Introducing Multivariate Connected Openings and Closings

Edwin Carlinet and Thierry Géraud

EPITA Research and Development Laboratory (LRDE)
14-16 rue Voltaire, F-94270 Le Kremlin-Bicêtre, France
edwin.carlinet@lrde.epita.fr

Abstract. Component trees provide a high-level, hierarchical, and contrast invariant representation of images, suitable for many image processing tasks. Yet their definition is ill-formed on multivariate data, e.g., color images, multi-modality images, multi-band images, and so on. Common workarounds such as marginal processing, or imposing a total order on data are not satisfactory and yield many problems, such as artifacts, loss of invariances, etc. In this paper, inspired by the way the Multivariate Tree of Shapes (MToS) has been defined, we propose a definition for a Multivariate min-tree or max-tree. We do not impose an arbitrary total ordering on values; we only use the inclusion relationship between components. As a straightforward consequence, we thus have a new class of multivariate connected openings and closings.

Keywords: Openings and closings · Component trees · Hierarchical representation · Color images · Multivariate data · Connected operators

1 Introduction

Mathematical Morphology (MM) offers a large toolbox to design image processing filters that serve as powerful building blocks and enable the user to rapidly build their image processing applications. Operators based on MM have many advantages both from the practitioner and the computer scientist view points: **1.** They are fast to compute, i.e., an erosion/dilation with a structuring element is as fast as a linear filter with a kernel of the same size. **2.** They are numerically stable, i.e., a morphological filter preserves the domain of the values (no rounding is necessary and there is no risk of overflow/underflow during computation). **3.** They are contrast-invariant, i.e., they are suitable for processing low-contrasted images as well as dark/bright images. This property is fundamental to process images with light-varying conditions.

The foundation for MM-based filters relies on complete lattices formed by digital images [20] which are well defined in the binary and grayscale cases for which a natural order exists. On the contrary, there is still no consensus for defining a *natural* order on vectors. The multivariate case is even more subtle when extending more advanced MM filters such as those requiring a total *ranking* of values. For example, *connected* operators, that only change the values of

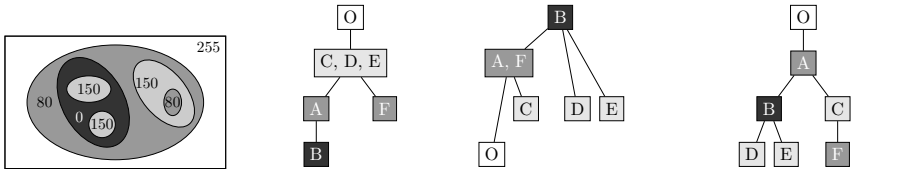

(a) Image (b) Min-Tree (c) Max-Tree (d) Tree of Shapes

Fig. 1: An image (a), and its morphological component trees (b) to (d).

connected components, have desirable contour-preserving properties and form a widely-used class of filters. Those filters require the values to be totally ordered.

To tackle this problem, many attempts have been done to define a "sensitive" total order which are reviewed in Section 2. In this paper, we adopt a different approach, following a *property-based* methodology. After a reminder about connected component-trees in Section 2 and how they relate to connected operators, we explain in Section 3 our requirements for a multivariate component tree and its construction process and its usage to extend connected filters to multivariate data. Its properties are studied in Section 4, and we conclude in Section 5.

Due to limited space, this paper only introduces the Multivariate Component Tree (MCT) and the resulting multivariate openings and closings. As a consequence, experiments with these new tools are kept for later.

2 Mathematical Background and State of the Art

2.1 Trees for Morphological Connected Openings and Closings

Connected operators are widely used in Mathematical Morphology for their properties. A connected operator ψ shares with all morphological filters the desirable property to be contrast change invariant. More formally, given a strictly increasing function $\rho : \mathbb{R} \to \mathbb{R}$ and an image u, ψ must verify $\rho(\psi(u)) = \psi(\rho(u))$. Contrary to structural openings/closings, they may not require prior knowledge about the geometry of what has to be removed. In the binary case, connected operators can only remove some connected sets of pixels. This extends for grayscale images with the threshold-set decomposition principle, connected operators are those that remove (*i.e*, merge) some flat-zones or change their gray level. As a consequence connected operators do not move object boundaries. Opening/closing by reconstruction [29, 22] and area openings [30] are some widely used examples of connected filters. They were then extended to *attribute* filters [2] to express more complex forms of filtering, and Salembier et al. [23] proposed a versatile structure, namely the Max-Tree, that brought the potential of MM approaches to a higher level.

