University of Groningen

Partition-induced connections and operators for pattern analysis

Ouzounis, Georgios K.; Wilkinson, Michael H.F.

Published in:
Pattern recognition

DOI:
[10.1016/j.patcog.2009.10.002](https://doi.org/10.1016/j.patcog.2009.10.002)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2010

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Ouzounis, G. K., & Wilkinson, M. H. F. (2010). Partition-induced connections and operators for pattern analysis. *Pattern recognition*, 43(10), 3193-3207. <https://doi.org/10.1016/j.patcog.2009.10.002>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Partition-induced connections and operators for pattern analysis

Georgios K. Ouzounis^{a,b,1}, Michael H.F. Wilkinson^{a,*}

^a Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, P.O. Box 407, 9700 AK Groningen, The Netherlands

^b School of Medicine, Democritus University of Thrace, University General Hospital of Alexandroupoli, 68100 Alexandroupoli, Greece

ARTICLE INFO

Article history:

Received 25 January 2008

Received in revised form

8 September 2009

Accepted 7 October 2009

Keywords:

Image analysis

Mathematical morphology

Connected filters

Connectivity classes

Diatoms

ABSTRACT

In this paper we present a generalization on the notion of image connectivity similar to that modeled by *second-generation* connections. The connected operators based on this new type of connection make use of image partitions aided by mask images to extract path-wise connected regions that were previously treated as sets of singletons. This leads to a redistribution of image power which affects texture descriptors. These operators find applications in problems involving contraction-based connectivities, and we show how they can be used to counter the over-segmentation problem of connected filters. Despite restrictions which prevent extensions to gray-scale, we present a method for gray-scale spectral analysis of biomedical images characterized by filamentous details. Using connected pattern spectra as feature vectors to train a classifier we show that the new operators outperform the existing contraction-based ones and that the classification performance competes with, and in some cases outperforms methods based on the standard 4- or 8-connectivity. Finally, combining the two methods we enrich the texture description and increase the overall classification rate.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

In image analysis it is often desirable to sort objects based on their structural characteristics, typically expressed by means of some attribute measure. The pattern spectrum [1] is a commonly used method that features ordered attribute classes that keep track of the amount of image detail or power (measured in number of pixels, or sum of gray levels) that falls within their range of attribute values.

Pattern spectra can be computed from granulometries [2–8] which are ordered sets of morphological operators adhering to some properties discussed later. The operators can either be structural or connected filters, each allowing a limited range of image details to pass. A recent comparison between the two filter types favored granulometries based on connected filters [9]. The method presented for computing pattern spectra, compared with other existing methods was shown to be rotation invariant, significantly less sensitive to noise, to allow for multi-dimensional spectra to be computed based on strict size and shape attributes, and its computation time to be independent of the number of scales or shape classes used.

Connected filters are shape preserving operators [4,10] which work by removing or retaining connected image regions known as *connected components* in their entirety, without edge distortion. If filtering is based on some component attribute they are referred to as *attribute* or *grain filters* [2,4,5]. Such filters have been used among other areas, in biomedical imaging for filament enhancement and area/volume filtering [11–13]. Connected operators rely on some notion of image connectivity, commonly the 4 and 8 pixel adjacency relations [14]. With connectivity expressed in a set-theoretic framework [3,15], several generalizations were introduced which overcame topological constraints of earlier formalisms. An example is the *second-generation connectivity* [10,16–20] in which operators associated to it can access families of sets that account for connected components that are not strictly 4 or 8 connected. Typically we refer to clusters or subregions of connected components according to standard connectivity or for combinations of the two. An efficient scheme allowing for all three cases has been recently introduced and is termed *mask-based connectivity* [20]. According to this, the connectivity of an image given the standard 4 and 8 connection can be obtained by a second image (the *connectivity mask*) that commonly results from the application of some operator on the original. Like earlier formulations, connectivity openings associated to mask-based connectivity extract the connected components of interest and handle the remaining structures as singleton sets. Generating singletons is a feature that has been used to counter the *leakage* problem of connected operators [6]. Leakage results from thin elongated paths connecting different objects in an image that should be treated individually, and is usually caused by background texture,

* Corresponding author.

E-mail addresses: gouzoun@med.duth.gr, Georgios.Ouzounis@jrc.ec.europa.eu (G.K. Ouzounis), m.h.f.wilkinson@rug.nl (M.H.F. Wilkinson).

¹ Current address: Global Security And Crisis Management Unit, IPSC – Joint Research Centre– European Commission, T.P. 267, Via E. Fermi 2749, I-21027 Ispra (VA), Italy.

noise or other image details. The treatment of these paths as groups of singleton sets though, is known under certain conditions to generate problems in both filtering and segmentation.

It is a condition known as the *over-segmentation* problem of connected operators, and is formally presented in [21]. Over-segmentation in this sense, as opposed to the equivalent notion in image segmentation, refers explicitly to the interpretation of certain image regions as sets of singletons. It can either be countered by region merging methods or prevented as in [22]. In the latter case, a second-generation connected operator was presented to handle points previously treated as singletons, in a regular connectivity fashion. That is, to extract path-wise connected structures in regions outside the connectivity mask, i.e. regions in which the same pixel has lower intensity in the mask image compared to the original. The operator in [22] employs a single mask image to shape the connected components, both those bounded by the mask and those outside it. Moreover, though not shown, families of these operators associate to a rather restrictive connectivity scheme of limited use.

An example of over-segmentation is given in Fig. 1 where a CT slice of the lungs is shown. To isolate and further process the sphere-like structures observed within the lungs, a second-generation connectivity opening with a mask given by an opening or an erosion of the original, is sufficient. This however distorts the edges heavily as shown in Fig. 1(b), and moreover does not allow us to automatically extract the tubular structures on which any of these spheres is attached. That is because all thin and elongated structures are converted to singletons as shown in Fig. 1(c).

In this paper we revise the idea of such an operator and present a similar one, only relying on an image partition instead of a single mask. This is through the newly presented connectivity scheme, termed π -connectivity (Section 3) which allows further flexibility compared to mask-based second-generation connectivity. In this new framework connectivity masks are used to represent the classes of a partition. This allows the explicit treatment of image regions that were inaccessible with earlier connectivity schemes. Returning to the problem of over-segmentation, the new operator involving the π -connection, handles points outside the mask as path-wise connected structures in a similar way to that in [22]. If the mask is the result of an anti-extensive operator such as an erosion or an opening (also called as a *contractive operator*) on the original image, these remaining structures are mostly thin elongated segments with size and shape dependant to the structuring element (SE) used. Thus, adjusting the SE accordingly, allows for a whole new set of object features to be extracted. To harness this texture information we investigate means to compute π -connected pattern spectra. Their computation is not trivial since we will show that the specific usage of π -connected operators does not allow for their extension to gray-scale using the existing filtering rules (Section 3.4). To address this issue, we first formalize the Max-Tree-based method from [9] in a mathematical statement (Section 4). This is a

method for computing granulometries and further connected pattern spectra without iterative filtering, using a tree-based image representation structure; the Max-Tree [6]. From this we deduce an expression for gray-scale connected pattern spectra based on the threshold decomposition of gray-scale images. Customizing this for our needs is done by replacing the standard connected operator with the respective π -connected operator.

To demonstrate the potential of this new class of operators we demonstrate on two medical images, the types of structures we can extract and experiment with texture based classification of diatom images using the spectrum as a feature vector. We compare its performance with that achieved using pattern spectra based on standard and contraction-based second-generation connectivities (Section 6). The results are discussed in Section 7.

2. Theory

2.1. Connections and connected operators

In mathematical morphology the concept of connectivity is defined by the notion of *connectivity classes* [3,15]:

Definition 1. Let E be an arbitrary space. A connectivity class or *connection* C is any family in $\mathcal{P}(E)$ that satisfies:

- (1) $\emptyset \in C$ and for all $x \in E$, $\{x\} \in C$,
- (2) for any $\{A_i\} \subseteq C$ for which $\bigcap A_i \neq \emptyset \Rightarrow \bigcup A_i \in C$.

This means that both the empty set and singleton sets, denoted as $\{x\}$, are connected, and any union of elements of C which have a non-empty intersection is also connected. The members of C are called *connected sets* and are element groupings of E .

