REPORT_RAPPORT_

PNA

Probability, Networks and Algorithms

Probability, Networks and Algorithms

Path Openings and Closings

Henk J.A.M. Heijmans, Michael Buckley, Hugues Talbot

CWI is the National Research Institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organization for Scientific Research (NWO).
CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Path Openings and Closings

Henk J.A.M. Heijmans

CWI

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

Michael Buckley

CSIRO Mathematical and Information Science
Locked Bag 17, North Ryde, NSW 1670
Australia

Hugues Talbot

CSIRO Mathematical and Information Science
Locked Bag 17, North Ryde, NSW 1670
Australia

ABSTRACT

This paper lays the theoretical foundations to path openings and closings. The traditional morphological filter used for the analysis of linear structures in images is the union of openings (or the intersection of closings) by linear segments. However structures in images are rarely strictly straight, and as a result a more flexible approach is needed. An extension to the idea of using straight line segments as structuring elements is to use constrained paths, i.e. discrete, one-pixel thick successions of pixels oriented in a particular direction, but in general forming curved lines rather than perfectly straight lines. However the number of such paths is prohibitive and the resulting algorithm by simple composition is inefficient. In this paper we propose a way to compute openings and closings over large numbers of constrained, oriented paths in an efficient manner, suitable for building filters with applications to the analysis of oriented features, such as for example texture.

1. INTRODUCTION

Practitioners of mathematical morphology are familiar with the importance of the structuring element in morphological and algebraic openings and closings. In spite of the infinite variety of available structuring elements, very few kinds of structuring elements are used in practice outside of a few specialized applications. The unit ball structuring elements of the discrete grid (e.g: diamond, square and hexagon) define a first family of common structuring elements, useful for basic filtering, granulometries, etc.

Probably the second most used structuring element family is generated by some instance of the discrete line segment, which is used when linear and oriented structures are present

in an application. This limited choice can be at least partly blamed on the dearth of truly efficient algorithms for more arbitrary structuring elements [6].

However most structures in real-world images are not perfectly straight, and therefore using line segments as structuring elements in openings and closings can be inadequate in the common situation where there exist narrow, locally oriented features in an image of interest. In this case one might be interested in using structuring elements that are themselves narrow and oriented, but not perfectly straight. Unfortunately generating useful morphological filters in the usual way by composition leads to computationally expensive and impractical algorithms.

In this paper we introduce the concept of path openings and closings, i.e. morphological filters that use families of structuring elements consisting of variously constrained paths, for which there exists algorithms as efficient as those using the usual families of straight line segments.

Path openings were originally proposed in [1] in an algorithmic, practical but incomplete manner. Here we are more concerned with laying down the theoretical foundations of these useful filters.

2. A GENTLE REMINDER ON MORPHOLOGY

In the sequel we assume basic familiarity of the reader with the general concepts of mathematical morphology. A general reference work is the recent book by Soille [5]. Two other useful resources are [2, 3].

2.1 BINARY IMAGES

It is generally accepted that the definition and investigation of morphological operators requires a partial ordering[1] on the space of images, henceforth denoted by L. This partial ordering, generally denoted by '$\leq$', induces a partial ordering on the family of morphological operators on $\mathcal{L}$. Given two operators ψ_1, ψ_2, we write $\psi_1 \leq \psi_2$ if and only if $\psi_1(X) \leq \psi_2(X)$, for every image $X \in \mathcal{L}$. An operator $\psi : \mathcal{L} \to \mathcal{L}$ is said to be *increasing* if $X \leq Y$ implies that $\psi(X) \leq \psi(Y)$ for every $X, Y \in \mathcal{L}$.

Throughout the remainder of this subsection we assume that $\mathcal{L} = \mathcal{P}(E)$, the subsets of E modeling the binary images with domain E. The partial ordering under consideration is now the set inclusion. A mapping $\delta : \mathcal{P}(E) \to \mathcal{P}(E)$ is called a *dilation* if it distributes over union, i.e.,

$$\delta\left(\bigcup_{i \in I} X_i\right) = \bigcup_{i \in I} \delta(X_i).$$

It follows immediately that δ is completely determined by the family of sets $A(x) = \delta(\{x\})$ where $x \in E$. Namely

$$\delta(X) = \bigcup_{x \in X} \delta(\{x\}) = \bigcup_{x \in X} A(x).$$

In the case of an infinite translation invariant domain, say $E = \mathbb{R}^d$ (but the same can be said for $\mathbb{Z}^d$), one often chooses $A(x) = A_x$, the translate of a given set $A \subseteq \mathbb{R}^d$, called *structuring element*, over the vector x. If, however, E is a mask outside which the image is unknown

---

[1] In fact, the modern setting of mathematical morphology requires the underlying image space to be a complete lattice [2].

(the computer screen being a typical example), and $A \subseteq \mathbb{R}^d$ is a structuring element, we might choose $A(x) = A_x \cap E$, for $x \in E$. In other words, $A(x)$ is obtained by choosing a fixed structuring element A and restricting the translate A_x to points inside the window E. This choice is an elegant way to handle the 'border problem' arising from the fact that we only have local knowledge of the image of interest.

Given a dilation δ on $\mathcal{P}(E)$, the *reciprocal dilation* $\check{\delta}$ is defined by

$$y \in \check{\delta}(\{x\}) \iff x \in \delta(\{y\}) \,.$$

Note that because of the aforementioned property of dilations, $\check{\delta}$ is entirely characterised by its behaviour on singletons.

2.2 OPENINGS AND CLOSINGS

An operator ψ on the partially ordered set $\mathcal{L}$ is anti-extensive if $\psi(X) \leq X$ for all $X \in \mathcal{L}$. It is extensive if $\psi(X) \geq X$ for all $X \in \mathcal{L}$. The operator is called idempotent if $\psi\psi = \psi^2 = \psi$.

An operator $\psi : \mathcal{L} \to \mathcal{L}$ is called an *opening* if ψ is anti-extensive, increasing and idempotent. Throughout this paper openings will be denoted by α. The dual concept is that of a *closing* which is an operator that is extensive, increasing and idempotent. In this paper we will mostly be concerned with openings, but we emphasise here that the Duality Principle (see [2] for a precise formulation) yields analogous results for closings without any additional effort.

An important property of openings, and one which we will exploit throughout this paper, is the following: if α_i is an opening for every $i \in I$, then $\bigvee_{i \in I} \alpha_i$ given by

$$\left(\bigvee_{i \in I} \alpha_i\right)(X) = \bigcup_{i \in I} \alpha_i(X) \,, \tag{2.1}$$

is an opening, too. The best known example of an opening is the *structural opening*, which is obtained as the composition of the Minkowski subtraction and addition and which is given by the expression

$$X \circ A = \bigcup \{A_x \ : \ x \in E \text{ and } A_x \subseteq X\} \,.$$

In other words, $X\circ$ is the union of all translates of structuring elements that lie within the original image.

