.13
	Opening	26.71 \pm 19.58	29.65 \pm 31.84	56.35 \pm 43.30	0.54 \pm 0.10	0.67 \pm 0.14
	Closing	23.29 \pm 17.55	29.82 \pm 31.43	53.00 \pm 41.32	0.58 \pm 0.13	0.65 \pm 0.15
	Erosion	20.41 \pm 18.04	31.24 \pm 32.17	50.94 \pm 42.01	0.56 \pm 0.11	0.64 \pm 0.13
	Dilation	19.29 \pm 15.70	20.71 \pm 26.85	37.29* \pm 29.24	0.54 \pm 0.11	0.69 \pm 0.15
	Normal	1.45 \pm 1.50	1.25 \pm 1.68	1.80 \pm 1.47	0.87 \pm 0.25	-
	Opening	1.20 \pm 1.32	0.70 \pm 1.22	1.40 \pm 1.64	0.86 \pm 0.27	-
	Closing	1.45 \pm 1.90	1.60 \pm 2.54	2.45 \pm 3.55	0.86 \pm 0.23	-
	Erosion	1.90 \pm 1.80	2.55 \pm 3.39	2.75 \pm 2.65	0.88 \pm 0.17	-
	Dilation	2.55 \pm 2.82	3.30 \pm 4.66	2.65 \pm 2.89	0.86 \pm 0.20	-
	Normal	6.00 \pm 4.83	2.60 \pm 2.37	4.80 \pm 3.49	0.54 \pm 0.28	-
	Opening	5.65 \pm 4.49	3.20 \pm 3.02	4.55 \pm 3.65	0.56 \pm 0.28	-
	Closing	3.20 \pm 2.38	1.90 \pm 2.59	3.70 \pm 3.51	0.41 \pm 0.30	-
	Erosion	9.15 \pm 9.30	2.70 \pm 3.21	8.25 \pm 6.88	0.57 \pm 0.26	-
	Dilation	5.70 \pm 5.02	3.90 \pm 3.73	4.50 \pm 2.56	0.54 \pm 0.28	-
	Normal	4.1 \pm 6.89	2.9 \pm 3.70	5.3 \pm 6.21	0.60 \pm 0.31	-
	Opening	5.35 \pm 4.46	2.6 \pm 3.39	5.45 \pm 5.16	0.55 \pm 0.32	-
	Closing	4.75 \pm 6.09	2.95 \pm 3.44	5.40 \pm 4.51	0.54 \pm 0.30	-
	Erosion	5.00 \pm 4.93	2.15 \pm 2.92	5.65 \pm 5.48	0.55 \pm 0.35	-
	Dilation	4.30 \pm 4.79	3.25 \pm 4.73	4.05 \pm 4.61	0.53 \pm 0.37	-
	Normal	2.81 \pm 3.84	2.48 \pm 3.57	3.44 \pm 4.08	0.94 \pm 0.07	-
	Opening	3.00 \pm 3.16	2.07 \pm 3.21	3.67 \pm 4.04	0.92 \pm 0.07	-
	Closing	1.85 \pm 2.14	1.33 \pm 2.32	2.15 \pm 2.32	0.92 \pm 0.07	-
	Erosion	1.63 \pm 1.88	1.63 \pm 3.19	2.44 \pm 3.32	0.92 \pm 0.07	-
	Dilation	1.19 \pm 1.55	2.04 \pm 3.24	2.33 \pm 3.53	0.92 \pm 0.07	-

to its lower performances. While the skeleton from Menten et al. demonstrates topological accuracy on binary images, its lower performances may be due to the reparametrization trick used in this method. This trick involves a sensitive set of parameters which can be hard to tune for the specific characteristics of the task and dataset. Here it is applied with the default parameters proposed by Menten et al. In this experiment, training and inference times were comparable across skeletonization methods, despite our operator being slower in standalone validation, suggesting its computational overhead is minimal within CNNs.

Therefore, these findings suggest that SoftMorph operators can effectively enhance topological performance in segmentation tasks by extracting accurate morphological features that can be integrated into loss functions or used as the final layer of a neural network. It can achieve improved results without requiring any parameter tuning.