Every set $X \subseteq E$ can be written as the union of pairwise disjoint connected sets of maximal extent, C_i . Maximality in this sense means that given a set C_i there can be no other set $C_j \supset C_i$ such that $C_j \subseteq X$ and $C_j \in C$. A set $C_i \subseteq X$ also called a *connected component* (see Fig. 2(d)) or a *grain* of X , given a point $x \in C_i$ is addressed by a *connectivity opening* which is an operator defined as

$$\Gamma_x(X) = \bigcup \{A_i \in C | x \in A_i \text{ and } A_i \subseteq X\}. \quad (1)$$

With all A_i containing x in their intersection, their union $\Gamma_x(X)$ is also connected and furthermore, $\Gamma_x(X) = \emptyset$ if $x \notin X$.

Connectivity openings are algebraic openings and therefore are *anti-extensive*, *increasing* and *idempotent* operators. For any set X each property implies the following:

- (1) Anti-extensivity: $\Gamma_x(X) \subseteq X$.
- (2) Increasingness: if $X \subseteq Y \Rightarrow \Gamma_x(X) \subseteq \Gamma_x(Y)$.
- (3) Idempotence: $\Gamma_x(\Gamma_x(X)) = \Gamma_x(X)$.

Fig. 1. Example of over-segmentation in the study of lung diseases. The original image (a); the connectivity mask m which is a replica of the original after an opening (b), and the top hat after contrast enhancement showing which structures are converted to singletons (c).

Fig. 2. Connected image regions—top row: original image (a); a connectivity mask by a closing on the original (b); a partition of (a) with its classes separated by the gray line (c). Bottom row and for a point $x \in X$, on the dragon's head: the corresponding connected component (d); the clustering-based m -connected component (e); the π -connected component (f).

The operator Γ_x is explicitly related to a connectivity class $\mathcal{C}$ if satisfying the set of conditions given by Serra [3] (also in [4,10]) in the following theorem:

Theorem 1. *The datum of a connectivity class $\mathcal{C}$ on $\mathcal{P}(E)$ is equivalent to the family $\{\Gamma_x | x \in E\}$ of openings on x such that:*

- (1) every Γ_x is an algebraic opening,
- (2) for all $x \in E$, we have $\Gamma_x(X) = \{x\}$,
- (3) for all $X \subseteq E$, $x, y \in E$, $\Gamma_x(X)$ and $\Gamma_y(X)$ are equal or disjoint,
- (4) for all $X \subseteq E$, and all $x \in E$, we have $x \notin X \Rightarrow \Gamma_x(X) = \emptyset$.

Concluding, it can be seen that connectivity openings characterize uniquely the connectivity class they are associated with and there is a one-to-one correspondence between the two.

2.2. Second-generation connectivity

The definition of connectivity by means of connectivity classes allows several generalizations. Second-generation connectivity is such an example where from a given “parent” class $\mathcal{C}$ we derive a “child” class given some image transformation captured by the corresponding connectivity opening. This concept is modeled by two types of connections, the *clustering* and *contraction*-based connectivity classes.

When clustering, disconnected components according to $\mathcal{C}$ satisfying some structural criteria, most commonly the distance separating them, are extracted as a single entity. By contrast, in a contractive transformation object regions that fail some structural criteria, most commonly the local width, are converted to singletons, splitting wide object regions connected by narrow bridges apart. The two connectivity transformations can be combined in a single framework known as *mask-based* or *m-connectivity* [20] in which the grains of a mask image M are used to selectively carry out clusterings or contractions on the connected components of the original.

A mask-based connectivity class is defined as follows:

$$\mathcal{C}^M = \{\emptyset\} \cup \mathcal{S} \cup \{A \subseteq E | \exists x \in E : A \subseteq \Gamma_x(M)\} \quad (2)$$

and a connectivity opening from the corresponding family $\{\Gamma_x^M(X) | x \in E\}$ as

$$\Gamma_x^M(X) = \begin{cases} \Gamma_x(M) \cap X & \text{if } x \in X \cap M, \\ \{x\} & \text{if } x \in X \setminus M, \\ \emptyset & \text{otherwise.} \end{cases} \quad (3a)$$

$$\Gamma_x^M(X) = \begin{cases} \{x\} & \text{if } x \in X \setminus M, \\ \emptyset & \text{otherwise.} \end{cases} \quad (3b)$$

$$\Gamma_x^M(X) = \emptyset \quad \text{otherwise.} \quad (3c)$$

The term $\mathcal{S}$ denotes the family of all singleton sets of E . The family of operators Γ_x^M essentially “masks” the desired members of $\mathcal{C}$ to $\mathcal{C}^M$ by selecting all subsets of X found within the grains of M . An important feature of the definition above is that there are no

assumptions as to how M should be generated. This eliminates constraints in the ways the image domain can be connected. An example of a clustering-based m -connected component is shown in Fig. 2(e).

2.3. Attribute filters

The notion of a connected filter in mathematical morphology describes a mapping $\psi: \mathcal{P}(E) \rightarrow \mathcal{P}(E)$ that is increasing and idempotent [3,4,23]. The connectivity opening is a trivial example and based on it we can define a number of other connected filters that work by imposing constraints on the connected components it returns. Constraints are commonly expressed in the form of attribute criteria to accept or to reject connected components based on some attribute measure. Attribute criteria A are put in place by means of a trivial opening Γ_A . The latter is defined as an operator $\Gamma_A: \mathcal{C} \rightarrow \mathcal{C}$ which if applied on a connected component $C \in \mathcal{C}$ yields C if $A(C)$ is true, and $\emptyset$ otherwise. Obviously, $\Gamma_A(\emptyset) = \emptyset$. Attribute criteria are often expressed as

$$A(C) = \text{Attr}(C) \geq \lambda \quad (4)$$

with $\text{Attr}(C)$ some real-valued attribute of C , and λ an attribute threshold.

Definition 2. The binary attribute opening Γ^A of a set X with an increasing criterion A is given by

$$\Gamma^A(X) = \bigcup_{x \in X} \Gamma_A(\Gamma_x(X)). \quad (5)$$

An example is the area opening [24,25].

Attribute-based connected operators may also be based on shape criteria rather than size. They are generally non-increasing operators which are scale, rotation and translation invariant. A shape operator that is also idempotent defines a *shape filter* and an example is the *attribute thinning* Φ^A [2,4]. An example of a shape criterion is the non-compactness (also referred to as elongation) criterion [11,13] given by

$$\text{Attr}(C) = I(C)/A^2(C), \quad (6)$$

in which $I(C)$ is the moment of inertia and $A(C)$ the area of a component C . This attribute is equivalent to the first moment invariant of Hu [26].

Attribute filters can be applied on sets characterized by some generalized notion of connectivity by replacing Γ_x in (5) with the appropriate connectivity opening, e.g. in the mask-based second-generation case by Γ_x^M from (3). For cases involving the handling of contractions, such filters present a drawback known as *over-segmentation* [21,22]. It has been shown that an attribute opening using a contraction-based connectivity [18] reduces to performing the standard attribute opening on M , unless the criterion has been set such that Γ^A is the identity operator [21]. This is summarized into the following:

Theorem 2. The attribute opening Γ_M^A for a contraction-based connectivity with an increasing, shift invariant criterion A is

$$\Gamma_M^A(X) = \begin{cases} X & \text{if } A(\{x\}) \text{ is true,} \\ \Gamma^A(M) & \text{otherwise.} \end{cases} \quad (7a)$$

$$(7b)$$

where Γ^A is the underlying attribute opening from (5). It is evident that if Γ^A is not the identity operator, then all the singleton sets generated by the connectivity opening of (3) fail the attribute criterion hence filtering X reduces to filtering M instead. Over-segmentation affects any region of X not overlapping with a grain of M and applies equally to attribute thinnings.

2.4. Granulometries and pattern spectra

Attribute filters lend themselves to a commonly used pattern analysis tool, known as *granulometry* [2,6,27–29]. Size granulometries have been described intuitively, as sets of sieves of different grades, each allowing details of certain size classes to pass [9]. More formally:

Definition 3. A binary size granulometry is a set of operators $\{\Gamma_r\}$ with r from some totally ordered set R , with the following three properties:

$$\Gamma_r(X) \subseteq X, \quad (8)$$

$$X \subseteq Y \Rightarrow \Gamma_r(X) \subseteq \Gamma_r(Y), \quad (9)$$

$$\Gamma_r(\Gamma_s(X)) = \Gamma_{\max(r,s)}(X), \quad (10)$$

for all $r, s \in R$.

The first two properties state that Γ_r is anti-extensive and increasing, and the third implies idempotence. This summarizes essentially the definition of a size granulometry to an ordered set of openings, each of which converts an image to a new image in which features smaller than a particular size are absent.