With the Max-Tree, filtering an image is as simple as pruning some branches and removing some nodes. Also, it enables much more powerful image filterings with pruning and non-pruning strategies [25], and with second-generation connectivities [17]. In [8], the authors propose a closely related structure, the Tree of Shapes (ToS), to support self-dual connected operators. The *grain filter* [7] is the self-dual counterpart of the *area* opening and removes *extremal* connected

components. Most importantly, while morphological trees support fast and advanced filters, they are also hierarchical representations of the image (the reader can refer to Sec. 4.3 of Ronse [21]). As such, they enable a multi-scale image analysis and bring us to an higher level of image understanding. Since then, they have been used for image simplification, segmentation, shape-based image retrieval, compression, image registration and more. To cap it all, those trees are fast to compute [4].

More formally, let an image $u : \Omega \to E$ defined on a domain Ω and taking values on a set E equipped with an ordering relation $\leq$. Let $[u < \lambda]$ (resp. $[u > \lambda]$) with $\lambda \in E$ be a threshold set of u (also called respectively lower cut and upper cut) defined as $[u < \lambda] = \{x \in \Omega, u(x) < \lambda\}$. We denote by $\mathcal{CC}(X), X \in \mathcal{P}(\Omega)$ the set of connected components of X. If $\leq$ is a *total* relation, any two connected components $X, Y \in \mathcal{CC}([u < \lambda])$ are either disjoint or nested. The set $\mathcal{S} = \mathcal{CC}([u < \lambda])$ endowed with the inclusion relation forms a tree called the *Min-Tree*. Its dual tree, defined on the upper cuts $\mathcal{S} = \mathcal{CC}([u > \lambda])$, is called the *Max-Tree* (see Figs. 1b and 1c). The last morphological tree, the *Tree of Shapes (ToS)*, depicted in Fig. 1d, is based on the fusion of the Min- and Max-Trees after having filled the holes of their components. Indeed, given the hole-filling operator $\mathcal{H}$, a *shape* is any element of $\mathcal{S} = \{\mathcal{H}(\Gamma), \Gamma \in \mathcal{CC}([u < \lambda])\}_\lambda \cup \{\mathcal{H}(\Gamma), \Gamma \in \mathcal{CC}([u > \lambda])\}_\lambda$. Two shapes being also either nested or disjoint, $(\mathcal{S}, \subseteq)$ also forms a tree.

In the rest of this paper, the Min-Tree, Max-Tree and ToS could be used interchangeably, as we will implicitly consider a set of connected components $\mathcal{S}$ endowed with the inclusion relation (*i.e.*, the cover of $(\mathcal{S}, \subseteq)$) to denote the corresponding tree. Also, without loss of generality, we will consider $E = \mathbb{R}^n$ throughout this paper, and we will note $u = \langle u_1, u_2, \cdots, u_n \rangle$ where u is a multi-band image and u_k are scalar images.

2.2 Connected Openings and Closings for Multi-band Images

A widely spread solution to extend morphological operators to multi-band images is to process the image channel-wise and finally recombine the results. Marginal processing is subject to the well-known false color problem as it may create new values that were not in the original images. False colors may or may not be a problem by itself (e.g. if the false colors are perceptually close to the original ones) but for image simplification it may produce undesirable artifacts.

Since the problem of defining multivariate connected operators lies in the absence of a total order between values, many attempts have been done to define a total order on vectorial data. Two strategies are mainly used: *conditional (C-) ordering* that gives priorities to some (or all) of the vector components, *reduced (R-) ordering* that defines a ranking projection function and orders vectors by their rank (a total pre-order).

Commonly used ranking functions are the l_1-norm or the luminance in a given color space. It makes sense if we assume that the geometric information is mainly held by the luminance [7] but it is not that rare to face images where edges are only available in the color space. In other words, this strategy is not sufficient if the geometric information cannot be summed up to a single dimension. In Tushabe and Wilkinson [24], Perret et al. [19], Naegel and Passat [15],

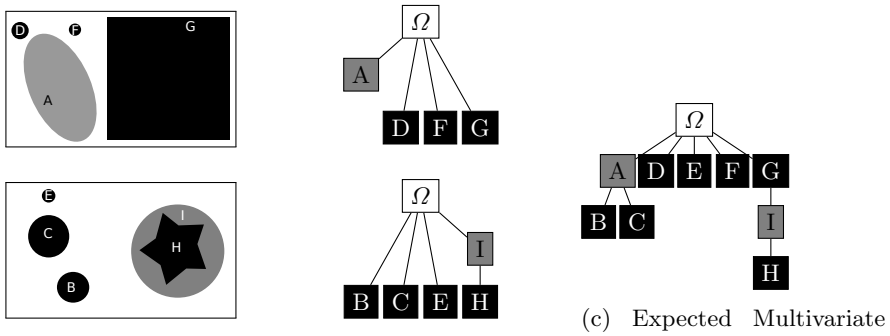

(a) Band u_1 and u_2 (b) Min-Trees of u_1 and u_2 (c) Expected Multivariate Min-Tree for $u = \langle u_1, u_2 \rangle$

Fig. 2: A multi-band image, marginal Min-Tree and the Multivariate Min-Tree that fulfills our requirements.

authors used another widely used R-ordering, the distance to a reference set of values to extend min- and max- trees to multivariate data for image compression or astronomical object detection. This approach is well-founded whenever the background is uniform or defined as set of values. R-orderings are usually combined with C-orderings (typically a lexicographic cascade) to get a strong ordering (as in [9], to extend grain filters on colors).