A special case that is of interest in the context of this paper is the supremum of openings given by (2.1) where $\alpha_i(X) = X \circ A_i$, and the structuring element A_i is a line segment with length L and direction θ_i. The resulting opening, called *linear opening*, is useful for the analysis of images that contain oriented linear features. Examples are given toward the end of this paper.

If we take the union of all structural openings using an arbitrary connected structuring element with a given area, then we obtain the so-called *area opening*. In practice, the family of structuring elements involved may be huge, suggesting that the concept of an area opening is quite impractical. However, there does exist a very efficient algorithm due to Vincent [7] both in the binary and the grey-scale case.

2.3 THE GREY-SCALE CASE

So far we have been dealing exclusively with the binary case. But, as we will explain in this section, all concepts developed so far carry over immediately to the grey-scale case. We first briefly recall the level set construction of increasing morphological operators for grey-scale functions. We refer to [2, Chapters 10-11] for a comprehensive account. For simplicity we restrict ourselves to the case where the grey-level set equals $T = \{0, 1, \ldots, 255\}$, but, as explained in [2], the construction holds in many other cases as well. We denote by $\mathrm{Fun}(E, T)$ the grey-scale functions $I : E \to T$.

Recall that the *level set* (or *threshold set*) of a function I is given by

$$X_t(I) = \{x \in E \ : \ I(x) \geq t\}, \ t = 0, 1, \ldots, 255.$$

Given an increasing set operator ψ on $\mathcal{P}(E)$, there exists an extension Ψ to the set of grey-scale functions $\mathrm{Fun}(E, T)$ defined by means of

$$\Psi(I)(x) = \max\{0 \leq t \leq 255 \ : \ x \in \psi(X_t(I))\}.$$

It can be shown that the following relation holds:

$$X_t(\Psi(I)) = \psi(X_t(I)), \ t = 0, 1, \ldots, 255.$$

One calls Ψ the *flat function operator generated by ψ*. Several elementary facts regarding flat function operators have been established in [2]. We mention only the ones that we will use in the sequel:

(a) If Ψ_1, Ψ_2 are generated by ψ_1, ψ_2, then

    — $\Psi_2\Psi_1$ is generated by $\psi_2\psi_1$;

    — $\Psi_1 \vee \Psi_2$ is generated by $\psi_1 \vee \psi_2$;

    — $\Psi_1 \wedge \Psi_2$ is generated by $\psi_1 \wedge \psi_2$.

    In fact these last two properties hold for an arbitrary family of operators.

(b) If ψ is an opening, then Ψ is an opening as well.

3. PATHS AND DILATIONS

Let E be a given set of points with a binary adjacency relation '$\mapsto$' indicating the presence of an edge: thus $x \mapsto y$ means that that there is an edge going from x to y. The relation '$\mapsto$' is, in general, neither reflexive nor symmetric[2]. Because of the non-symmetry, the graph given by the vertices E and the adjacency relation $\mapsto$ is a *directed graph*. We show some examples in Fig. 1.

The examples in (a)-(c) are periodic, or more precisely, they are restrictions of a periodic adjacency relation on $\mathbb{Z}^2$ to a finite window. If $x \mapsto y$, we call y a successor of x and x a predecessor of y. Using the adjacency relation we can define a dilation on $\mathcal{P}(E)$ by putting

$$\delta(\{x\}) = \{y \in E \ : \ x \mapsto y\}.$$

---

[2] 'Reflexive' would mean that $x \mapsto x$ for every $x \in E$. 'Symmetric' would mean that $x \mapsto y$ iff $y \mapsto x$, for every $x, y \in E$.
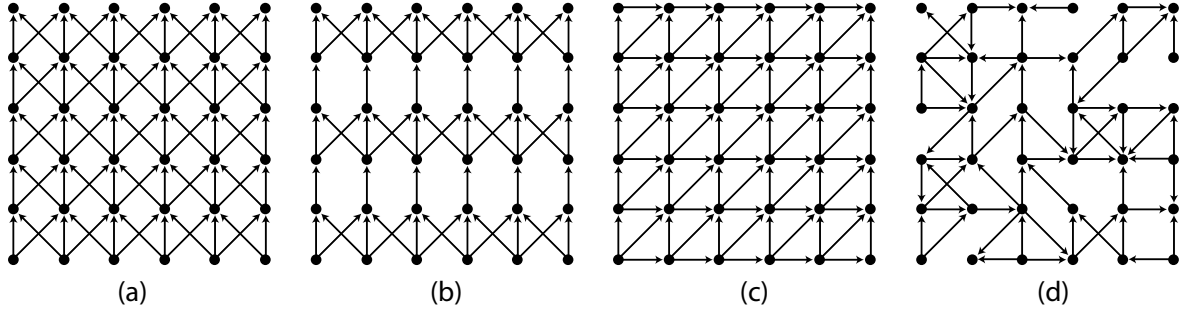
Figure 1: *Directed graphs.*

In other words, the dilation of a subset $X \subseteq E$ comprises all points which have a predecessor in X. These concepts are further illustrated in Fig. 2. Here b_1, b_2, b_3 are successors of a and $\delta(\{a\}) = \{b_1, b_2, b_3\}$. Furthermore, a_1, a_2, a_3 are the predecessors of b and $\breve{\delta}(\{b\}) = \{a_1, a_2, a_3\}$.
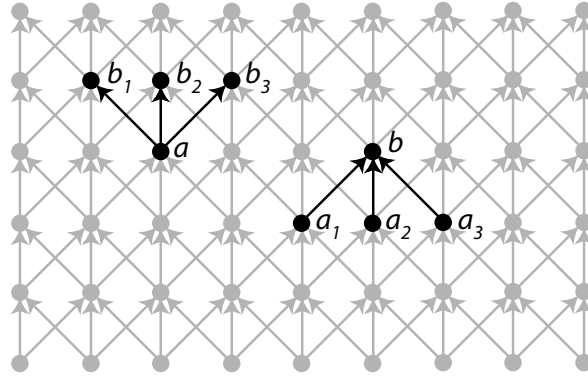
Figure 2: b_1, b_2, b_3 *are successors of* a *and* a_1, a_2, a_3 *are the predecessors of* b.

The L-tuple $\boldsymbol{a} = (a_1, a_2, \ldots, a_L)$ is called a *path of length* L if $a_k \mapsto a_{k+1}$, or equivalently, if

$$a_{k+1} \in \delta(\{a_k\}), \text{ for } k = 1, 2, \ldots, L - 1.$$

Henceforth we refer to such a path as a δ-path of length L. It is evident that $\boldsymbol{a} = (a_1, a_2, \ldots, a_L)$ is δ-path if and only if the reverse path $\breve{\boldsymbol{a}} = (a_L, a_{L-1}, \ldots, a_1)$ is a $\breve{\delta}$-path and obviously, both paths have the same length L. We denote the set of all δ-paths of length L by Π_L and the set of all $\breve{\delta}$-paths of length L by $\breve{\Pi}_L$. Given a path $\boldsymbol{a}$ in E, we denote by $\sigma(\boldsymbol{a})$ the set of its elements:

$$\sigma(a_1, a_2, \ldots, a_L) = \{a_1, a_2, \ldots, a_L\}.$$

The set of δ-paths of length L contained in a subset X of E is denoted by $\Pi_L(X)$, i.e.,

$$\Pi_L(X) = \{ \boldsymbol{a} \in \Pi_L \ : \ \sigma(\boldsymbol{a}) \subseteq X \} \,,$$

and the $\breve{\delta}$-paths of length L in X by $\breve{\Pi}_L(X)$.