Granulometries have been used to define pattern spectra [1] for image analysis. A pattern spectrum is a binning tool for connected components based on their attributes. The bins may correspond to size or shape classes and each contains an amount of image detail measured as a number of pixels. The image detail is essentially the set of connected components with attribute values that are within the range of each bin.

The pattern spectrum $s_r(X)$ is obtained by applying the size granulometry $\{\Gamma_r\}$ to a binary image X and is defined as

$$(s_r(X))(u) = - \left. \frac{d\zeta(\Gamma_r(X))}{dr} \right|_{r=u}, \quad (11)$$

where ζ denotes the Lebesgue measure in $\mathbb{R}^n$ which is simply the area $A(X)$ for $n=2$.

Shape operators insensitive to size information are also used to define granulometries [9]. This requires omitting the second property of Definition 3 and instead include a condition ensuring scale invariance as follows:

$$\Phi_r(tX) = t(\Phi_r(X)), \quad \forall t > 0, \quad (12)$$

in which tX denotes scaling of set X by t . Thus a shape granulometry consists of operators Φ which are anti-extensive, idempotent and scale invariant. Furthermore, shape pattern spectra can be defined in a way analogous to size pattern spectra [9].

3. Partition-induced connections

3.1. Partitions and connections

The notion of a partition like that of a connection, describes element groupings on E . The formal definition as given in [17] is the following:

Definition 4. Let E be an arbitrary set. A partition $\mathbf{P}$ of E is a mapping $x \rightarrow \mathbf{P}(x)$ from E into $\mathcal{P}(E)$ such that

- (1) for all $x \in E: x \in \mathbf{P}(x)$,
- (2) for all $x, y \in E: \mathbf{P}(x) = \mathbf{P}(y)$ or $\mathbf{P}(x) \cap \mathbf{P}(y) = \emptyset$.

$\mathbf{P}(x)$ is called the class of the partition of origin x . The two conditions indicate that classes $\mathbf{P}(x)$ occupy the whole space E and that two distinct classes have no common point.

Partition classes as opposed to connected components, do not necessarily contain elements from the foreground sets only. Because of this, establishing a relation with a connection requires the use of connectivity openings which naturally separate background from foreground components [3,30].

Definition 5. Given a partition $\mathbf{P}$ of the space E , all the subsets of each class $\mathbf{P}(x), x \in E$, of the partition generate a family conditionally closed under union given by

$$\mathcal{C}^\pi = \{A \cap \mathbf{P}(x), x \in E \text{ and } A \in \mathcal{P}(E)\}. \quad (13)$$

We call $\mathcal{C}^\pi$ a *partition-induced* (pi) or π -connection and the associated operators, π -connectivity openings. It follows that for a set $A \subseteq E$, the connected component given by $\Gamma_x^\pi(A)$ is simply

$$\Gamma_x^\pi(A) = A \cap \mathbf{P}(x). \quad (14)$$

An example is shown in Fig. 2(f). Serra [30] concludes with the following theorem linking the notion of a partition with that of connection.

Theorem 3. Let $\mathcal{C}^\pi$ be a connection on $\mathcal{P}(E)$ associated to the family of connectivity openings $\{\Gamma_x^\pi | x \in E\}$. For each set $A \subseteq E$ the connectivity openings $\{\Gamma_x^\pi | x \in E\}$ subdivide A according to the largest possible partition into members of $\mathcal{C}^\pi$. This operation is increasing in that if $A \subseteq A'$, then any connected component of A is upper-bounded by a connected component of A' .

An example is shown in Fig. 2(c) and (f), where the partition on the original image X controls the connected components of $\mathcal{C}^\pi$.

3.2. Countering over-segmentation with π -connections

The connectivity openings associated to contraction-based or mask-based second-generation connections return singleton sets that account for foreground elements of the original set X that correspond to the background in the connectivity mask M or $\psi(X)$ (where ψ typically an erosion or an opening). Attribute filters based on such connectivity openings yield over-segmented sets as discussed in Section 2.3 and furthermore disregard structural information from objects in regions given by $X \setminus M$. In this section, aided by the concept of partitions, we introduce a connectivity opening that addresses elements in these regions as connected components thus allowing meaningful attributes to be assigned to them and to process them further. This prevents over-segmentation at the operator level independent of the mask image.

Consider a partition of E such that given any arbitrary set $A \subseteq E$, its classes are given by

$$\mathbf{P}_A(x) = \begin{cases} \Gamma_x(A) & \text{if } x \in A, \\ \{x\} & \text{otherwise.} \end{cases} \quad (15a)$$

$$(15b)$$

The proof that $\mathbf{P}_A$ is a valid partition is trivial. Given a mask image M resulting from some operator applied on X , substituting $\mathbf{P}_M(x)$ in (14) yields

$$\begin{aligned} \Gamma_x^\pi(X) &= X \cap \mathbf{P}_M(x) \\ &= \begin{cases} \Gamma_x(M) \cap X & \text{if } x \in X \cap M, \\ \{x\} & \text{if } x \in X \setminus M, \\ \emptyset & \text{otherwise.} \end{cases} \end{aligned} \quad (16a)$$

$$(16b)$$

$$(16c)$$

This is the mask-based connectivity opening discussed in Section 2.2, derived in a far simpler way than in [20]. The objective is to replace the term returning singleton sets with a more specific function to extract components in $X \setminus M$.

Proposition 1. Let $\mathcal{C}$ be a connection of E associated with the family $\{\Gamma_x | x \in E\}$ of connectivity openings. The mapping of $x \rightarrow \mathbf{P}_M^x(x)$ from E

to $\mathcal{P}(E)$ is a partition whose classes are given by

$$\mathbf{P}_M^x(x) = \begin{cases} \mathbf{P}_M(x) & \text{if } x \in M, \\ \mathbf{P}_{X \setminus M}(x) & \text{otherwise.} \end{cases} \quad (17a)$$

$$(17b)$$

Proof. First we show that the classes of $\mathbf{P}_M^x$ cover E , i.e. $\bigcup_{x \in E} \mathbf{P}_M^x(x) = E$.

$$\bigcup_{x \in E} \mathbf{P}_M^x(x) = \left(\bigcup_{x \in M} \mathbf{P}_M^x(x) \right) \cup \left(\bigcup_{x \notin M} \mathbf{P}_M^x(x) \right). \quad (18)$$

The first term is trivial since $\bigcup_{x \in M} \mathbf{P}_M^x(x) = \bigcup_{x \in M} \Gamma_x(M) = M$. For the second term we identify two subcases:

$$\bigcup_{x \notin M} \mathbf{P}_M^x(x) = \left(\bigcup_{x \in X \setminus M} \mathbf{P}_M^x(x) \right) \cup \left(\bigcup_{x \notin X \cup M} \mathbf{P}_M^x(x) \right). \quad (19)$$

Substituting (15) with the appropriate subscript we get

$$\bigcup_{x \in X \setminus M} \Gamma_x(M^c) \cap \Gamma_x(X) = \bigcup_{x \in X \setminus M} \Gamma_x(X \setminus M) = X \setminus M$$

for the first, and

$$\bigcup_{x \notin X \cup M} \Gamma_x(M^c) \cap \{x\} = \bigcup_{x \notin X \cup M} \{x\} = E \setminus (X \cup M)$$

for the second subcase. Summarizing, (18) yields

$$\bigcup_{x \in E} \mathbf{P}_M^x(x) = (X \setminus M) \cup M \cup (E \setminus (X \cup M)) = E. \quad (20)$$

For the last part of the proof we are required to show that the classes of the partition $\mathbf{P}_M^x$ are equal or disjoint; that is for any two points of origin $x, y \in E \Rightarrow \mathbf{P}_M^x(x) = \mathbf{P}_M^x(y)$ or $\mathbf{P}_M^x(x) \cap \mathbf{P}_M^x(y) = \emptyset$. We identify the following three cases:

- (1) if $x, y \in M$ then $\mathbf{P}_M^x(x) = \mathbf{P}_M(x)$ and $\mathbf{P}_M^x(y) = \mathbf{P}_M(y)$ which are either disjoint or equal because $\mathbf{P}_M$ is a partition,
- (2) if $x \in M$ and $y \notin M$ then $\mathbf{P}_M^x(x) \subseteq M$ and $\mathbf{P}_M^x(y) \subseteq (E \setminus M)$, and $\mathbf{P}_M^x(x) \cap \mathbf{P}_M^x(y) = \emptyset$,
- (3) if $x, y \notin M$ then $\mathbf{P}_M^x(x) = \mathbf{P}_{X \setminus M}(x)$ and $\mathbf{P}_M^x(y) = \mathbf{P}_{X \setminus M}(y)$ which are either disjoint or equal because $\mathbf{P}_{X \setminus M}$ is a partition.