V. CONCLUSION

We presented SoftMorph, a family of differentiable probabilistic and quasi-probabilistic morphological operators for deep learning frameworks. SoftMorph successfully bridges the gap to translate any morphological operation defined on any SE in its soft counterpart while maintaining the computational complexity of the original binary operator. We have

demonstrated that probabilistic filters can be defined as the expectation of the binary filter and represented as a multilinear polynomial. We have also shown that the factorized form of the original binary filter can be approximated as a quasi-probabilistic filter using fuzzy logic. Besides, these operators replicate the exact output of binary operators on binary images while enabling gradient-based optimization and handling probabilistic maps. Integrating some basic morphological operations can improve the topological performances of segmentation networks when inserted as the final layer or within loss functions. Possible improvements include the optimization of the 2D and 3D skeletonization algorithms, the definition of morphological operations with non-flat (fuzzy) SEs, and learning new morphological operations and their associated optimal SE. Overall, this work opens avenues for defining new task-specific morphological operations. Their application could be extended to other image analysis applications, neural network architectures, and new loss functions requiring the extraction of morphological features.

REFERENCES

- [1] J. Serra, P. Soille, and M. A. Viergever, *Mathematical Morphology and Its Applications to Image Processing*, ser. Computational Imaging and Vision. Dordrecht: Springer Netherlands, 1994, vol. 2.
- [2] F. Y. Shih, *Image Processing and Mathematical Morphology: Fundamentals and Applications*. Boca Raton: CRC Press, Jan. 2017.

TABLE IV

2D U-NET SEGMENTATION PERFORMANCE ON DRIVE DATASET. (*) DENOTES STATISTICALLY SIGNIFICANT DIFFERENCES WITH SOFT-SKELETON AND MENTEN ET AL. METHODS HAVING THE SAME α VALUE USING THE WILCOXON RANK TEST.

LOSS \ METRICS			$\beta_0 \downarrow$	$\beta_1 \downarrow$	Euler \downarrow	Dice \uparrow	cDice \uparrow
SoftDice (baseline)			45.35 \pm 6.58	35.15 \pm 16.32	80.50 \pm 17.99	0.82 \pm 0.02	0.86 \pm 0.03
cDice	$\alpha=0.7$	Soft-Skeleton	17.90 \pm 4.63	27.2 \pm 14.54	44.60 \pm 14.92	0.79 \pm 0.01	0.88 \pm 0.02
		Menten et al.	18.15 \pm 5.14	38.40 \pm 14.20	56.55 \pm 14.37	0.76 \pm 0.04	0.85 \pm 0.03
		Ours	14.70 \pm 4.37	27.95 \pm 15.28	42.65 \pm 14.68	0.81 \pm 0.01	0.89 \pm 0.02
	$\alpha=0.5$	Soft-Skeleton	24.70 \pm 4.69	32.05 \pm 14.57	56.75 \pm 15.41	0.82 \pm 0.02	0.88 \pm 0.03
		Menten et al.	25.20 \pm 6.34	34.35 \pm 15.10	59.55 \pm 15.75	0.78 \pm 0.04	0.87 \pm 0.03
		Ours	16.50* \pm 5.72	32.90 \pm 15.04	44.40 \pm 16.83	0.83 \pm 0.02	0.89* \pm 0.03
cDice + final layer	Ours $\alpha=0.7$	Dilation	11.30* \pm 4.74	30.35 \pm 15.16	41.65 \pm 16.94	0.80 \pm 0.02	0.90* \pm 0.03
		Erosion	10.15* \pm 4.33	32.75 \pm 14.49	42.90 \pm 15.87	0.82* \pm 0.01	0.89 \pm 0.03
	Ours $\alpha=0.5$	Dilation	11.10* \pm 4.71	30.55 \pm 15.67	41.55* \pm 16.08	0.81 \pm 0.02	0.90* \pm 0.03
		Erosion	11.10* \pm 4.23	36.85 \pm 15.09	47.95 \pm 15.82	0.83 \pm 0.02	0.89 \pm 0.03