In our approach, we first provided the adjacency relation and we defined the dilation afterward. It is straightforward to see that we might as well start with a dilation δ on $\mathcal{P}(E)$ and define the adjacency relation as

$$x \mapsto y \text{ if } y \in \delta(\{x\}) \,.$$

Both approaches are equivalent and it is a matter of taste which one is taken.

Note that in the first three examples in Fig. 1, the adjacency relation is periodic away from the borders. The fact that we can choose any dilation as a starting point enables us to handle the border in consistent and flexible manner. Note also the following major difference between the adjacency relation in Fig. 1 (a) and (c) and the one in (b). In (a) and (c), the graph structure is translation invariant with respect to any translation (again away from the borders), whereas in (b), this is only true if translation takes place over an even number of rows. Rephrased in terms of the dilation, this means that the structuring element is different at odd and at even rows. We will briefly address the issue of choosing the adjacency relation in Section 10.

4. Path opening

We define the operator $\alpha_L(X)$ as the union of all δ-paths of length L contained in X:

$$\alpha_L(X) = \bigcup \{ \sigma(\boldsymbol{a}) \ : \ \boldsymbol{a} \in \Pi_L(X) \} \,.$$

It is easy to see that α_L is an opening, and we call it the *path-opening*. We can define the reciprocal path-opening $\breve{\alpha}_L(X)$ in a similar way. Since $\boldsymbol{a} \in \Pi_L(X)$ iff $\breve{\boldsymbol{a}} \in \breve{\Pi}_L(X)$ and $\sigma(\boldsymbol{a}) = \sigma(\breve{\boldsymbol{a}})$ we get immediately that

$$\alpha_L = \breve{\alpha}_L \,.$$

It is obvious that $\alpha_1 = \mathsf{id}$. We can show that that

$$\alpha_{L+1} \leq \alpha_L \text{ for } L \geq 1 \,.$$

To prove this, assume that $x \in \alpha_{L+1}(X)$. Thus there is a δ-path $(a_1, a_2, \ldots, a_{L+1})$ of length $L + 1$ which contains x and lies inside X. But then both δ-paths $(a_1, a_2, \ldots, a_L)$ and $(a_2, a_3, \ldots, a_{L+1})$ of length L lie inside X and at least one of them must contain the point x. This proves that $x \in \alpha_L(X)$, too. In Fig. 3 we show an example of a path opening where $L = 6$.

4.1 Path decomposition

By definition, $x \in \alpha_L(X)$ iff there exists a δ-path $\boldsymbol{a} \in \Pi_L(X)$ that contains x, i.e., $x = a_k$ for some k between 1 and L. In that case we have

$$(a_1, a_2, \ldots, a_{k-1}, x) \in \Pi_k(X) \ \text{ and } \ (x, a_{k+1}, \ldots, a_L) \in \Pi_{L-k+1}(X) \,. \tag{4.1}$$
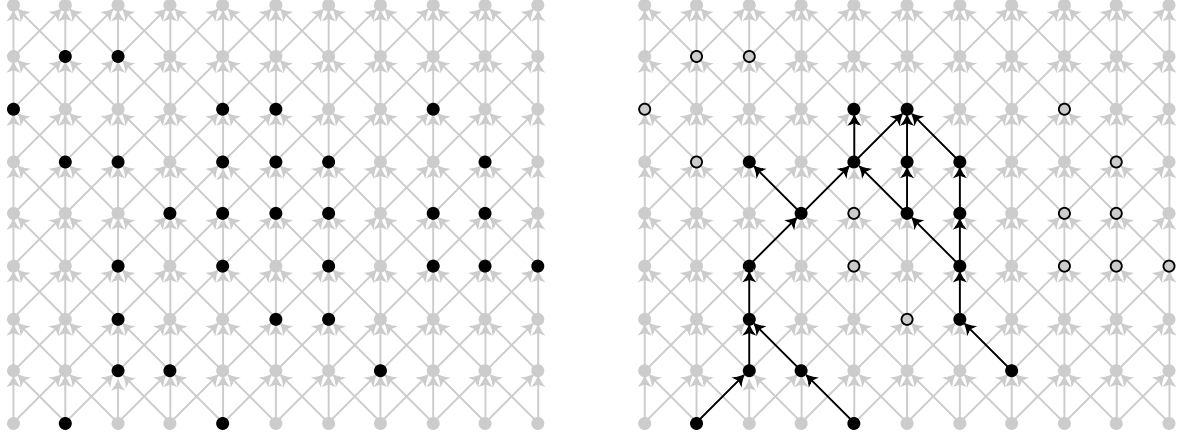
Figure 3: *A set $X \subseteq E$ (black points at the left) and its opening $\alpha_6(X)$ (black points at the right). Unfilled points at the right have been discarded.*

The first condition can be rewritten as

$$(x, a_{k-1}, a_{k-2}, \ldots, a_1) \in \breve{\Pi}_k(X) \,. \tag{4.2}$$

We define the operator ψ_k as

$$\psi_k(X) = \{a_1 \;:\; \boldsymbol{a} \in \Pi_k(X)\} \,,$$

that is, $\psi_k(X)$ contains the first point of every δ-path of length k in X. The operator $\breve{\psi}_k$ is defined analogously. Obviously, $\psi_1 = \breve{\psi}_1 = \mathsf{id}$.

Now the first condition in (4.1), which is equivalent to (4.2), can be written as $x \in \breve{\psi}_k(X)$, and the second condition can be written as $x \in \psi_{L-k+1}(X)$. Combined, they give

$$x \in \psi_{L-k+1}(X) \cap \breve{\psi}_k(X) \,,$$

and we have shown that

$$\alpha_L = \bigvee_{k=1}^{L} \left(\psi_k \wedge \breve{\psi}_{L-k+1} \right) \,, \tag{4.3}$$

where we have interchanged k and $L - k + 1$.

Note that $\alpha_2 = \mathsf{id} \wedge (\delta \vee \breve{\delta})$, which is known in the literature as the *annular opening* [2, 4].