Thus, in all cases the we have equal or disjoint sets concluding that $\mathbf{P}_M^x$ is a valid partition. This yields a partition-induced connection in the form of (13) with connectivity openings given by $\Gamma_x^\pi(X) = X \cap \mathbf{P}_M^x(x)$ or more explicitly:

$$\Gamma_x^\pi(X) = \begin{cases} \Gamma_x(M) \cap X & \text{if } x \in X \cap M, \\ \Gamma_x(X \setminus M) & \text{if } x \in X \setminus M, \\ \emptyset & \text{otherwise,} \end{cases} \quad (21a)$$

$$(21b)$$

$$(21c)$$

The classes of $\mathbf{P}_M^x$ can also be set to occupy coarser regions. An example is by setting $\mathbf{P}_M^x(x) = \mathbf{P}_{E \setminus M}(x), \forall x \notin M$ in which the corresponding connectivity opening returns a cluster of all regions in $X \setminus M$. Furthermore, since there are no constraints as to what the partitions should be in relation to the original image or how should they be computed, any set of masks can be used. In the case of a single mask this means that the condition $M \subseteq X$ is not essential. If there are mask regions with higher intensity compared to the original image, then local clusters are extracted according to (21b). For the purposes of the current work, however, we employ the partition $\mathbf{P}_M^x$ as is defined in (17). This means a single mask which is a subset of the original image for regions in $X \cap M$ and a second image computed as $X \setminus M$ for all the remaining structures. More explicitly, in all our experiments we use a mask that is the result of an opening on the original image. $\square$

3.3. π -Connected attribute filters

Attribute filters making use of contraction-based, π -connectivity as opposed to m -connectivity openings, i.e. both parameterized with masks $M \subseteq X$, have the advantage of dealing with meaningful structures in regions given by $X \setminus M$. These structures are usually thin elongated segments like the filamentous protrusions of the binary neuron image of Fig. 3(a). The mask in this example, shown in (b), is created from an anti-extensive opening with a disc structuring element on (a). The aforementioned structures are shown explicitly in (c). A non-compactness filter based on (3) would remove all pixels in these regions, seen as “leaking” paths, unless it is set to be the identity operator. Fig. 3(d) shows the result of this filter for $\lambda = 1.2$ —both the disc-like neuron soma and all singletons are removed. Operating the same attribute thinning, on the connected components given by (21) instead, removes compact structures and allows the extraction of the dendrites from the soma. Fig. 3(e) and (f) shows that for $\lambda = 5$ and 9, respectively. That is, together with preventing the over-segmentation of sets, the connectivity opening in (21) can extract image structures previously inaccessible for further processing.

The effects of over-segmentation on similar examples with gray-scale images are demonstrated in Fig. 1 and in [22]. Thin/small structures that appear at higher gray levels often contribute to the object sharpness thus removing them causes severe blurring and edge distortion. Extending π -connected attribute filters to gray-scale is not trivial and the problem remains to be solved. Despite the limitations discussed next, π -connectivity openings/filters find use in pattern analysis of gray-scale images and provide richer spectra when compared to contraction/mask-based connectivity openings. Note that binary granulometries based on π -connectivity openings or other binary attribute filters relying on them can be trivially defined since the operator properties confirmed in the previous subsection and verified in [22], conform with Definition 3 [9].

3.4. Gray-scale limitations

Connected operators extend to gray-scale quite readily [2,6]. A requirement, however, is that for threshold sets at each gray level either the same connectivity class is used or that the connectivity classes form a connectivity pyramid [20]. For anti-extensive filters

this means that the connectivity class used at level h is a subset of that at any level $h' < h$, which guarantees that any connected component of level h is also a connected set at level h' . This property is violated by the π -connectivity opening that we propose. An example is shown in the schematic of Fig. 4 where given a gray-scale image f shown in (a), and a connectivity mask m shown in (b) such that $m < f$, the operator handling regions in $T_h(f) \setminus T_h(m)$ (shown in (c)) extracts connected components which are not nested along the intensity range H .

The bottom row of Fig. 4 shows three images (d)–(f); the decompositions of f from h_0 to h_2 under standard, contraction-based m - and π -connectivity, respectively. The last case shows clearly that although the squares which are the non-singleton components according to m -connectivity, are nested appropriately, the bridging regions of (c) which are missing from the connectivity mask, violate this nesting property and as such

$$\Gamma_X(T_{h_2}(f)) \setminus \Gamma_X(T_{h_2}(m)) \not\subseteq \Gamma_X(T_{h_1}(f)) \setminus \Gamma_X(T_{h_1}(m)) \quad (22)$$

for $h_1 < h_2$. The shadowed areas in the middle plane point the two regions which generate this nesting conflict. This lack of nesting prevents building *Max-Trees* [6] of connected components of threshold sets at all gray levels. Instead, a more general graph, allowing multiple parents would be needed. More on *Max-Trees* and their significance in image filtering and pattern-spectrum computation are discussed in the next section.

4. Gray-scale pattern analysis

4.1. Gray-scale pattern spectra using *Max-Trees*

Existing methods for computing connected pattern spectra for a gray-scale images require that the corresponding granulometries extend to gray-scale [2,9,29]. Under this condition, the gray-scale pattern spectrum is given by replacing the Lebesgue measure with the integral of f (sum of the gray levels) over the image domain. In the discrete case, like with binary images, computing $s_{\gamma_r}(f)$ requires a repetitive filtering by each γ_r , in ascending order of r . At each filter step the sum of gray-levels s_r of the resulting image is computed and the pattern spectrum value at r is given by subtracting s_r from s_{r-} , with $r-$ the scale immediately preceding r . In the case of π -connectivity, however, the lack of a direct gray-scale extension prevents repetitive

Fig. 3. Attribute thinning of neurons—top row: original image X (a); the mask image M (b); the image structures in $X \setminus M$ (c). Bottom row—the results of the non-compactness filter in M (d), in $X \setminus M$ with $\lambda = 5$ (e) and with $\lambda = 9$ (f).

Fig. 4. Gray-scale decompositions—top row: the original 3-level image f (a); the mask image m by a structural opening with a square SE (b); the superposition of sets in $T_h(f) \setminus T_h(m)$ for all h (c). Bottom row: the decomposition of (a) based on standard connectivity (d), on contraction-based m -connectivity (e) and on contraction-based π -connectivity (f). The hatched regions of $T_h(f)$ on the middle level of (f) demonstrate why threshold superimposition is not possible; the decomposition is not increasing.

filtering. To compute a gray-scale (pseudo) pattern spectrum based on π -connected operators we look into methods that do not require filtering. Urbach et al. used such methods [9] based on connected operators and Max-Trees.

The Max-Tree [6] is a rooted, unidirected tree in which the node hierarchy corresponds to the nesting of peak components given a gray-scale image. A peak component P_h at level h is a connected component of the thresholded image $T_h(f)$ while a flat-zone [5] at level h is a set containing all the pixels of a peak component which are at level h in f . Each tree node C_h^k (k is the node index) contains the sum of the pixels found in all the flat-zones of a given peak component at level h . In addition each node except for the root, points towards its parent $C_{h'}^{k'}$ with $h' < h$. The root node is defined at the minimum level $h_{\min}$ and contains the set of pixels belonging to the background. The algorithm, used primarily for anti-extensive attribute filtering, runs a three-stage process in which the construction of the tree and the computation of node attributes is independent of filtering and image restitution. During the construction stage every pixel visited contributes to the auxiliary data buffer associated to the node it belongs to. Once a node is finalized, its parent inherits these data and recomputes its attribute. Inheritance in the case of attributes such as area/volume is a simple addition while for more complicated attributes such as the non-compactness measure of (6) the accumulation relies on more sophisticated attribute handling functions described in [9,20]. An example of a simple Max-Tree structure is shown in Fig. 5.