More advanced strategies have been designed to build a more "sensitive" total ordering that depends on image content. Velasco-Forero and Angulo [27, 28] use machine learning techniques to partition the value space, then a distance to clusters allows to build an ordering. In [13], manifold learning is used to infer a ranking function of values and in [12] a locally-dependent ordering is computed on spatial windows. Lezoray and Elmoataz [11], Lézoray [10] combine both ideas for a manifold learning in a domain-value space capturing small dependencies between a pixel and its neighbors during the construction of the total order. More recently, keeping with content-based ordering approaches, Veganzones et al. [26] proposed an R-ordering based an indexing of the leaves of the image binary partition tree. A review of vector orderings applied to MM can be found in [1].

In [18], the authors propose to deal directly with the partial ordering of values and manipulate the underlying graph structure. While theoretically interesting, a component-graph is algorithmically harder to deal with and the complexity of the structure (in terms of computation time) compels the authors to perform the filtering only locally [16].

Finally, in [5], we introduced the Multivariate Tree of Shapes (MToS) as a novel structure to represent color images. While dedicated to extend the Tree of Shapes, the basic idea lies in a sensible strategy to merge several ToS and thus, can be transposed to merge any set of component trees.

3 Extension of Morphological Trees for Multivariate Data

3.1 Requirements for a Multivariate Component Tree

The construction of the Multivariate Component Tree (MCT) is designed from the following premises. First, we have a channel-wise prior about the objects we

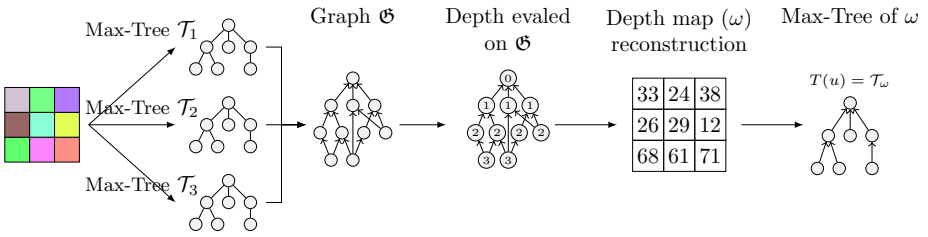

Fig. 3: Scheme of the Multivariate Component Tree construction process.

want to detect and their background (*i.e*, we know channel-wise the contrast "direction"). Second, we have no prior about how objects are spread among the channels, they may be visible in a single, in some or in many channels. Before going further into details of the *what is* a MCT, and *how to get it*, let us start with the *what for* and a description of the properties we want it to have:

Single-band equivalence. On a single-band image, the Multivariate Component Tree must be the same as the normal component tree.

Preservative behavior. Given a connected component C of $\mathcal{S}$, if C is either nested or disjoint to every other connected components, then it must appears in the Multivariate Component Tree. This property actually covers the first one. It is illustrated in Fig. 2 with two Min-Trees; since there is no overlap between components from the two channels of u, the Multivariate Min-Tree is intuitive. Note that the "merging" is purely based on the inclusion, there is no *less-than* relation (in terms of values) between the components from u_1 and u_2. Indeed, in this example, I becomes a child of G while its value in u_2 is greater than in u_1. Being agnostic about the value ordering when merging trees is closely related to the marginal contrast change invariance described hereafter.

Invariance to any marginal increasing change of contrast. For any family of increasing function $\{\rho_1, \rho_2, \cdots \rho_n\}$, $\mathcal{T}(\langle \rho_1(u_1), \rho_2(u_2), \cdots \rho_n(u_n)\rangle) = \mathcal{T}(u)$. This property enforces a fundamental property of morphological representation: well-contrasted and low-contrasted objects are considered equally. This property is twofold. First it enables to consider channels that have different dynamics whereas linear-based approaches generally require a proper prior data normalization. Second, it makes the representation robust when a change of exposure appears only in some channels.

3.2 Multivariate Component Tree Construction

The merging process of the marginal component trees consists in 2 main steps: the construction of a depth map ω and the deduction of a tree from ω. These two steps, designed to get the properties described in Section 3.1, are described hereafter.