- [3] D. Sinha and E. R. Dougherty, "Fuzzy mathematical morphology," *Journal of Visual Communication and Image Representation*, vol. 3, no. 3, pp. 286–302, Sep. 1992.
- [4] I. Bloch and H. Maitre, "Fuzzy mathematical morphologies: A comparative study," *Pattern Recognition*, vol. 28, no. 9, pp. 1341–1387, Sep. 1995.
- [5] S. Nomura, K. Yamanaka, O. Katai, H. Kawakami, and T. Shiose, "A novel adaptive morphological approach for degraded character image segmentation," *Pattern Recognition*, vol. 38, no. 11, pp. 1961–1975, Nov. 2005.
- [6] Y. Xia, D. Feng, and R. Zhao, "Morphology-based multifractal estimation for texture segmentation," *IEEE Transactions on Image Processing*, vol. 15, no. 3, pp. 614–623, Mar. 2006, conference Name: IEEE Transactions on Image Processing.
- [7] E. Said, D. Nassar, G. Fahmy, and H. Ammar, "Teeth segmentation in digitized dental X-ray films using mathematical morphology," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 2, pp. 178–189, Jun. 2006, conference Name: IEEE Transactions on Information Forensics and Security.
- [8] L. Yucheng and L. Yubin, "An Algorithm of Image Segmentation Based on Fuzzy Mathematical Morphology," in *2009 International Forum on Information Technology and Applications*, vol. 2, May 2009, pp. 517–520.
- [9] J. Angulo, "A Mathematical Morphology Approach to Cell Shape Analysis," in *Progress in Industrial Mathematics at ECMI 2006*, H.-G. Bock, F. De Hoog, A. Friedman, A. Gupta, H. Neunzert, W. R. Pulleyblank, T. Rusten, F. Santosa, A.-K. Tornberg, V. Capasso, R. Mattheij, H. Neunzert, O. Scherzer, L. L. Bonilla, M. Moscoso, G. Platero, and J. M. Vega, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, vol. 12, pp. 543–547, series Title: Mathematics in Industry.
- [10] R. J. Feehs and G. R. Arce, "Multidimensional Morphological Edge Detection," in *Visual Communications and Image Processing II*, vol. 0845. SPIE, Oct. 1987, pp. 285–292.
- [11] R. Krishnapuram and S. Gupta, "Morphological methods for detection and classification of edges in range images," *Journal of Mathematical Imaging and Vision*, vol. 2, no. 4, pp. 351–375, Dec. 1992.
- [12] J. Angulo and J. Serra, "Automatic analysis of DNA microarray images using mathematical morphology," *Bioinformatics*, vol. 19, no. 5, pp. 553–562, Mar. 2003.
- [13] A. Asano, Y. Kobayashi, Chie Muraki, and M. Muneyasu, "Optimization of gray scale morphological opening for noise removal in texture images," in *The 2004 47th Midwest Symposium on Circuits and Systems, 2004. MWCAS '04.*, vol. 1. Hiroshima, Japan: IEEE, 2004, pp. 1_313–1_316.
- [14] D. Schonfeld and J. Goutsias, "Optimal morphological pattern restoration from noisy binary images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 1, pp. 14–29, Jan. 1991, conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [15] R. Peters, "A new algorithm for image noise reduction using mathematical morphology," *IEEE Transactions on Image Processing*, vol. 4, no. 5, pp. 554–568, May 1995, conference Name: IEEE Transactions on Image Processing.
- [16] R. M. Stringhini, D. Welfer, M. C. d'Ornellas, and D. F. T. Gamarra, "A Mathematical Morphology-Based Filter for Noise Reduction and Detail Preservation in Low-Dose Dental CT Images," *Studies in Health Technology and Informatics*, vol. 264, pp. 253–257, Aug. 2019.