4.2 Recursive structure of the operators ψ_k

We will prove below that the following relations hold:

$$\psi_{k+1} = \mathsf{id} \wedge \breve{\delta}\psi_k \quad \text{and} \quad \breve{\psi}_{k+1} = \mathsf{id} \wedge \delta\breve{\psi}_k \,. \tag{4.4}$$

The decomposition of α_L in (4.3) together with the iterative formulas in (4.4) provide an efficient algorithm for the computation of the path opening α_L. Algorithm 2 in Section 9

provides an illustration for a vertical path on a square grid with the adjacency of Fig 1(a). This algorithm has complexity $O(L)$ which is similar to that of the more common union of openings by line segments, as shown in [1].

We will prove only the first identity in (4.4) as the second is nothing but its reciprocal version. To prove '$\leq$' assume that $x \in \psi_{k+1}(X)$. This means that there exist $a_2, \ldots, a_{k+1}$ such that $(x, a_2, \ldots, a_{k+1}) \in \Pi_{k+1}(X)$. Now $(a_2, \ldots, a_{k+1}) \in \Pi_k(X)$ and $x \in \breve{\delta}(\{a_2\})$. Since $a_2 \in \psi_k(X)$ this yields that $x \in \breve{\delta}(\psi_k(X))$, and we conclude that $x \in X \cap \breve{\delta}(\psi_k(X))$.

To prove '$\geq$', let $x \in (\mathsf{id} \wedge \breve{\delta}\psi_k)(X)$, i.e., $x \in X$ and $x \in \breve{\delta}(\{y\})$ with $y \in \psi_k(X)$. The latter means that there exist $a_2, \ldots, a_k$ such that $(y, a_2, \ldots, a_k) \in \Pi_k(X)$. Now $(x, y, a_2, \ldots, a_k) \in \Pi_{k+1}(X)$, which yields that $x \in \psi_{k+1}(X)$.

The path opening α_L depends strongly upon the dilation, or equivalently, the adjacency relation. This is clearly seen in Fig. 4 where we have computed the opening $\alpha_5(X)$ of a set X for three different adjacencies.
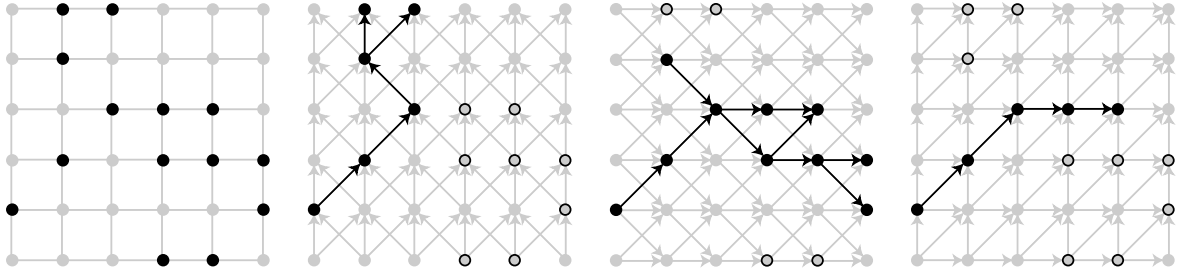

Figure 4: *A set $X \subseteq E$ (left) and its opening $\alpha_5(X)$ for three different adjacencies.*

In Section 2 we have seen that a union of openings is an opening. Therefore we can, for example, take the union of the first two openings in Fig. 4, i.e., two figures in the middle, to get an opening that allows both horizontal and vertical oriented paths. Note however, that this is not the same as combining both adjacencies into one and computing the opening with respect to this new adjacency.

Finally we note that

$$\psi_k \psi_l = \psi_{k+l-1}, \text{ for } k, l \geq 1,$$

which implies in particular that $\psi_2^k = \psi_{k+1}$, that is

$$\psi_{k+1} = (\mathsf{id} \wedge \breve{\delta})^k, \text{ for } k \geq 1.$$

The proof of these semi-group identities is rather straightforward and left as an exercise for the reader.

5. OPENING TRANSFORM

Often, we are interested in all openings $\alpha_L(X)$ of a set X for a range of values of L rather than for a single value only. For example, it is quite common that we do not know beforehand which L to choose in a particular application. In such cases it may be more efficient to

compute the so-called *opening transform* of the image. Given a set $X \subseteq E$ and an ordered family of openings $\mathcal{A} = \{\alpha_L\}$, the opening transform A_X of X with respect to $\mathcal{A}$ is a function mapping the domain E into $\mathbb{Z}_+ = \{0, 1, 2, \ldots\}$ such that its threshold sets correspond with the various openings $\alpha_L(X)$: see (5.2) below for an exact formulation.

Throughout this section we assume that there exist only finite paths. More precisely we assume that there exists an integer $N \geq 1$ such that $\delta^N(\{x\}) = \emptyset$, for every $x \in E$, i.e., $\delta^N(E) = \emptyset$. Note, however, that this assumption does not necessarily mean that E is finite. Furthermore, we take N to be the smallest integer with this property. Thus the maximal length of a path in E is N. Define $\lambda(x)$ as the maximal length of a δ-path with begin-point x:

$$
\begin{aligned}
\lambda(x) &= \max\{L \geq 1 \ : \ \exists\, \boldsymbol{a} \in \Pi_L \text{ such that } a_1 = x\} \\
&= \max\{L \geq 1 \ : \ x \in \check{\delta}^{L-1}(E)\}\,.
\end{aligned}
$$

Obviously, if $x \mapsto y$ then $\lambda(x) \geq \lambda(y) + 1$. Moreover, it is not difficult to prove that

$$
\lambda(x) = 1 + \max\{\lambda(y) \ : \ x \mapsto y\}\,,
$$

where the maximum is taken to be zero if x has no successors. Similarly $\check{\lambda}(x)$ is the maximal length of a path with endpoint x . Then

$$
\Lambda(x) = \lambda(x) + \check{\lambda}(x) - 1\,, \tag{5.1}
$$

is the length of the longest path that contains x.

To understand why this is correct, assume that $\lambda(x) = m$ and $\check{\lambda}(x) = n$. Thus we know that there exist paths of the form $(a_1, \ldots, a_{n-1}, x)$ and $(x, b_2, \ldots, b_m)$ in E. The concatenation $(a_1, \ldots, a_{n-1}, x, b_2, \ldots, b_m)$ of these two paths has length $n + m - 1$. It is also obvious that there cannot be a longer path in E containing x, for such a path could then be broken at position x yielding two paths $(a_1, \ldots, a_{n'-1}, x)$ and $(x, b_2, \ldots, b_{m'})$ with either $m' \geq m$ or $n' \geq n$ or both.

Define the disjoint partition $E_1, E_2, \ldots, E_N$ of E by

$$
E_L = \{x \in E \ : \ \lambda(x) = L\}\,.
$$

In the figures at the left in Fig. 5 we depict this partition for two different adjacency relations; here the arrows indicate the relation $x \mapsto y$.

Now we define a function $F_X : E \to \mathbb{Z}_+$ by means of algorithm 1 which resembles a geodesic propagation algorithm.