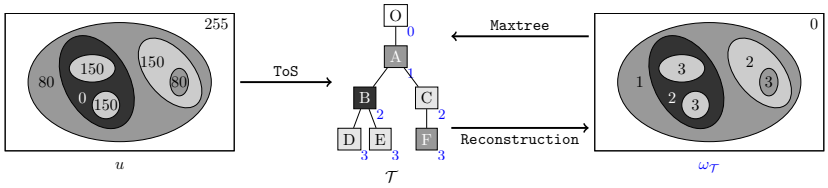

Fig. 4: Equivalence between a hierarchy (here the ToS) and the Max-Tree of its *depth* map.

Getting a tree out of a depth map. Before going further into details, let us explain and motivate the construction of the *depth* map. The depth conveys the object inclusion level. The lighter an object appears in the depth map, the deeper it stands in the object hierarchy. In [6], we showed that given a hierarchy and its corresponding *depth* map, one can recover the initial hierarchy by computing its max-tree, as illustrated in Fig. 4. The interest of the *depth* map lies in the way it abstracts away the underlying value ordering relation between objects. Indeed, no matter the input hierarchy, whether is it a Min-tree (based on the inclusion of *lower* components), Max-Tree (based on the inclusion of *upper* components) or a ToS (based on the inclusion of *hole-filled lower and upper* components), it is equivalent to the Max-Tree computed on its *depth* map. Figure 5 shows the *depth* maps $\omega_{\mathcal{T}min}$, $\omega_{\mathcal{T}max}$, and $\omega_{\mathcal{T}tos}$ of an image computed from its Min-Tree, its Max-Tree and its ToS. As one can see, dark objects appear bright in (b), bright objects appear bright in (c) and the most inner shapes appear bright in (d). Actually, we can show that for a gray-level image u, $\omega_{\mathcal{T}min} = \rho(255 - u)$ and $\omega_{\mathcal{T}max} = \rho(u)$ with ρ some increasing contrast change. As a consequence, it is straightforward that the Max-Tree of $\omega_{\mathcal{T}}$ is $\mathcal{T}$ itself.

The *depth* map leads to an alternate representation of the image content. Instead of having a representation based on the brightness, we now have a pixel-wise interpretation of the *inclusion level* of objects. In other words, if a pixel appears brighter than its neighbor in the *depth* map, this means that it belongs to a nested sub-object.

From marginal component trees to the graph of components. The component trees $\mathcal{T}_1, \mathcal{T}_2, \ldots, \mathcal{T}_n$ are computed on each band $u_1, u_2, \ldots, u_n$ of the

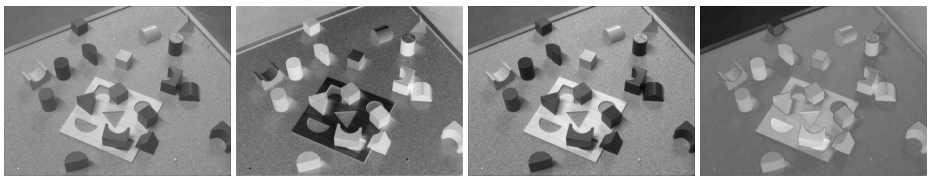

(a) Input u (b) $\omega_{\mathcal{T}min}$ (c) $\omega_{\mathcal{T}max}$ (d) $\omega_{\mathcal{T}tos}$

Fig. 5: An input gray-level image (a) and its corresponding *depth* map for the Min-Tree (b), Max-Tree (c) and ToS (d) hierarchies. The dynamic of the depth image has been stretched to [0-255].

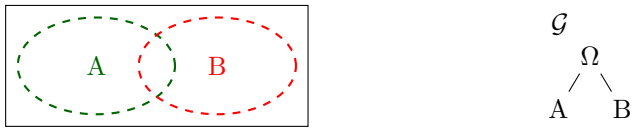

Fig. 6: A GoC that is an invalid morphological tree. Left: two marginal components from two channels are overlapping. Right: The corresponding Graph of Components (GoC) $\mathfrak{G}$. The points of $A \cap B$ belong to both nodes A and B, so the tree $\mathfrak{G}$ is not valid.

input image. Each tree is associated with its set of components $\mathcal{S}_1$, $\mathcal{S}_2$, ..., $\mathcal{S}_n$. Let $\mathcal{S} = \bigcup \mathcal{S}_i$, we call the Graph of Components (GoC) $\mathfrak{G}$ the cover of $(\mathcal{S}, \subseteq)$. The GoC, depicted in Fig. 3, is actually the inclusion graph of all the connected components computed marginally.

It is worth mentioning even if the GoC is actually a tree, it is not a valid morphological tree. Indeed, in "standard" morphological hierarchies (min-/max-trees) and their extension (the component-graph [18]), for any point x, there exists a *single* smallest component that contains x. As a consequence, a point belongs to a *single* node in the structure. In the GoC, a point may belong to several nodes. For example, in Fig. 6, the points in $(A \cap B)$ belong to both nodes A and B, but $(A \cap B)$ does not exist as a node in any marginal tree. Thus, we cannot just extract a tree (e.g. the minimum spanning tree) from the GoC as it will not be valid.