- [17] S. Mukhopadhyay and B. Chanda, "A multiscale morphological approach to local contrast enhancement," *Signal Processing*, vol. 80, no. 4, pp. 685–696, Apr. 2000.
- [18] J. C. M. Román, V. R. Fretes, C. G. Adorno, R. G. Silva, J. L. V. Noguera, H. Legal-Ayala, J. D. Mello-Román, R. D. E. Torres, and J. Facon, "Panoramic Dental Radiography Image Enhancement Using Multiscale Mathematical Morphology," *Sensors*, vol. 21, no. 9, p. 3110, Jan. 2021, number: 9 Publisher: Multidisciplinary Digital Publishing Institute.
- [19] Y. Kimori, "Mathematical morphology-based approach to the enhancement of morphological features in medical images," *Journal of Clinical Bioinformatics*, vol. 1, no. 1, p. 33, Dec. 2011.
- [20] S. Amutha, D. Ramesh Babu, M. Ravi Shankar, and N. Harish Kumar, "Mammographic image enhancement using modified mathematical morphology and Bi-orthogonal wavelet," in *2011 IEEE International Symposium on IT in Medicine and Education*, vol. 1, Dec. 2011, pp. 548–553.
- [21] H. yong, T. yin, and L. huijuan, "Noise image restoration based on mathematical morphology," in *The 2nd International Conference on Information Science and Engineering*, Dec. 2010, pp. 3840–3843, iSSN: 2160-1291.
- [22] Zhao Yu-qian, Gui Wei-hua, Chen Zhen-cheng, Tang Jing-tian, and Li Ling-yun, "Medical Images Edge Detection Based on Mathematical Morphology," in *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*. Shanghai, China: IEEE, 2005, pp. 6492–6495.
- [23] K. Lidayová, H. Frimmel, C. Wang, E. Bengtsson, and O. Smedby, "Chapter 12 - Skeleton-based fast, fully automated generation of vessel tree structure for clinical evaluation of blood vessel systems," in *Skeletonization*. Academic Press, Jan. 2017, pp. 345–382.
- [24] R. E. Jurdi, C. Petitjean, P. Honeine, V. Cheplygina, and F. Abdallah, "A Surprisingly Effective Perimeter-based Loss for Medical Image Segmentation," in *Proceedings of the Fourth Conference on Medical Imaging with Deep Learning*. PMLR, Aug. 2021, pp. 158–167.
- [25] S. Shit, J. C. Paetzold, A. Sekuboyina, I. Ezhov, A. Unger, A. Zhyka, J. P. W. Pluim, U. Bauer, and B. H. Menze, "cldice - a novel topology-preserving loss function for tubular structure segmentation," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2021, pp. 16 555–16 564.
- [26] I. Bloch, S. Blusseau, R. Pino Pérez, É. Puybareau, and G. Tochon, "On some associations between mathematical morphology and artificial intelligence," in *Discrete Geometry and Mathematical Morphology*, J. Lindblad, F. Malmberg, and N. Sladoje, Eds. Cham: Springer International Publishing, 2021, pp. 457–469.
- [27] K. Nogueira, J. Chanussot, M. D. Mura, and J. A. D. Santos, "An introduction to deep morphological networks," *IEEE Access*, vol. 9, pp. 114 308–114 324, 2021.
- [28] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [29] S. K. Kumar, "GD doesn't make the cut: Three ways that non-differentiability affects neural network training," Jan. 2024, arXiv:2401.08426 [cs].
- [30] J. Oh and L. Chaparro, "Adaptive fuzzy morphological filtering of images," in *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP '98 (Cat. No.98CH36181)*, vol. 5, May 1998, pp. 2901–2904 vol.5, iSSN: 1520-6149.
- [31] M. Nakashizuka, "A design method for morphological filters with approximations of min/max operators," in *Proceedings : APSIPA ASC*