Algorithm 1.

```
F_X = 0 on E        /* initialisation */
for k = 1 to N
    for x ∈ E_k ∩ X
        F_X(x) = 1 + max{F_X(y) : y ↦ x}
    end
end
```

The second column of Fig. 5 shows the function F_X for a given image X (grey pixels). In a similar way we can define $\breve{F}_X$ by using the partition $\breve{E}_1, \breve{E}_2, \ldots, \breve{E}_N$.

The following lemma shows that $\psi_k(X)$ can be obtained by thresholding of F_X.

Lemma 5.1. *With the definitions given before we have*

$$\psi_k(X) = \{x \in E \ : \ F_X(x) \geq k\} \quad and \quad \breve{\psi}_k(X) = \{x \in E \ : \ \breve{F}_X(x) \geq k\},$$

for $k = 1, 2, \ldots, N$.

Proof. Let $x \in E_k$ and suppose that $F_X(x) = l$. Obviously, $l \leq k$ and there must exist a δ-path $x = a_1, a_2, \ldots, a_l$ in $\Pi_l(X)$ such that $a_i \in E_{k-i+1}$ and $F_X(a_i) = l - i + 1$. This implies that $x \in \psi_l(X)$. Conversely, if $x \in \psi_l(X)$, then there is a path $x = a_1, a_2, \ldots, a_l$ in $\Pi_l(X)$. Now if $x \in E_k$, where $k \geq l$, then $a_i \in E_{k-i+1}$ and $F_X(a_i) = l - i + 1$. $\qquad\square$

This lemma can be used to prove the following result.

Proposition 5.2. *The function $A_X = F_X + \breve{F}_X - 1$ is the opening transform of X, that is*

$$\alpha_L(X) = \{x \in E \ : \ A_X(x) \geq L\}, \tag{5.2}$$

for every $L \geq 1$ and $X \subseteq E$.

Proof. We use the expression for α_L in (4.3) which says that $x \in \alpha_L(X)$ implies that $x \in \psi_k(X) \cap \breve{\psi}_{L-k+1}(X)$ for some $k = 1, 2, \ldots, L$. Therefore, $F_X(x) \geq k$ and $\breve{F}_X(x) \geq L - k + 1$, which yields that $A_X(x) \geq L$. This proves '$\subseteq$' in (5.2). The converse is proved similarly. $\quad\square$

In Fig. 5, the algorithm for the opening transform is shown for two different adjacencies, namely (a) and (c) in Fig. 1.

6. Border issues

Throughout this section, E will be a finite rectangular window within $\mathbb{Z}^2$, and the adjacency on E is the restriction of a periodic adjacency on $\mathbb{Z}^2$ like in Fig. 1 (a)-(c). The *inward boundary* of E, denoted by ∂E, is the set of points in E which have a predecessor outside E:

$$\partial E = \{x \in E \ : \ \exists\, y \in E^c \text{ such that } y \mapsto x\},$$

where E^c is the complement of E. The *outward boundary* of E, denoted by $\breve{\partial} E$, is given by

$$\breve{\partial} E = \{x \in E \ : \ \exists\, y \in E^c \text{ such that } x \mapsto y\},$$

There are various ways to deal with the border problem:

(a) We can simply ignore the existence of the borders and treat paths that contain boundary points in the same way as any other path. In fact, this is the choice that we have implicitly made so far.

(b) The other extreme is to set the length of a path that crosses the border to $+\infty$, meaning that all points on such a path are contained in every opening $\alpha_L(X)$. In fact, this choice means that we extend X outside E by adding all points in E^c.

function λ

1	1	1	1	1	1
2	2	2	2	2	2
3	3	3	3	3	3
4	4	4	4	4	4
↖	↑	↗	5	5	5
6	↓	6	6	6	6

function F_X

	1	1			
	2				
		3	1	1	
	4		4	2	2
5					3
				1	4

function $\check{F}_X$

	5	5			
	4				
		3	4	4	
	2		1	3	3
1					2
				1	1

opening transform A_X

	5	5			
	5				
		5	4	4	
	5		4	4	4
5					4
				1	4

partition E_k

6	5	4	3	2	1
7	6	5	4	3	2
8	7	6	5	4	3
9	8	7	6	5	4
10	↑	↗	7	6	5
11	↓	→	8	7	6

function F_X

	2	1			
	3				
		3	2	1	
	4		3	2	1
5					2
				4	3

function $\check{F}_X$

	2	3			
	1				
		3	4	5	
	2		1	2	4
1					3
				1	2

opening transform A_X

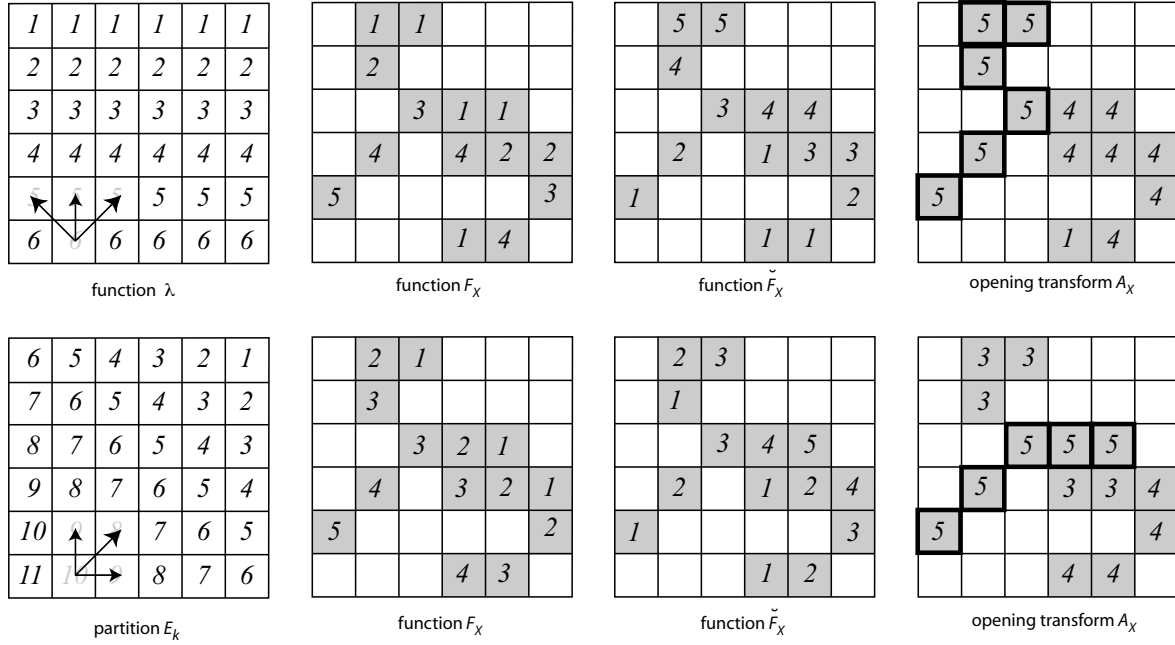|					
	3	3			
	3				
		5	5	5	
	5		3	3	4
5					4
				4	4

Figure 5: *Computation of the opening transform for two different adjacencies indicated by the arrows in the first column, which also shows the partition of E. The second and third columns depict the functions F_X and $\check{F}_X$ for the set X represented by the grey pixels. The right column shows the opening transform A_X and the opened set $\alpha_5(X)$ represented by the pixels with the thick boundaries.*

(c) An intermediate option is to try to compensate for the points cut off by restricting to a finite window by replacing the computed length L of a path that has a begin-point in ∂E or an endpoint in $\check{\partial} E$ by $h(L)$. One possible choice for h would be $h(L) = 2L$. Such a choice could be justified by the presumption that on average only half of the path falls inside the window. Another possibility is to add a fixed compensation to the length of a border-crossing path, i.e., $h(L) = L + L_0$. Note, that one might use different compensation functions for paths that start on ∂E and end on $\check{\partial} E$.