It is also worth mentioning that the graph of components $\mathfrak{G}$ is different from the component-graph introduced by Passat and Naegel [18]. The latter is the inclusion graph of all connected components based on a partial ordering $\prec$ (*i.e*, the set $\{\mathcal{CC}([u \prec \lambda]), \lambda \in \mathbb{R}^n\}$) while $\mathfrak{G}$ is the inclusion graph of connected components computed marginally (*i.e*, the set $\bigcup_{k=1}^{n} \{\mathcal{CC}([u_k < \lambda]), \lambda \in \mathbb{R}\}$). $\mathfrak{G}$ is thus a sub-graph of the component-graph.

Note that the graph is a "complete" representation of the input image u that can be reconstructed from $\mathfrak{G}$. Indeed, without loss of generality, suppose that $\mathfrak{G}$ is built from $\mathcal{T}_1, \cdots, \mathcal{T}_n$ being Max-Trees. For a component A of $\mathfrak{G}$, let $\lambda_i(A) = \min\{u_i(x), x \in A\}$, then the input can be reconstructed with:

$$u(x) = \left\langle \max_{X:\, x \in X} \lambda_1(X), \max_{X:\, x \in X} \lambda_2(X), \cdots, \max_{X:\, x \in X} \lambda_n(X) \right\rangle. \tag{1}$$

Constructing a *depth* map from the graph of components. As said previously, getting a morphological tree out of the GoC is not as simple as it seems (e.g. with a MST algorithm) as one has to ensure that pixels do not belong to disjoint branches of the tree. We have also seen that the *depth* map is an interesting intermediate representation, which provides pixel-wise description of the organization of the scene (in terms of inclusion) and enables us to extract a tree out of it. As a consequence, our objective is now to build a *depth* map from the GoC. Let ρ be the *depth* of a node A in $\mathfrak{G}$, *i.e*, $\rho(A)$ is the length of the longest path of a component A from the root. Let $\omega : \Omega \to \mathbb{R}$ defined as:

$$\omega(x) = \max_{X \in \mathcal{S}, x \in X} \rho(X). \tag{2}$$

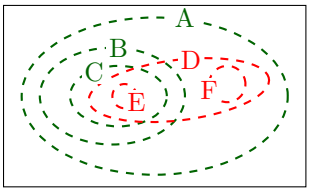

(a) A 2-band image u and its components.

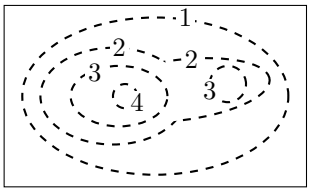

(c) ω image built from $\mathfrak{G}$

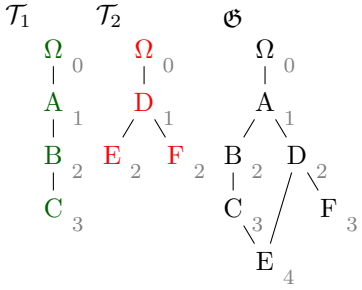

(b) The marginal Max-Trees $\mathcal{T}_1$, $\mathcal{T}_2$ and the GoC $\mathfrak{G}$. The depth appears in gray near the nodes.

(d) The Max-Tree $\mathcal{T}_\omega$

Fig. 7: Merging trees with a depth map; in this simple example, we suppose that u verifies: $u_1(\Omega) < u_1(A) < u_1(B) < u_1(C)$ and $u_2(\Omega) < u_2(D) < u_2(E) < u_2(F)$.

The map ω associates each point x with the depth of the deepest component containing x . Let $\mathbb{C} = \bigcup_{h\in\mathbb{R}} \mathcal{CC}([\omega \geq h])$. $(\mathbb{C}, \subseteq)$ is actually the max-tree of ω. The method is illustrated in Fig. 7 where a two-channel image has overlapping components from red and green components. The trees $\mathcal{T}_1$ and $\mathcal{T}_2$ are computed marginally and merged into the GoC $\mathfrak{G}$ for which we compute the depth of each node (Fig. 7b). Those values are reported in the image space, pixel-wise (Fig. 7c). This is the step which decides which components are going to merge; here B and D are set to the same value. This choice is based on the level of inclusion and no longer on the original pixel values.

While the graph is a complete representation of u, once $\mathfrak{G}$ is flattened to ω, it is not possible to go backward and recover the original connected components. This is the only lossy step part of the process.

Filtering and reconstruction with the Multivariate Component Tree. Morphological trees enables to perform openings, closings and extremal filterings in two steps: **1.** a pruning where nodes are removed and pixels are re-assigned; **2.** a restitution where pixels get their final values.