- 2009 : Asia-Pacific Signal and Information Processing Association, 2009 Annual Summit and Conference. Asia-Pacific Signal and Information Processing Association, 2009 Annual Summit and Conference, International Organizing Committee, oct 2009, pp. 841–844.
- [32] G. Franchi, A. Fehri, and A. Yao, “Deep morphological networks,” *Pattern Recognition*, vol. 102, p. 107246, Jun. 2020, publisher: Elsevier.
- [33] R. Mondal, P. Purkait, S. Santra, and B. Chanda, “Morphological Networks for Image De-raining,” in *Discrete Geometry for Computer Imagery*, M. Couprie, J. Cousty, Y. Kenmochi, and N. Mustafa, Eds. Cham: Springer International Publishing, 2019, pp. 262–275.
- [34] M. J. Menten, J. C. Paetzold, V. A. Zimmer, S. Shit, I. Ezhov, R. Holland, M. Probst, J. A. Schnabel, and D. Rueckert, “A skeletonization algorithm for gradient-based optimization,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 21 394–21 403.
- [35] I. Bloch, “Chapter 3 - fuzzy skeleton and skeleton by influence zones: a review,” in *Skeletonization*, P. K. Saha, G. Borgefors, and G. Sanniti di Baja, Eds. Academic Press, 2017, pp. 71–87.
- [36] N. H. Nguyen, “U-Net based skeletonization and bag of tricks,” in *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. Montreal, BC, Canada: IEEE, Oct. 2021, pp. 2105–2109.
- [37] O. Panichev and A. Voloshyna, “U-Net Based Convolutional Neural Network for Skeleton Extraction,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. Long Beach, CA, USA: IEEE, Jun. 2019, pp. 1186–1189.
- [38] L. Guzzi, M. A. Zuluaga, F. Lareyre, G. Di Lorenzo, S. Goffart, A. Chierici, J. Raffort, and H. Delingette, “Differentiable Soft Morphological Filters for Medical Image Segmentation,” in *MICCAI 2024 - Medical Image Computing and Computer Assisted Intervention*, Marrakesh, Morocco, Oct. 2024.
- [39] J. Angulo, “Pseudo-morphological Image Diffusion Using the Counter-Harmonic Paradigm,” in *Advanced Concepts for Intelligent Vision Systems*, J. Blanc-Talon, D. Bone, W. Philips, D. Popescu, and P. Scheunders, Eds. Berlin, Heidelberg: Springer, 2010, pp. 426–437.
- [40] J. Masci, J. Angulo, and J. Schmidhuber, “A Learning Framework for Morphological Operators Using Counter-Harmonic Mean,” in *Mathematical Morphology and Its Applications to Signal and Image Processing*, C. L. L. Hendriks, G. Borgefors, and R. Strand, Eds. Berlin, Heidelberg: Springer, 2013, pp. 329–340.
- [41] D. Mellouli, T. M. Hamdani, M. B. Ayed, and A. M. Alimi, “Morph-CNN: A Morphological Convolutional Neural Network for Image Classification,” in *Neural Information Processing*, D. Liu, S. Xie, Y. Li, D. Zhao, and E.-S. M. El-Alfy, Eds. Cham: Springer International Publishing, 2017, pp. 110–117.
- [42] D. Mellouli, T. M. Hamdani, J. J. Sanchez-Medina, M. Ben Ayed, and A. M. Alimi, “Morphological convolutional neural network architecture for digit recognition,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 9, pp. 2876–2885, 2019.
- [43] F. Y. Shih, Y. Shen, and X. Zhong, “Development of Deep Learning Framework for Mathematical Morphology,” *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 33, no. 06, p. 1954024, Jun. 2019, publisher: World Scientific Publishing Co.
- [44] Y. Shen, F. Y. Shih, X. Zhong, and I.-C. Chang, “Deep Morphological Neural Networks,” *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 36, no. 12, p. 2252023, Sep. 2022.
- [45] T. Aouad and H. Talbot, “Binary morphological neural network,” in *2022 IEEE International Conference on Image Processing (ICIP)*, 2022, pp. 3276–3280.
- [46] R. Mondal, S. Santra, S. S. Mukherjee, and B. Chanda, “Morphological Network: How Far Can We Go with Morphological Neurons?” Dec. 2022, arXiv:1901.00109 [cs, stat].
- [47] N. S. T. Hirata and G. A. Papakostas, “On Machine-Learning Morphological Image Operators,” *Mathematics*, vol. 9, no. 16, p. 1854, Jan. 2021, number: 16 Publisher: Multidisciplinary Digital Publishing Institute.
- [48] R. Mondal, M. S. Dey, and B. Chanda, “Image Restoration by Learning Morphological Opening-Closing Network,” *Mathematical Morphology - Theory and Applications*, vol. 4, no. 1, pp. 87–107, Jan. 2020, publisher: De Gruyter Open Access.
- [49] R. Pihlak and A. Riid, “Morphological cross entropy loss for improved semantic segmentation of small and thin objects,” *Procedia Computer Science*, vol. 192, pp. 582–591, 2021, knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 25th International Conference KES2021.