(d) A possibility which is easy to implement is to enlarge the window E with a border B of thickness L_0. Denote by α'_L the corresponding path opening on $\mathcal{P}(E')$, where $E' = E \cup B$ denotes the enlarged window. Thus we can compute α'_L according to the algorithm given in the previous sections. Define the opening α_L on $\mathcal{P}(E)$ by

$$\alpha_L(X) = \alpha'_L(X \cup B) \cap E.$$

In Fig. 6 we show that in this case, two paths which were originally disjoint may considered to be part of the same path which lies partially outside the window.
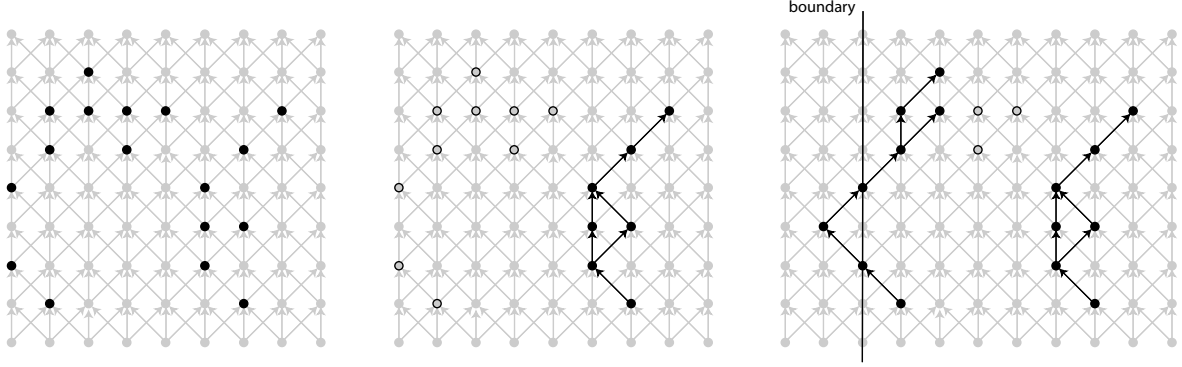
Figure 6: *Opening $\alpha_6(X)$ of set X (left) according to option (a) (middle) and (d) (right).*

7. INCOMPLETE PATH OPENINGS

The path-opening $\alpha_L(X)$ of a set X comprises the union of all length-L paths contained inside X. We can relax this condition by demanding that only k out of L vertices of the path lie inside X, thus yielding a so-called *incomplete path-opening* $\alpha_L^k(X)$. Incomplete paths can be used to robustify path openings against noise. We present a formal definition below.

Define $\Pi_L^k(X)$ as the collection of length-L paths in E which contain at least k points inside X:

$$\Pi_L^k(X) = \{\boldsymbol{a} \in \Pi_L \ : \ |\sigma(\boldsymbol{a}) \cap X| \geq k\}.$$

Note that this definition only makes sense for $0 \leq k \leq L$, and that

$$\Pi_L(X) = \Pi_L^L(X) \subseteq \Pi_L^{L-1}(X) \subseteq \cdots \subseteq \Pi_L^1(X) \subseteq \Pi_L^0(X) = \Pi_L.$$

Define the *incomplete path-opening*

$$\alpha_L^k(X) = \bigcup\{\sigma(\boldsymbol{a}) \cap X \ : \ \boldsymbol{a} \in \Pi_L^k(X)\}.$$

It is obvious that

$$\alpha_L^L \leq \alpha_L^{L-1} \leq \cdots \leq \alpha_L^0,$$

and that

$$\alpha_L^L = \alpha_L \text{ and } \alpha_L^0(X) = \{x \in X \ : \ \Lambda(x) \geq L\},$$

where $\Lambda(x)$ was defined in (5.1). Putting

$$\bar{E}_L = \{x \in E \ : \ \Lambda(x) \geq L\},$$

we get that

$$\alpha_L^0(X) = X \cap \bar{E}_L.$$

Furthermore, we define

$$\psi_L^k(X) = \{a_1 \ : \ \boldsymbol{a} \in \Pi_L^k(X)\}.$$

We have

$$\psi_L^L = \psi_L \text{ and } \psi_L^0(X) = \{x \in E \; : \; \breve{\lambda}(x) \geq L\} = \bigcup_{k \geq L} \breve{E}_k \,.$$

Henceforth we use the convention that $\psi_L^k \equiv \emptyset$ if $k > L$. We will now express ψ_{L+1}^{k+1} in terms of ψ_L^{k+1} and ψ_L^k. Observe that $x \in \psi_{L+1}^{k+1}(X)$ if there exists $\boldsymbol{a} = (a_1, \ldots, a_L)$ such that $(x, a_1, \ldots, a_L) \in \Pi_{L+1}$ and either $x \in X$ and $\boldsymbol{a} \in \psi_L^k(X)$ or $\boldsymbol{a} \in \psi_L^{k+1}(X)$. We have shown that

$$\psi_{L+1}^{k+1} = (\mathsf{id} \wedge \breve{\delta}\psi_L^k) \vee \breve{\delta}\psi_L^{k+1} \,.$$

Since $\psi_L^{k+1} \leq \psi_L^k$ we can also write

$$\psi_{L+1}^{k+1} = (\mathsf{id} \vee \breve{\delta}\psi_L^{k+1}) \wedge \breve{\delta}\psi_L^k \,.$$

Note that for $k = L$ we obtain $\psi_{L+1}^{L+1} = \mathsf{id} \wedge \breve{\delta}\psi_L^L$, which coincides with the expression derived in (4.4).