Let $\mathbb{P}$ be a binary predicate that tells if a node is to be removed. For example, $\mathbb{P}(X) = \bar{X} > \alpha$ denotes a predicate that retains all components whose size is above α (grain filter, area opening and closing). Similarly to Eq. (1), the reconstruction after filtering the Max-Tree $\mathcal{T} = (\mathcal{S}, \subseteq)$ of a scalar image u is:

$$\tilde{u}(x) = \max_{X\in\mathcal{S},\, x\in X \,|\, \mathbb{P}(X)} \lambda(X). \tag{3}$$

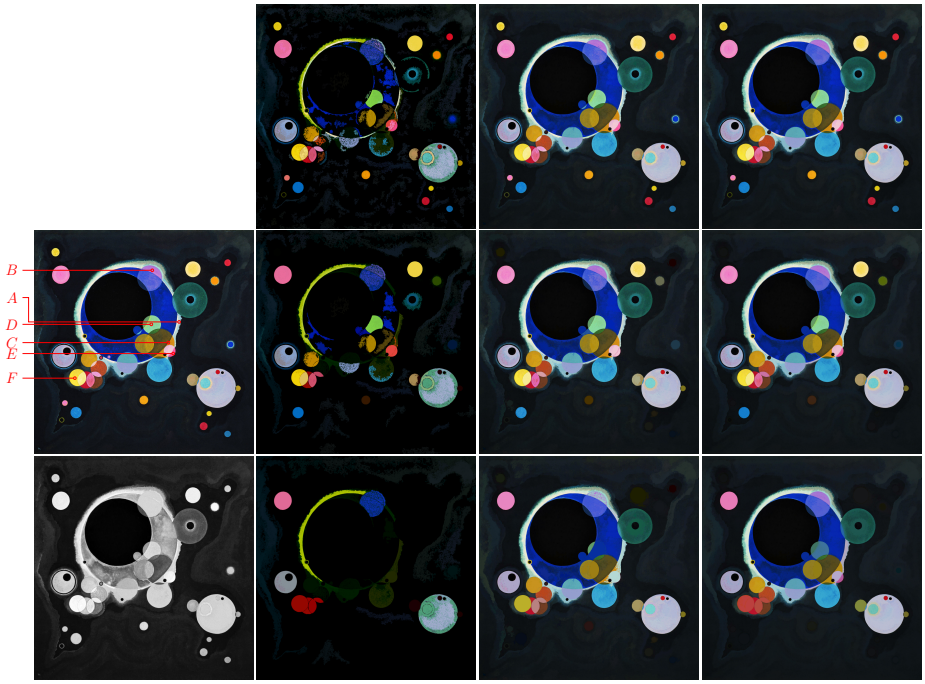

Fig. 8: Comparison of a marginal opening and multivariate openings with several restitution strategies. (a) Original image (top) and its depth map ω_u (bottom). (b) Multivariate opening with restitution strategy *Rinf*. (c) Multivariate opening with restitution strategy *RNC*. (d) Marginal opening. Openings are of sizes 500, 4000 and 16000 (top-down).

Equation (3) actually describes the **direct** filtering, the pixels are affected to the level of the *last surviving* node in the hierarchy. This is one of the four filtering strategies described in [14]. The others -namely *min*, *max*, *subtractive*- only make sense for a non-increasing predicate, that is a predicate that may keep a node while its parent is going to be deleted. Interested reader should refer to [14] for more details on filtering strategies.

Last, we need to define the restitution for multivariate data. One could simply define the restitution as in Eq. (1) with a predicate constraint as shown in Eq. (5):

$$\Gamma(x) = \{X \in \mathcal{S}_w \mid x \in X, \mathbb{P}(X)\} \tag{4}$$

$$\tilde{u}(x) = \left\langle \max_{X \in \Gamma(x)} \lambda_1(X), \max_{X \in \Gamma(x)} \lambda_2(X), \cdots, \max_{X \in \Gamma(x)} \lambda_n(X) \right\rangle. \tag{5}$$

Simply stated, each pixel is assigned the infimum of the pixels of its node. The painting *Several Circles*[1] Fig. 8 does not reveal how the filters behave on "real"

[1] *Einige Kreise*, from Vasily Kandinsky, 2016 (Solomon R. Guggenheim Museum, New York, ©2018 Artists Rights Society, New York/ADAGP, Paris).

natural images but enables us to illustrate the troubles we can face when filtering colors. Fig. 8a shows that neighboring regions with non-comparable colors (like those of A, B and C) merge at the same level.

The restitution strategy from Eq. (5) (let us call it *Rinf*) leads to several artifacts as shown in Fig. 8b. First, this restitution may affect pixels that belong to components that should be preserved. Second, it creates *false* colors that were not present in the original image (as in region D).