- [50] I. M. Copi, C. Cohen, and K. McMahon, *Introduction to Logic*, 14th ed. New York: Routledge, Jul. 2016.
- [51] V. Novák, I. Perfilieva, and J. Mockor, *Mathematical Principles of Fuzzy Logic*, ser. The Springer International Series in Engineering and Computer Science. Springer US, 1999.
- [52] J. Dombi, “A general class of fuzzy operators, the demorgan class of fuzzy operators and fuzziness measures induced by fuzzy operators,” *Fuzzy Sets and Systems*, vol. 8, no. 2, pp. 149–163, 1982.
- [53] L. Lam, S.-W. Lee, and C. Suen, “Thinning methodologies—a comprehensive survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 9, pp. 869–885, 1992.
- [54] M. G. Wagner, “Real-Time Thinning Algorithms for 2D and 3D Images using GPU processors,” *Journal of real-time image processing*, vol. 17, no. 5, pp. 1255–1266, Oct. 2020.
- [55] K. Palágyi and A. Kuba, “A 3D 6-subiteration thinning algorithm for extracting medial lines,” *Pattern Recognition Letters*, vol. 19, no. 7, pp. 613–627, May 1998.
- [56] J. Staal, M. Abramoff, M. Niemeijer, M. Viergever, and B. Van Ginneken, “Ridge-Based Vessel Segmentation in Color Images of the Retina,” *IEEE Transactions on Medical Imaging*, vol. 23, no. 4, pp. 501–509, Apr. 2004.
- [57] V. Mnih, “Machine learning for aerial image labeling,” Ph.D. dissertation, University of Toronto, 2013.
- [58] I. Krasin, T. Duerig, N. Aldrin, V. Ferrari, S. Abu-El-Haija, A. Kuznetsova, H. Rom, J. Uijlings, S. Popov, S. Kamali, M. Mallocci, J. Pont-Tuset, A. Veit, S. Belongie, V. Gomes, A. Gupta, C. Sun, G. Chechik, D. Cai, Z. Feng, D. Narayanan, and K. Murphy, “Openimages: A public dataset for large-scale multi-label and multi-class image classification.” *Dataset available from <https://storage.googleapis.com/openimages/web/index.html>*, 2017.
- [59] J. Wang, K. Markert, and M. Everingham, “Learning models for object recognition from natural language descriptions,” in *Proceedings of the British Machine Vision Conference*, 2009.
- [60] J. Paetzold, O. Schoppe, R. Al-Maskari, G. Tetteh, V. Efremov, M. I. Todorov, R. Cai, H. Mai, Z. Rong, A. Ertuerk, and B. H. Menze, “Transfer learning from synthetic data reduces need for labels to segment brain vasculature and neural pathways in 3D,” in *International Conference on Medical Imaging with Deep Learning—Extended Abstract Track*, Apr. 2019.
- [61] M. Antonelli, A. Reinke, S. Bakas, K. Farahani, A. Kopp-Schneider, B. A. Landman, G. Litjens, B. Menze, O. Ronneberger, R. M. Summers *et al.*, “The medical segmentation decathlon,” *Nature communications*, vol. 13, no. 1, p. 4128, 2022.
- [62] S. Van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, “scikit-image: image processing in python,” *PeerJ*, vol. 2, p. e453, 2014.
- [63] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.
- [64] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, “nnu-net: a self-configuring method for deep learning-based biomedical image segmentation,” *Nature methods*, vol. 18, no. 2, pp. 203–211, 2021.

VI. BIOGRAPHY SECTION

Lisa Guzzi is a Ph.D. student at Inria, researching AI-based medical image analysis, with a focus on the segmentation of the vascular system.

Maria A. Zuluaga is an associate professor at EURECOM and a visiting Senior Lecturer at King’s College London. She holds a chair at the 3IA Côte d’Azur. Her research focuses on the development of machine learning techniques that can be safely deployed in healthcare applications.

Fabien Lareyre MD, PhD, is head of the Department of Vascular Surgery at the Hospital of Antibes and focuses on digital innovation in vascular surgery.

Gilles Di Lorenzo MD, is a consultant in vascular medicine involved in clinical studies on AI in vascular diseases.

Sébastien Goffart is a Ph.D. student at CHU-Nice focusing on AI-based decision support system for Peripheral artery disease.

Andrea Chierici MD, is a consultant in visceral surgery and researcher on applications of AI in medical imaging analysis.

Juliette Raffort MD, PhD, is a consultant at the University Hospital of Nice and holds a chair at 3IA Côte d’Azur. She focuses on applications of AI in vascular diseases, with over 140 publications.

Hervé Delingette is a research director at Inria and chairholder at 3IA Côte d’Azur, focusing on computational medicine and machine learning.