Computing the pattern spectrum using Max-Trees becomes essentially an accumulation procedure. The method scans the tree structure by visiting all nodes from $h_{\min}$ to $h_{\max}$ and retrieves the attribute measures of the corresponding peak components. Using some binning function (see next subsection) this measure is used

to place the corresponding peak component to the appropriate spectral entry. The contribution of each peak component is given by the product of its area with the gray-scale difference from its parent. Each peak component belonging to a given class updates the class *energy counter* by accumulating its product to the existing value. We conclude Urbach's method to the following statement:

$$(S_{\gamma_r}(f))(u) = \sum_{h=h_{\min}}^{h_{\max}} \sum_{\substack{k: C_h^k \neq \emptyset \wedge \\ \text{Bin}(P_h^k) = u}} A(P_h^k) \Delta h^k, \quad (23)$$

where Δh^k is the gray-scale difference between the k th node at level h and its parent, and Bin the binning function. If we wish to compute the same sums on a level basis instead of using the Max-Tree dynamics given by the term Δh^k , the same expression reduces to

$$(S_{\gamma_r}(f))(u) = \sum_{h=h_{\min}}^{h_{\max}} \sum_{k: \text{Bin}(P_h^k) = u} A(P_h^k). \quad (24)$$

That is, for every level accumulate the area of all peak components whose attributes fall within the bounds of class u . Since π -connected operators are limited to binary sets only, using this formula we can compute a maximum of $H - 1$ Max-Trees, one for each binary image from the threshold decomposition of f . This is for structures in $T_h(f) \setminus T_h(m)$ since the spectrum entries for stable components in $T_h(m)$ are computed using (23). For each threshold set at level h (24) becomes

$$(S_{\gamma_r}(T_h(f)))(u) = \sum_{k: \text{Bin}(P_h^k) = u} A(P_h^k). \quad (25)$$

Fig. 5. An example of a Max-Tree representation of the simple four-level image on the left. The two peak components at h_1 and h_2 of the left root-path, contain one flat-zone each which is the outer ring in each case.

4.2. Binned 2D shape-size spectra

Multi-dimensional spectra have been used before to sort connected components based on several attribute measures [9]. For the purposes of this work we consider a joint 2D shape-size pattern spectrum that features the non-compactness attribute of (6). The method we present for its computation relies on the Max-Tree structure whose corresponding connected components are computed based on the contraction-based π -connectivity opening of (3). The procedure is summarized in the following:

Algorithm 1. Computing 2D shape-size binned pattern spectrum s using the Max-Tree for a contraction-based π -connectivity with N_a shape and N_b size classes.

- (1) Mask generation: Compute the opening transform of the input image for a given SE.
- (2) Stable components: Compute the Max-Tree of the gray-scale mask image.
- (3) Auxiliary data: As the Max-Tree is built, compute the area $A(P_h^k)$ and the moments of inertia $I(P_h^k)$ corresponding to each peak component P_h^k .
- (4) Spectrum initialization: Set the $N_a \times N_b$ elements of the spectrum array S to zero.
- (5) Spectrum update: For each node C_h^k :
 - Compute the size class r from the area $A(P_h^k)$.
 - Compute the shape class s from $I(P_h^k)/A^2(P_h^k)$.
 - Compute the gray-level difference Δh^k between the current node and its parent.
 - Add the product of Δh^k and $A(P_h^k)$ to $S[r, s]$.
- (6) Non-stable components: For all gray levels threshold both f and m and compute the binary mask-complement image. For each binary image:
 - Compute the binary Max-Tree.
 - Compute the auxiliary data and size/shape classes as above.
 - For each node of the tree update the spectrum as above.

The term *stable components* in the above algorithm refers to connected components in regions within the mask. By contrast, *unstable components* are those in regions $T_h(f) \setminus T_h(m)$.

In this algorithm if we were to compute a pattern spectrum based on contraction-based m -connectivity openings, there are

certain simplifications which can boost its performance. Since m -connected operators are used to construct gray-scale granulometries the need to threshold f and m would no longer exist. In fact, using the Dual-Input Max-Tree algorithm from [20] in step 2 would be sufficient for skipping step 6. In practice, however, since singleton sets, just like noise, do not contribute particularly valuable information, it is usually sufficient to compute the Max-Tree of the mask image only.

Connected components are mapped into bins or classes based on their attribute values. Each bin covers a range of values which is specified during the spectrum initialization. The mapper is a heuristic function and for the needs of our experiments we used the one presented in [9]. According to this, a component with an attribute value v is mapped into a class c given by

$$c = \left\lfloor \frac{\log_2(v) - \log_2(D_0)}{\log_2(D_1) - \log_2(D_0)} N_c \right\rfloor, \quad (26)$$

where $\lfloor \cdot \rfloor$ denotes the floor function, N_c the number of classes, and D_0 and D_1 the lower and upper bounds of the range of interest of the attribute values.

5. Examples of π -connected feature extraction

The three different types of connectivity discussed in this paper, the standard, mask-based second-generation and π -connectivity, partition any given image in ways that can either be totally different to each other even if using the same mask, or equivalent if the mask image equals the original. Using a mask image that is a subset of the original, i.e. dealing with a contraction-based m -connectivity, was shown to be responsible for the over-segmentation problem due to the lack of control in generating singleton sets. The π -connectivity opening of (21) prevents this, as was discussed in Section 3.2, and moreover allows the explicit handling of subregions of connected components according to standard connectivity, that were previously inaccessible.

The first two experiments that follow demonstrate this feature by extracting thin elongated structures that account for object regions outside the mask. Computing the π vs. the standard connected pattern spectrum and displaying the contents of selected bins, shows that in the case of π -connectivity a whole new set of features can be computed that allows for more detailed measurements on binary medical images. These features are mostly measured manually or simply estimated in everyday medical practice. An example is the

length of filaments or vessel segments, for which the Euclidean distance between two manually selected points is often seen as the most reliable estimate. The extracted structures are highlighted (white) and superimposed over the darkened binary counterpart of the original image, for better localization.

The third experiment compares π vs. m -connected pattern spectra to show how the over-segmentation associated to connected operators is suppressed. In all three cases we compute a 2D pattern spectrum that is 15×15 in size, with the x -axis accounting for values of the non-compactness attribute and the y -axis for values of the size attribute. Binning is done using the logarithmic mapper of (26).

The first case shows an axial CT slice of a patient's head, in which a large aneurysm is present, Fig. 6(a). The study of an aneurysm prior to its surgical treatment aims at estimating its size and the size of the vessel on which it appears, and at checking whether it is located at a vessel junction or not.

Computing these information requires that the aneurysm is handled separately from the vessels and other structures, i.e. as is the case with a π -connected pattern spectrum (Fig. 6(c)). To show the practical differences with a standard-connected pattern spectrum (Fig. 6(c)) we run a comparison in which we aim at three different types of structures; the skull, the aneurysm and the vessels. Images (d)–(f), in the middle row of Fig. 6, and (g)–(i) in the bottom row, show these structures by selecting one or more bins from the standard and π -connected pattern spectrum, respectively.

The second case, shown in Fig. 7(a), is of a patient suffering from pulmonary fibrosis. The evolution and/or treatment progress of this disease is evaluated from the inspection of regularly acquired CT scans. The procedure involves the measurement of certain key attributes on the filamentous tissue running across the pulmonary walls. The features of most interest are the length and width of the various segments. The use of pattern spectra

Fig. 6. CT scan showing a cranial aneurysm. Top row: original image (a); pattern spectrum of (a) based on standard connectivity (b) and π -connectivity (c). Middle row—features based on standard connectivity: bin(6,12) (d), bin(11,14) (e) and bin(14,10) (f). Bottom row—features based on π -connectivity: bins(6,13), (4,10), (7,11), (9,13) (g), bin(0,11) (h) and bins(11,8), (12,9), (13,9), (14,11) and (14,9) (i).

Fig. 7. Pulmonary fibrosis. Original image (a); the pattern spectrum of (a) based on standard (b) and π connectivity (c), respectively. Middle row—features based on standard connectivity: bin(7, 14) (d), bin(11, 8) (e) and bin(14, 7) (f). Bottom row—features based on π -connectivity: bin(7, 14) (g), bin(3, 8), (5, 6)(6, 7)(8, 5) (h) and bins(9, 7)(10, 8)(11, 6)(11, 7)(12, 6)(13, 5)(14, 7) (i).

configured with standard connectivity (Fig. 7(b)) yields information of little use since most of these filamentous segments are interconnected between them and to the main body. The middle row of Fig. 7(d)–(f) shows three dominant connected components containing much of the image energy. It can be seen that a large fraction of these segments are extracted as a single object in (d). By contrast, employing a π -connection with a mask given by an anti-extensive opening on the original image, allows the fragmentation of this tissue into smaller individual segments that give a richer pattern spectrum as shown in Fig. 7(c). The bottom row of Fig. 7 shows the main body handled separately in (g) and components from various bin groupings (h)–(i).