This second problem is shared with the *marginal* processing strategy shown in Fig. 8d. It creates false-colors because it combines marginal filtering results (e.g. in region F). It is also worth noting that, for the same reason, it creates flat zones that are smaller than the filtering grain size (e.g. the region C) and blurs object boundaries. That is why a marginal processing is a legitimate filtering approach when the quality is assessed according to the human perception.

The restitution strategy proposed in [24] tackles those problems by keeping the original pixel values when they belong to a node "living" and using the a color of the contour of the connected component when they belong to a "dead" one. Results of the nearest color strategy are shown in Fig. 8c. It combines the strength of a vectorial approach that produces "real" flat zones with no contour blurring and a quite qualitative color filtering similar to the one obtained by marginal filtering. Yet, there remain some restitution quality troubles for nodes which merge non-comparable regions as we need to choose a single color to reconstruct both regions. It yields color artifacts as in B, and C where orange and magenta pixels have been colorized in cyan from A.

4 Algebraic Properties and Complexity

First, trivially, the Multivariate Min-Tree and the Multivariate Max-Tree are dual by construction. The *preservative behavior* property is also ensured by construction. If a component A is either nested or disjoint to any other component, then this is a "bottleneck" node in $\mathfrak{G}$, in the sense that every path leading to a sub-component of A needs to pass through this node. As a consequence, every pixels in A will have a depth ω greater than $\rho(A)$ (*i.e*, in $\mathcal{CC}[\omega > \rho(A)]$) and there exists a node in T_ω corresponding to A.

The openings described previously meet the prerequisites described in Section 3.1. Due to the lack of space, we only give the intuition of the proofs. The *single-band equivalence* is straightforward; if there is a single input tree $\mathcal{T}_1$, the graph $\mathfrak{G}$ is thus a tree and $\mathcal{T}_\omega = \mathcal{T}_1$. Whichever restitution strategy is then used, they all are equivalent in this case and produce the same output as the regular attribute opening. The *marginal-contrast-change invariance* is also straightforward; since each tree T_i is contrast change invariant for the i-th channel (and does not depend on the others), so does the graph $\mathfrak{G}$. Also, since the rest of the process only depends on the inclusion relation between components, the whole process is thus marginal-contrast-change invariant.

Morphological properties for multivariate openings and closings based on $\mathcal{T}_\omega$ depends on the restitution strategy. With the strategy *Rinf*, the openings (resp. closings) are marginally anti-extensive (resp. extensive). However, in the general case they are neither idempotent, nor marginally anti-extensive.

From a computational standpoint, in the classical case of images with low-quantized values (for instance, with values encoded on 8 bit in every channel / band), the marginal trees can be computed in linear time w.r.t. the number N of pixels [4]. The most expensive part is the graph computation which is $O(n^2.H.N)$, where n is the number of channels, H the maximal depth of the trees, and N the number of pixels.

5 Conclusion

This paper introduced preliminary ideas to extend connected filters on multi-band images[2]. This approach relies on the extension of component trees to data for which no natural total order exists. Instead of imposing an arbitrary total ordering, the method tries to combine and merge connected components from several marginal component trees. As such, it can be seen as a local and context-dependent ordering of the pixels based on their level of inclusion in each hierarchy. Beyond multivariate connected filters, our method produces a single hierarchical representation of the image, the MCT, that can be used to study "shapes" in images [3]. As a further work, we plan to figure out a way to prioritize some bands while keeping the idea of "ordering by inclusion level". We also plan to transform the value space before using the MCT (whether by a principal component analysis, or by projection in a given color-space) to study the effect of the data correlation on the merging process used by the MCT.

References

1. E. Aptoula and S. Lefèvre. A comparative study on multivariate mathematical morphology. *Pattern Recognition*, 40(11):2914–2929, 2007.
2. E. J. Breen and R. Jones. Attribute openings, thinnings, and granulometries. *Computer Vision and Image Understanding*, 64(3):377–389, 1996.
3. F. Cao, J.-L. Lisani, J.-M. Morel, P. Musé, and F. Sur. *A Theory of Shape Identification*, volume 1948 of *Lecture Notes in Mathematics*. Springer, 2008.
4. E. Carlinet and T. Géraud. A comparative review of component tree computation algorithms. *IEEE Trans. on Image Processing*, 23(9):3885–3895, 2014.
5. E. Carlinet and T. Géraud. MToS: A tree of shapes for multivariate images. *IEEE Trans. on Image Processing*, 24(12):5330–5342, 2015.
6. E. Carlinet, T. Géraud, and S. Crozet. The tree of shapes turned into a max-tree: A simple and efficient linear algorithm. In *Proc. of the Intl. Conf. on Image Processing (ICIP)*, pages 1488–1492, Athens, Greece, October 2018.
7. V. Caselles and P. Monasse. Grain filters. *JMIV*, 17(3):249–270, 2002.
8. V. Caselles, B. Coll, and J.-M. Morel. Topographic maps and local contrast changes in natural images. *International Journal of Computer Vision*, 33(1):5–27, 1999.
9. B. Coll and J. Froment. Topographic maps of color images. In *Proc. of the Intl. Conf. on Pattern Recognition (ICPR)*, volume 3, pages 609–612, 2000.
10. O. Lézoray. Complete lattice learning for multivariate mathematical morphology. *Journal of Visual Communication and Image Representation*, 35:220–235, 2016.
11. O. Lezoray and A. Elmoataz. Nonlocal and multivariate mathematical morphology. In *Proc. of the Intl. Conf. on Image Processing (ICIP)*, pages 129–132, 2012.