Now consider a point $x \in \alpha_L^k(X)$, where $1 \leq k \leq L$. Thus there is a path $\boldsymbol{a} \in \Pi_k$ with $x \in \sigma(\boldsymbol{a})$ such that $|\sigma(\boldsymbol{a}) \cap X| \geq k$. Assume that $a_l = x$. Now $\boldsymbol{a}$ is the concatenation of the sequences $\boldsymbol{b} = (a_1, \ldots, a_{l-1}, x)$ and $\boldsymbol{c} = (x, a_{l+1}, \ldots, a_L)$. Define $j = |\sigma(\boldsymbol{b}) \cap X|$, hence $|\sigma(\boldsymbol{c}) \cap X| = k + 1 - j$. We conclude that

$$x \in \breve{\psi}_l^j(X) \cap \psi_{L-l+1}^{k-j+1}(X) \,.$$

Since the length of $\boldsymbol{b}$ and $\boldsymbol{c}$ is l and $L + 1 - l$, respectively, we have

$$1 \leq j \leq l \text{ and } k + 1 - j \leq L + 1 - l \,.$$

Furthermore we have

$$j \leq k \,,$$

and we conclude that

$$\alpha_L^k = \bigvee_{l=1}^{L} \bigvee_{j=\max\{1, k+l-L\}}^{\min\{l, k\}} \left(\breve{\psi}_l^j \wedge \psi_{L-l+1}^{k-j+1} \right) \,, \tag{7.1}$$

for $1 \leq k \leq L$. Observe that this expression reduces to the one in (4.3) if $k = L$.

Denote by $\theta(L, k)$ the number of terms in (7.1). One can show that

$$
\begin{aligned}
\theta(2p - 1, k) &= k(2p - k), \text{ for } k = 1, \ldots, p \\
\theta(2p - 1, k) &= \theta(2p - 1, 2p - k), \text{ for } k = p + 1, \ldots, 2p - 1 \,,
\end{aligned}
$$

and that

$$
\begin{aligned}
\theta(2p, k) &= k(2p - k + 1), \text{ for } k = 1, \ldots, p \\
\theta(2p, k) &= \theta(2p, 2p - k + 1), \text{ for } k = p + 1, \ldots, 2p \,.
\end{aligned}
$$

For example, if L is even and the degree of relaxation $r = L - k$ is small, then the number of terms in (7.1) is approximately $(r + 1)L$ so for fixed L the computational complexity is proportional to $r + 1$.

8. CYCLE OPENINGS

In most of the examples considered so far, the adjacency did not allow cyclic paths (henceforth called *cycles*) $\boldsymbol{a} = (a_1, \ldots, a_L)$ with $a_L = a_1$. Moreover, the assumption $\delta^N(E) = \emptyset$ made in Section 4 explicitly excludes the existence of a cycle. In this section we will briefly discuss the case where cycles may exist. Moreover, we define a new family of openings which preserves only cycles in a set. We start with some definitions.

We say that that the dilation δ (or alternatively, the adjacency relation '$\mapsto$') is *p-recurrent* if

$$x \in \delta^p(\{x\}) \quad \text{for every } x \in E,$$

and if p is the smallest integer for which this holds. This means that for every $x \in E$ there exist $y_1, \ldots, y_{p-1}$ such that $x \mapsto y_1 \mapsto y_2 \mapsto \cdots \mapsto y_{p-1} \mapsto x$. In other words, every x lies on a cycle of length p. We show two examples in Fig. 7.
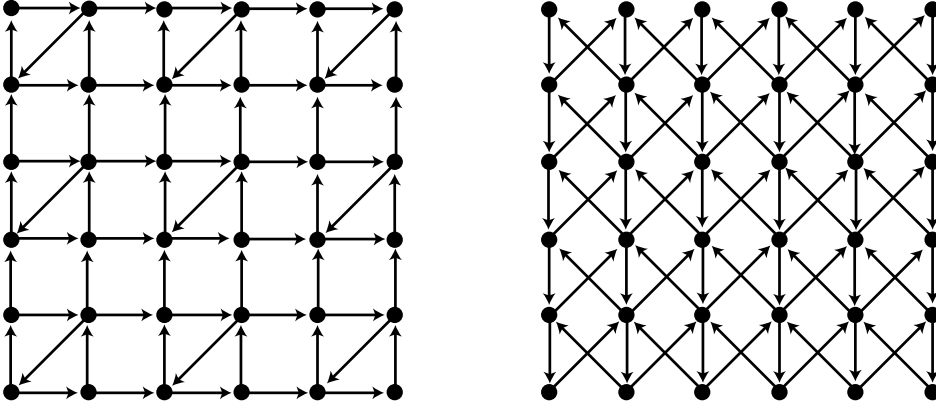

Figure 7: *A 3-recurrent and a 4-recurrent adjacency relation.*

If δ is not p-recurrent for any integer p, then we say that δ is *non-recurrent*. It is evident that δ is p-recurrent if and only if $\check{\delta}$ is p-recurrent. Similarly, δ is non-recurrent if and only if $\check{\delta}$ is non-recurrent.

Define $\alpha_L^\circ(X)$ as the union of all length-L cycles contained in X, and $\alpha^\circ(X)$ as the the union of all cycles (with length ≥ 2) contained in X:

$$\alpha^\circ(X) = \bigcup_{L \geq 2} \alpha_L^\circ(X).$$

Recall that $\check{\psi}_2(X) = X \cap \delta(X)$ removes begin-points of a path and, dually, that $\psi_2(X) = X \cap \check{\delta}(X)$ removes end-points. Thus applying $\omega = \psi_2 \wedge \check{\psi}_2$ we remove all begin- and end-points of paths within the set and after iteration we are left with the cycles:

$$\alpha^\circ(X) = \bigvee_{k \geq 1} \omega^k(X).$$

9. THE GREY-SCALE CASE

In this section we use the framework discussed in § 2.3 to extend the results developed in the previous sections to the grey-scale case, with the exception of the opening transform which is not defined in the grey-scale case to the best of our knowledge.

First we define the grey-scale analogue of paths of length L in the grey-scale image I. This involves an additional parameter t representing the grey-level of the path:

$$\Pi_L^t(I) = \{\boldsymbol{a} \in \Pi_L \ : \ I(a_k) \geq t, \ k = 1, 2, \ldots, L\}.$$

It is easy to see that

$$\Pi_L^t(I) = \Pi_L(X_t(I)).$$

Let Ψ_k be the flat extension of ψ_k defined by means of the level set construction, and let $\mathcal{A}_L$ be the flat extension of α_L.