The last experiment shows the effects of three different types of connectivity on the computation of connected pattern spectra. The pattern spectrum of a gray-scale image of a diatom (more on diatoms in Section 6) is computed based on standard, contraction-based second-generation m -connectivity, and based on π -connectivity. For the last two cases we use the same mask image m which is the result of an anti-extensive opening on the original. The redistribution of image power as a function of the SE radius is monitored for four different radii. Images in Fig. 8(a) and (b) show the original image and the structures corresponding to the bin(3, 14) of the pattern spectrum based on standard connectivity in (c).

The top row of Fig. 9 shows the pattern spectra based on m -connectivity for masks computed with SE radii from 3 to 9, incremented by 2 in each case. The bottom row shows the equivalent pattern spectra based on π -connectivity.

This experiment shows clearly the effect of over-segmentation in m -connected pattern spectra where the very first bin in each case (bin(0, 0)—accounting mostly for singletons and noise) is saturated while texture information is poorly represented. As the radius of the SE increases the spectrum becomes poorer since more structures are converted to sets of singletons, contributing to the very first bin alone. By contrast, in the case of π -connected spectra, we observe that bin(0, 0) is approximately equivalent in contribution to that in the standard connected spectrum. Moreover, as the radius increases, more structures are fragmented and the spectrum becomes richer.

6. Diatom identification experiments

The experiments described in this section aim at highlighting the significance of structures discarded by contraction-based m -connectivity openings. We chose a diatom image classification problem for this purpose using the π -connected pattern spectra of the image set as feature vectors. Similar experiments were conducted in the past using pattern spectra based on standard connected operators [31]. We follow similar procedures to allow comparisons between the two methods and report on the overall classification performance.

6.1. The ADIAC diatom image database

Diatoms are a large and ecologically important group of unicellular or colonial algae which are found in almost all aquatic

Fig. 8. Pattern spectrum of a diatom image based on standard connectivity. The original image (a); the structures corresponding to the bin(3, 14) (b) and the pattern spectrum (c). In the pattern spectrum the origin is in the top, left hand corner, the logarithm of the area runs along the vertical axis, the logarithm of the non-compactness along the horizontal, both according to (26).

Fig. 9. Diatom image pattern spectra, using the same binning as in Fig. 8(c). Top row shows the pattern spectrum based on m -connectivity for four different mask images described in the text. The bottom row shows the corresponding π -connected pattern spectra using the same masks. Note how in the top row image content is shifted to the bin at the origin, which corresponds to singletons.

habitats. Their silica cell walls consist of two halves called *valves* and together with the pattern of pores (internal valve texture also called *ornamentation*) and other valve markings, provide the information needed for species or *taxa* identification.

The experiments that follow make use of two sets of diatom images obtained from the publicly available ADIAC database which can be found at <http://www.ualg.pt/adiac/pubdat/pubdat.html>. The first set referred to as *mixed genera* consists of 781 images representing 37 distinct taxa, and the second, the *Sellaphora pupula*, of 120 images from six different subspecies of *Sellaphora*. For both sets each *taxa* or subspecies is represented by at least 20 images. Moreover, for each of the 8-bit gray-scale images a contour file is given to mask out regions outside the diatom valve. Acquisition and preprocessing methods as well as image features and other details are available in [31].

6.2. Experimental methods and parametrization

For each of the two image data-sets we replicate the experimental procedure followed by Urbach et al. [9] only instead of computing connected pattern spectra based on the standard connectivity, we experiment with m and π -connectivities. We compute the respective spectra for connectivity openings with five different sizes of circular structuring elements starting from radius 3 until 11 incremented each time by 2 pixels. Prior to each run, we compute the bin extrema D_0 and D_1 for each of the two classes as in (26). The values are obtained by a scan through the entire image database. The entries of the size class are scaled with the pixel width associated with each image and so are the extrema.

In each experiment we produce a feature vector of 600 elements. The first 300 correspond to processes on the original image and the remaining 300 on the inverted copy. This is done to capture information from both bright and dark patterns in the images. Each set of 300 elements corresponds to a 2D matrix

mapped in a lexicographic order to a 1D vector. For each matrix, x -dim. always refers to the non-compactness attribute and y -dim. to the size attribute. The 600-long vector is a concatenation of two such vectors and is complemented by some additional information to meet the classifier's input specifications.

6.3. The C4.5 decision tree classifier

To carry out meaningful comparisons we employed the same decision tree classifier built with the C4.5 algorithm as in [9,31]. To compensate for classifier instabilities using a single decision tree, we use *bootstrap aggregation* or *bagging* with the same procedure reported in [9]. Briefly this can be summarized into the following. Firstly, for each image database we divide the total number of images into two subsets, the training and the test set. The latter one contains roughly 25% of the total images per class. To generate the decision tree forest, we select randomly a number of images from the training set which we group into 25 smaller subsets. A single decision tree is built for each set separately, a process which is known as bootstrapping. An accuracy measure described in [9] is then used to evaluate each of the decision trees followed by a majority vote on their outcome (aggregation). We repeat this procedure over 10 times on newly created training sets and obtain the overall classification performance by averaging the individual outputs. An error estimate is computed using cross-validation.

6.4. Experiments

The first experiment trains the classifier using π -connected pattern spectra. We run the same experiment with six different spectral arrangements. That is, we used three sets of extrema and two different spectral layouts, i.e. 15×20 and 20×15 . The sets of extrema are the ones provided by Urbach et al., the absolute

extrema computed by our scan routine and the extended extrema which are the same as the absolute only excluding small particles by multiplying D_0 of the size class by 3. Table 1 lists Urbach's extrema which are given for the *mixed genera* data-set only and the set of absolute extrema that we have computed. Note that in Urbach's experiments D_0 for the size class is also multiplied by 3 to reduce the effects of noise.

The highest performance in all scales for both image data-sets was achieved using the extended extrema, i.e., excluding singletons which over-flood the first bin of the spectrum. To avoid this in the second experiment where we train the classifier with m -connected pattern spectra, we compute ordinary Max-Trees from the gray-scale connectivity masks associated to each input image. This essentially discards the middle term of (3) and resolves the over-flood issue. The classification performances for both image data-sets together with the error estimates are listed in Tables 2 and 3. We mark with bold numbers the best result in each case. We observe that the best classification performance is given for a 20×15 spectral layout in both data-sets and for relatively small radii of the structuring elements used.

Table 1
Pattern spectrum extrema for both image data-sets.

SE radius	Mixed genera			
	Size		Shape	
	D_0	D_1	D_0	D_1
3–9	0.000983	2050.0	1.0	446.68
11	0.000983	2050.0	1.0	395.907
Urbach	0.003	7198	1.0	328.1
<i>Sellaphora pupula</i>				
3–11	0.0025	325.731	1.0	153.963

Table 2
Classification performance for the mixed genera data-set.

SE radius	π -connectivity				SE radius	m -connectivity			
	20×15		15×20			20×15		15×20	
	μ (%)	σ	μ (%)	σ		μ (%)	σ	μ (%)	σ
3	88.47	3.79	89.20	4.47	3	88.37	3.82	88.27	3.63
5	90.70	2.71	87.70	3.10	5	86.86	3.16	86.16	4.88
7	90.64	5.02	90.05	4.52	7	84.75	6.21	85.08	6.68
9	89.89	3.82	89.35	3.10	9	82.38	5.18	84.00	5.21
11	89.45	4.41	89.79	3.47	11	81.02	4.48	82.64	7.62

By comparison: 4-connectivity used by Urbach et al. yields $91.1 \pm 1.6\%$ performance using 15×20 binning.

Table 3
Classification performance for the *Sellaphora pupula* data-set.

SE radius	π -connectivity				SE radius	m -connectivity			
	20×15		15×20			20×15		15×20	
	μ (%)	σ	μ (%)	σ		μ (%)	σ	μ (%)	σ
3	83.00	1.37	76.33	1.97	3	75.00	2.61	71.66	1.91
5	78.66	1.28	75.00	1.80	5	69.66	1.57	70.00	2.68
7	77.00	2.16	79.33	2.03	7	66.33	2.11	69.66	2.07
9	73.66	2.10	71.00	1.34	9	69.66	2.07	66.00	3.34
11	69.33	1.72	75.33	1.85	11	63.33	2.23	69.33	2.13

By comparison: 4-connectivity used by Urbach et al. yields $78.00 \pm 2.14\%$ performance using 15×20 binning.