[2] Some extra illustrations can be found on `https://publications.lrde.epita.fr/carlinet.19.ismm`

12. O. Lezoray, C. Meurie, and A. Elmoataz. A graph approach to color mathematical morphology. In *Proceedings of the IEEE International Symposium on Signal Processing and Information Technology*, pages 856–861, 2005.
13. O. Lezoray, C. Charrier, and A. Elmoataz. Rank transformation and manifold learning for multivariate mathematical morphology. In *Proc. of European Signal Processing Conference*, volume 1, pages 35–39, 2009.
14. A. Meijster, M. Westenberg, M. Wilkinson, and M. Hamza. Interactive shape preserving filtering and visualization of volumetric data. In *Proc. of the 5th IASTED Conf. on Computer Graphics and Imaging (CGIM)*, pages 640–643, 2002.
15. B. Naegel and N. Passat. Component-trees and multi-value images: A comparative study. In *Proc. of ISMM*, volume 5720 of *LNCS*, pages 261–271. Springer, 2009.
16. B. Naegel and N. Passat. Towards connected filtering based on component-graphs. In *Proc. of ISMM*, volume 7883 of *LNCS*, pages 353–364. Springer, 2013.
17. G. K. Ouzounis and M. H. Wilkinson. Mask-based second-generation connectivity and attribute filters. *IEEE Trans. on PAMI*, 29(6):990–1004, 2007.
18. N. Passat and B. Naegel. An extension of component-trees to partial orders. In *Proc. of the Intl. Conf. on Image Processing (ICIP)*, pages 3933–3936, 2009.
19. B. Perret, S. Lefèvre, C. Collet, and E. Slezak. Connected component trees for multivariate image processing and applications in astronomy. In *Proc. of the Intl. Conf. on Pattern Recognition (ICPR)*, pages 4089–4092, Aug. 2010.
20. C. Ronse. Why mathematical morphology needs complete lattices. *Signal processing*, 21(2):129–154, 1990.
21. C. Ronse. Ordering partial partitions for image segmentation and filtering: Merging, creating and inflating blocks. *Journal of Mathematical Imaging and Vision*, 49(1):202–233, 2014.
22. P. Salembier and J. Serra. Flat zones filtering, connected operators, and filters by reconstruction. *IEEE Trans. on Image Processing*, 4(8):1153–1160, 1995.
23. P. Salembier, A. Oliveras, and L. Garrido. Antiextensive connected operators for image and sequence processing. *IEEE TIP*, 7(4):555–570, 1998.
24. F. Tushabe and M. Wilkinson. Color processing using max-trees: A comparison on image compression. In *Proceedings of the International Conference on Systems and Informatics (ICSAI)*, pages 1374–1380, 2012.
25. E. R. Urbach and M. H. F. Wilkinson. Shape-only granulometries and grey-scale shape filters. In *Mathematical Morphology and Its Applications to Signal and Image Processing—Proceedings of the International Symposium on Mathematical Morphology (ISMM)*, volume 2002, pages 305–314, 2002.
26. M. A. Veganzones, M. Dalla Mura, G. Tochon, and J. Chanussot. Binary partition trees-based spectral-spatial permutation ordering. In *Proc. of the Intl. Symp. on Mathematical Morphology (ISMM)*, pages 434–445. Springer, 2015.
27. S. Velasco-Forero and J. Angulo. Supervised ordering in $\mathcal{R}_p$: Application to morphological processing of hyperspectral images. *IEEE Trans. on Image Processing*, 20(11):3301, 2011.
28. S. Velasco-Forero and J. Angulo. Random projection depth for multivariate mathematical morphology. *IEEE Journal of Selected Topics in Signal Processing*, 6(7): 753–763, 2012.
29. L. Vincent. Morphological grayscale reconstruction in image analysis: Applications and efficient algorithms. *IEEE Trans. on Image Processing*, 2(2):176–201, 1993.
30. L. Vincent. Morphological area openings and closings for grey-scale images. In *Shape in Picture: Mathematical Description of Shape in Grey-level Images*, pages 197–208. Springer, 1994.