6.5. Performance optimization using combined methods

The second of the two morphological based methods for diatom feature extraction reported in [31] uses contour analysis by morphological curvature scale spaces [32–34]. We use this in combination with the π -connected pattern spectra presented in this paper and the method of Urbach et al. [9] to optimize the classification performance in the case of the mixed genera data-set.

Urbach's method which processes comparatively larger structures on the diatom valves achieves a classification accuracy of 91.1% in a 15×20 spectral layout. Using the π -connected pattern spectrum alone, we fail to raise this figure further while a small improvement appears (91.46%) if concatenating the two vectors, the original by Urbach and ours, into a 1200-long new one. The increase is limited most probably due to feature correlations which degrade the classifier's stability (error estimate is 5.49%). To reduce this, we create for each image a 600-long vector made of the average values between the respective members of the five feature vectors (one for each scale). We subtract each new member from the corresponding member in Urbach's feature vector and concatenate the resulting vector with Urbach's original. Note that in order to compute an average vector, the binning and spectral layout for each scale must be the same. As such we do not use our optimal setup but instead the results obtained using Urbach's spectral extrema for a 15×20 layout. Using this new set of feature vectors (referred to as *combined spectra*) the classifier achieves a prediction accuracy of 92.92% with a slightly reduced error estimate of 3.61%.

A look through the individual connected components associated to different spectral bins reveals that contour structures are poorly represented in both pattern spectrum-based methods. To account for fragmentations and incomplete boundaries we employ the method of Jalba et al. [33,34]. For each image this contour-based method yields a 66-long vector. Concatenating our vector of the combined spectra with Jalba's we raise the classifier's prediction rate to 95.18% with a considerably lower

error estimate of 1.92%. Furthermore, if instead of the average vector in the combined spectra we use the best performing vector with Urbach's extrema and in a 15×20 layout and concatenate this with Jalba's vector this figure raises slightly further to 95.78% at the expense of an increased σ of 2.52%. This small increase in the error estimate is expected since in the first case the average vector ensures better stability. The significance of contour information can also be seen if combined with each of the two spectral-based methods separately. A summary of these results and the best classification rates achieved using combinations of other not-necessarily morphological methods is given in Table 4.

7. Discussion of results

In the first set of experiments we train the classifier with π -connected pattern spectra. We see that in the case of the mixed genera data-set, we obtain a rather stable performance throughout the five scales and for both types of spectral layouts. This is due to differences in ornamentation between diatom species which result in fragmentations at different scales. Having a limited subset of diatoms affected by this operation at a given scale, the method fails to contribute in improving further the classification accuracy. The success rate of the classifier, however, compares with that obtained using Urbach's method.

By contrast, in the case of the *Sellaphora pupula* data-set where different subspecies differ little in ornamentation, we see major

variations in classification accuracy as the scale changes. Notably, the larger the radius of the structuring element used further the drop in success rate. This suggests that for the specific species the fragmentation which occurs at the first scale separates thin elongated features from larger structures providing a more accurate characterization of the ornamentation. This when compared with Urbach's method run on the entire data-set as opposed to a limited subset reported in [31], yields a gain of 5% in success rate.

Both data-sets show preference in higher bin resolution for the non-compactness attribute. This suggests that the fragmentation of the diatom ornamentation contributes more to shape rather than to size information.

The second set of experiments makes use of m -connected pattern spectra. The observation of classifier instabilities when including large numbers of singletons at the very first bin of the 2D spectrum led to their exclusion. This essentially reduces to computing the standard connected pattern spectrum on the contracted connectivity masks. In both data-sets we see a progressive decline in classification success as the size of the structuring element increases. This is expected due to the incremental loss of information. It is remarkable though that even for considerably large structuring elements (diameter of 23 pixels) this spectrum based method remains robust and yields a classification success which outperforms many of the other methods reported in [31]. For small values of SE radius much of the noise in the mask which is computed with a structural opening on the original image, is suppressed. As such the classification performance remains high and the feature extraction process essentially becomes equivalent to Urbach's method running on smoothed images.

The graphs in Fig. 10 illustrate how the difference between classification success rates using π vs. m -connected spectra changes over scale for each data-set.

As can be seen in Table 2, in the case of the mixed genera, the very small difference recorded in the first scale for the 20×15 layout suggests that the influence of ornament fragmentation there is minimal. Since in all methods we discard singletons and small objects with area up to 3 pixels, most of the structures in each $X \setminus M$ must be within this size range. The classification difference which appears as an increasing function of scale is upper bounded by a SE radius value above which each $M = \emptyset$ and consequently $X \setminus M = X$. For the *Sellaphora pupula* set this function is not increasing since for certain SE radii there can be common features between the subspecies that when removed or detached

Table 4

Classification performances on the mixed genera data-set using combinations of feature extraction methods.

Methods	Clas. performance		
	μ (%)	σ	Vector size
π -Conn. pattern sp. alone	90.7	2.71	600
Std. conn. pattern sp. alone	91.1	1.6	600
Std. pattern sp. and π -conn. pattern sp.	91.46	5.49	1200
Combined spectra	92.92	3.61	1200
Std. pattern sp. and contours	93.94	3.5	666
π -Conn. pattern sp. and contours	94.05	3.12	666
Combined spectra and contours	95.78	2.52	1266
All methods from [31]	96.9	1.2	329
Same with robust features only	95.5	1.5	17

The term *and* above indicates concatenation of vectors.

Fig. 10. Difference in classification performance when training the classifier with π and m -connected pattern spectra (left).

from the remaining ornamentation provide a set of more distinctive descriptors to the classifier. We see such an example for the SE radius 9 in the right graph of Table 3. The function, however, is upper bounded in the same way as with the first case. Concluding on this comparison, we see that the connected pattern spectra method using contracted masks, when based on π -connectivity outperforms the equivalent based on m -connectivity under all types of spectral configurations. This holds for both data-sets.

The last set of experiments targets the performance optimization using a combination of the two spectral-based methods together with features obtained by contour analysis using morphological curvature scale spaces [33,34]. Using the two spectral methods in a way described in the previous section to minimize the feature correlations yields a small performance gain of 1.82%. Note that a number of other methods were tried such as multi-scale and weighted multi-scale sums but none of them succeeded in overcoming Urbach's result of 91.1%. This suggests that texture based information, although considered the best feature descriptors from the comparison in [31], can reach a certain limit in multi-species classification success beyond which further features and of different nature are required. The morphological method of Jalba et al. [34] focuses on contour information instead of the diatom ornamentation and when used separately it reaches a classification success rate of up to 91.3% with $\sigma = 5.0\%$. The contour descriptors complement the combined spectrum-based method and as such reduce the error estimate while boosting the overall performance to 95.2%. This is comparable to the best performance reported in [31] by using all 17 methods which were available for this purpose.

8. Conclusions

In this paper, starting from Serra's work on image partitions [30] we have presented a new type of connection, the π -connectivity class aided by connectivity masks, which can be used in ways analogous to second-generation connectivity. The steps we use in our proof for establishing the π -connection provide an alternative way to prove the validity of the mask-based connectivity [20] and are applicable in establishing other types of connections trivially.

The strength of π -connected operators is in contraction based problems where the handling of path-wise connected regions otherwise treated as singletons, allows the assignment of meaningful attributes and thus further processing. This in part resolves the problem of over-segmentation [22,21], but due to limitations in extending π -connected operators to gray scale, developing efficient attribute filters remains a topic for further research. The same limitations prevent the introduction of gray-scale granulometries and therefore connected pattern spectra defined in the conventional way. Using the method from [9] we have introduced pseudo-pattern spectra and showed that these can be adopted trivially to compute a gray-scale spectrum based on π -connected operators. A brute-force algorithm is also given.

Classification experiments on two diatom image data-sets showed that the use of pattern spectra associated to contraction-based π -connected operators as feature vectors outperforms their counterparts associated to contraction-based m -connected operators. Comparisons were also made with spectra associated to standard connected operators. The results in the case of the *Sellaphora pupula* data-set indicate that the fragmentation of ornament structures enhances the differentiation between subspecies of the same family and yields higher classification success rate.

Comparing the classification performance of this method on *Sellaphora pupula* with other methods reported in literature we

achieve a similar rate to the best reported which again uses pattern spectra only based on the standard connectivity. This suggests that the spectral methods alone are limited. Combining the two methods and adding contour descriptors, however, yields a success rate comparable to the one based on all methods combined (reported in [31]). The obvious advantages in this case is the far smaller number of methods needed to reach this rate and not having the need of manually selecting the best performing features.

In future work we expect to increase these figures further by using different classifiers while resolving further feature correlations that can reduce the size of the feature vectors used. In addition, further work can be done in deriving appropriate filtering rules to extend the π -connected operators to gray-scale directly and thus implement more efficient algorithms for both filtering and pattern spectra.

Acknowledgement

This work is funded by the Netherlands Organization for Scientific Research under project number 612.065.202.

References

- [1] P. Maragos, Pattern spectrum and multiscale shape representation, IEEE Trans. Pattern Anal. Mach. Intell. 11 (7) (1989) 701–715.
- [2] E.J. Breen, R. Jones, Attribute openings, thinnings and granulometries, Comput. Vis. Image Understanding 64 (3) (1996) 377–389.
- [3] J. Serra (Ed.), Image Analysis and Mathematical Morphology. II: Theoretical Advances. Academic Press, London, 1988.
- [4] H.J.A.M. Heijmans, Connected morphological operators for binary images, Comput. Vis. Image Understanding 73 (1) (1999) 99–120.
- [5] P. Salembier, J. Serra, Flat zones filtering, connected operators, and filters by reconstruction, IEEE Trans. Image Process. 4 (8) (1995) 1153–1160.
- [6] P. Salembier, A. Oliveras, L. Garrido, Anti-extensive connected operators for image and sequence processing, IEEE Trans. Image Process. 7 (4) (1998) 555–570.
- [7] A. Meijster, M.H.F. Wilkinson, A comparison of algorithms for connected set openings and closings, IEEE Trans. Pattern Anal. Mach. Intell. 24 (4) (2002) 484–494.
- [8] A. Sofou, C. Tzafestas, P. Maragos, Segmentation of soil section images using connected operators, in: International Conference on Image Processing 2001, vol. 3, 2001, pp. 1087–1090.
- [9] E.R. Urbach, J.B.T.M. Roerdink, M.H.F. Wilkinson, Connected shape-size pattern spectra for rotation and scale-invariant classification of gray-scale images, IEEE Trans. Pattern Anal. Mach. Intell. 29 (2) (2007) 272–285.
- [10] C. Ronse, Set-theoretical algebraic approaches to connectivity in continuous or digital spaces, J. Math. Imaging Vision 8 (1) (1998) 41–58.
- [11] M.H.F. Wilkinson, M.A. Westenberg, Shape preserving filament enhancement filtering, in: W.J. Niessen, M.A. Viergever (Eds.), Proceedings of the MICCAI'2001, Lecture Notes in Computer Science, vol. 2208, Springer, Berlin, 2001, pp. 770–777.
- [12] M.A. Westenberg, J.B.T.M. Roerdink, M.H.F. Wilkinson, Volumetric attribute filtering and interactive visualization using the max-tree representation, IEEE Trans. Image Process. 16 (12) (2007) 2943–2952.
- [13] E.R. Urbach, M.H.F. Wilkinson, Shape-only granulometries and grey-scale shape filters, in: H. Talbot, R. Beare (Eds.), Mathematical Morphology: Proceedings of the 6th International Symposium on Mathematical Morphology, CSIRO Publishing, Collingwood, Australia, 2002, pp. 305–314.
- [14] T.Y. Kong, A. Rosenfeld, Digital topology: introduction and survey, Comput. Vision Graph. Image Process. 48 (3) (1989) 357–393.
- [15] G. Matheron, J. Serra, Strong filters and connectivity, in: J. Serra (Ed.), Image Analysis and Mathematical Morphology, vol. 2, Academic Press, London, 1988, pp. 141–157.
- [16] J. Serra, Connectivity on complete lattices, J. Math. Imaging Vision 9 (3) (1998) 231–251.
- [17] J. Serra, Connections for sets and functions, Fund. Inf. 41 (1–2) (2000) 147–186.
- [18] U. Braga-Neto, J. Goutsias, Connectivity on complete lattices: new results, Comput. Vis. Image Understanding 85 (1) (2002) 22–53.
- [19] U. Braga-Neto, J. Goutsias, Grayscale level connectivity: theory and applications, IEEE Trans. Image Process. 13 (12) (2004) 1567–1580.
- [20] G.K. Ouzounis, M.H.F. Wilkinson, Mask-based second-generation connectivity and attribute filters, IEEE Trans. Pattern Anal. Mach. Intell. 29 (6) (2007) 990–1004.
- [21] M.H.F. Wilkinson, Attribute-space connectivity and connected filters, Image Vis. Comput. 25 (4) (2007) 426–435.
- [22] G.K. Ouzounis, M.H.F. Wilkinson, Countering oversegmentation in partitioning-based connectivities, in: Proceedings of the International Conference on Image Processing, 2005, pp. 844–847.

- [23] H.J.A.M. Heijmans, Composing morphological filters, *IEEE Trans. Image Process.* 6 (5) (1997) 713–723.
- [24] L. Vincent, Morphological area openings and closings for greyscale images, in: *Proceedings of the NATO Shape in Picture Workshop*, Driebergen, The Netherlands, 1992, pp. 197–208.
- [25] F. Cheng, A.N. Venetsanopoulos, An adaptive morphological filter for image processing, *IEEE Trans. Image Process.* 1 (4) (1992) 533–539.
- [26] M.K. Hu, Visual pattern recognition by moment invariants, *IRE Trans. Inf. Theory* IT-8 (1962) 179–187.
- [27] G. Matheron, *Random Sets and Integral Geometry*, Wiley, New York, 1975.
- [28] J. Serra, *Image Analysis and Mathematical Morphology*, Academic Press, New York, 1982.
- [29] L. Vincent, Granulometries and opening trees, *Fund. Inf.* 41 (1–2) (2000) 57–90.
- [30] J. Serra, A lattice approach to image segmentation, *J. Math. Imaging Vision* 24 (1) (2006) 83–130.
- [31] J.M.H.D. Buft, M.M. Bayer (Eds.), *Automatic Diatom Identification*, Series in Machine Perception and Artificial Intelligence, vol. 51, World Scientific Publishing Co., Singapore, 2002.
- [32] A.C. Jalba, J.B.T.M. Roerdink, M.H.F. Wilkinson, Morphological hat-transform scale spaces and their use in pattern classification, *Pattern Recognition* 37 (5) (2004) 901–915.
- [33] A.C. Jalba, M.H.F. Wilkinson, J.B.T.M. Roerdink, Automatic diatom identification using contour analysis by morphological curvature scale spaces, *Mach. Vision Appl.* 16 (4) (2005) 217–228.
- [34] A.C. Jalba, M.H.F. Wilkinson, J.B.T.M. Roerdink, Shape representation and recognition through morphological curvature scale spaces, *IEEE Trans. Image Process.* 15 (2) (2006) 331–341.

About the Author—GEORGIOS K. OUZOUNIS obtained an M.Phil. degree in computer science from the Department of Computing and Electrical Engineering, Heriot Watt University, Edinburgh, Scotland, in 2001. He received his Ph.D. on image analysis and connected morphological filters from the Institute of Mathematics and Computing Science, University of Groningen, The Netherlands, in 2009. He is now heading the research on computer vision techniques for hepatic, pancreatic and biliary tract surgery supervision at the Department of Experimental Surgery, University Hospital of Alexandroupoli, Democritus University of Thrace, Greece. His research interests involve attribute filters and segmentation operators for 3D medical image analysis, real-time decision making systems and pattern analysis. He is currently working on very high resolution satellite imaging at the Global Security And Crisis Management Unit, IPSC - Joint Research Centre - European Commission.

About the Author—MICHAEL WILKINSON obtained an M.Sc. in astronomy from the Kapteyn Laboratory, University of Groningen (RUG) in 1993, after which he worked on image analysis of intestinal bacteria at the Department of Medical Microbiology, RUG. This work formed the basis of his Ph.D. at the Institute of Mathematics and Computing Science (Now Johann Bernoulli Institute, JBI), RUG, in 1995. He was appointed as researcher at the Centre for High Performance Computing (also RUG) working on simulating the intestinal microbial ecosystem on parallel computers. During that time he edited the book “Digital Image Analysis of Microbes” (John Wiley, UK, 1998) together with Frits Schut. After this he worked as a researcher at the JBI on image analysis of diatoms. He is currently Senior Lecturer at the JBI.