Proposition 9.1. *The operator Ψ_k and the opening $\mathcal{A}_L$ on* $\mathrm{Fun}(E, T)$ *are, respectively, given by*

$$
\begin{aligned}
\Psi_k(I)(x) &= \max\{t \in T \ : \ (x, a_2, \ldots, a_k) \in \Pi_k^t(I) \\
&\qquad \textit{for some } a_2, \ldots, a_k \in E\}, \\
\mathcal{A}_L(I)(x) &= \max\{t \in T \ : \ \exists\ \boldsymbol{a} \in \Pi_L^t(I) \textit{ with } x \in \sigma(\boldsymbol{a})\}.
\end{aligned}
$$

From the theory on flat function operators we know that the expressions in (4.3) and (4.4) carry over immediately to the function case. Let Δ and $\check{\Delta}$ be the extension of δ and $\check{\delta}$, respectively, to functions, i.e.,

$$
\begin{aligned}
\Delta(I)(x) &= \max\{I(y) \ : \ y \mapsto x\} \\
\check{\Delta}(I)(x) &= \max\{I(y) \ : \ x \mapsto y\}.
\end{aligned}
$$

Putting $J_k = \Psi_k(I)$ and $\check{J}_k = \check{\Psi}_k(I)$ (hence $J_1 = \check{J}_1 = I$) we obtain from (4.4) that

$$
\begin{aligned}
J_{k+1}(x) &= \min\left\{I(x), \check{\Delta}(J_k)(x)\right\} = \min\left\{I(x), \max_{x \mapsto y} J_k(y)\right\} \\
\check{J}_{k+1}(x) &= \min\left\{I(x), \Delta(\check{J}_k)(x)\right\} = \min\left\{I(x), \max_{y \mapsto x} \check{J}_k(y)\right\}.
\end{aligned}
$$

We consider the path opening associated with the North-South adjacency in Fig. 1(a). For a given vertex x, the adjacency $x \mapsto y$ characterises three vertices y which are denoted by $x_\nwarrow, x_\uparrow, x_\nearrow$. Similarly, the reciprocal adjacency $y \mapsto x$ characterises the vertices $x_\swarrow, x_\downarrow, x_\searrow$. With this notation we arrive at the following algorithm which, for a given input image I and path length L, computes the path opening $I^\alpha = \mathcal{A}_L(I)$.

Algorithm 2 (North-South path opening).

$$J_1 \leftarrow I; \quad \breve{J}_1 \leftarrow I;$$

for $1 \leq k \leq L - 1$ **do**

$$J_{k+1}(x) = \min\Big\{ I(x), \max\big(J_k(x_\swarrow), J_k(x_\downarrow), J_k(x_\searrow) \big) \Big\} \ ;$$

$$\breve{J}_{k+1}(x) = \min\Big\{ I(x), \max\big(\breve{J}_k(x_\nwarrow), \breve{J}_k(x_\uparrow), \breve{J}_k(x_\nearrow) \big) \Big\} \ ;$$

end for ;

for $1 \leq k \leq L$ **do**

$$J_k^\alpha(x) = \min\big(J_k(x), \breve{J}_{L-k+1}(x) \big)$$

end for ;

$$I^\alpha(x) = \max\big(J_1^\alpha(x), J_2^\alpha(x), \ldots, J_L^\alpha(x) \}$$

10. EXAMPLE

Figure 8 is an example of path closing compared with other methods. We chose to illustrate with a closing rather than an opening because of the better contrast in the printing process, but the same conclusions would apply to both.

Fig. 8(a) is the original 500×160 image. This is an image of DNA (the long thin structure) observed in a scanning electron microscope. The objective is to separate the DNA from the noisy background, and we use various closings as a pre-processing filtering.

Fig. 8(b) is the result of applying a closing by intersection of 44 segments of length 23 pixels, each in a different direction, approximately uniformly oriented (subject to the digital grid). As can be seen this closing filters the background is mostly filtered out but so is the DNA.

Fig. 8(c) is a path closing with path length of 33 pixels. While this is longer than that of the straight segments family in (b), the shape and contrast of the DNA is well preserved, while the background is reasonably filtered out.

Fig. 8(d) is an area closing with parameter 50 square pixels. The DNA is well preserved as in (c) but the background is not as filtered out.

As can be seen, the path closing was able to better preserve the shape of the object of interest than the closing with segments, while removing more of the unwanted background than the area closing, which was the intended behaviour.

For this example, we ran algorithm 2 four times and took the infimum of the four results. The final result is still therefore a closing. In the first two instances, the adjacency graph was (1) vertical, each pixel being connected to the 3 pixels immediately above (N, NW and NE directions) and the 3 pixels below (S, SW and SE directions) ; and (2) horizontal, each pixel being connected to the 3 pixels immediately to the right and the 3 pixels immediately to the left, each of the those in the basic 3×3 neighborhood of the square grid.

In the last two instances, the adjacency graph was diagonal: instance (3) in the NE-SW direction, the central pixel in the 3×3 neighborhood being connected to the three pixels in the N, NE and E direction on the one hand, and to the three pixels in the S, SW and W direction. The last instance (4) was in the NW-SE direction symmetrically.

The result of this combination was to choose a family of paths such that at each point the entirety of the path was contained in a 90 degree angle double-ended cone, either vertically or diagonally. In some sense this captures the idea of a family of oriented, but flexible structuring elements.

It is of course possible to modify the adjacency graphs in order to constrain these paths more or less.

11. CONCLUSIONS

In this paper we have explored the theory of path openings and closings on binary and grey-level images. Path openings are openings over large number of connected or disconnected paths, which extend the useful notions of openings by unions of line segments by allowing the use of oriented, narrow but non-straight segments as a family of structuring elements. Because of the oriented nature of the family of structuring elements used, the resulting operators are more constrained than area openings.

Path openings and closing essentially allow practitioners to close the gap between openings by line segments (which are constrained and anisotropic) and area openings (which are unconstrained and isotropic). The framework developed in this paper allows for paths which behave more closely like one or the other, which should be useful as well, by varying the adjacency relation.

We have provided a workable algorithm for such path openings with low complexity, which makes the computation of such paths possible.
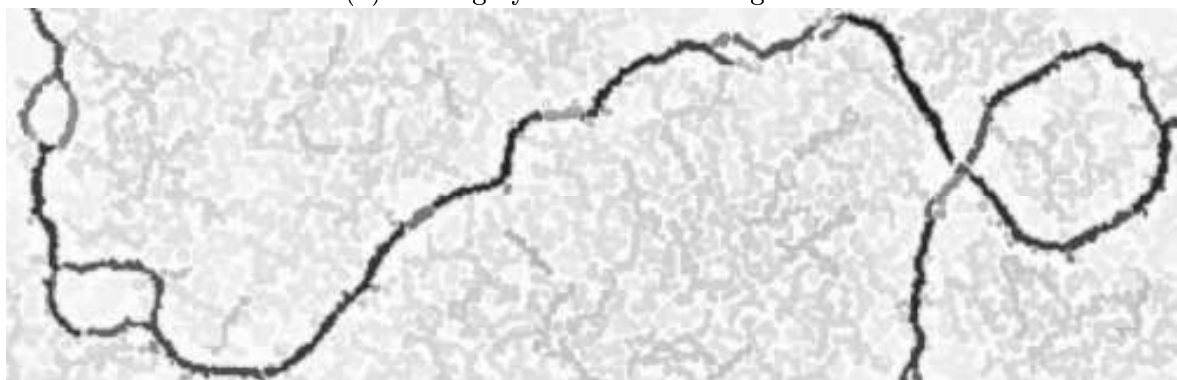
Finally we have explored the questions of how to deal with border effects, how to compute path opening transforms (only in the binary case) and how to extend path openings to incomplete paths, which could provide a degree of robustness against noise.
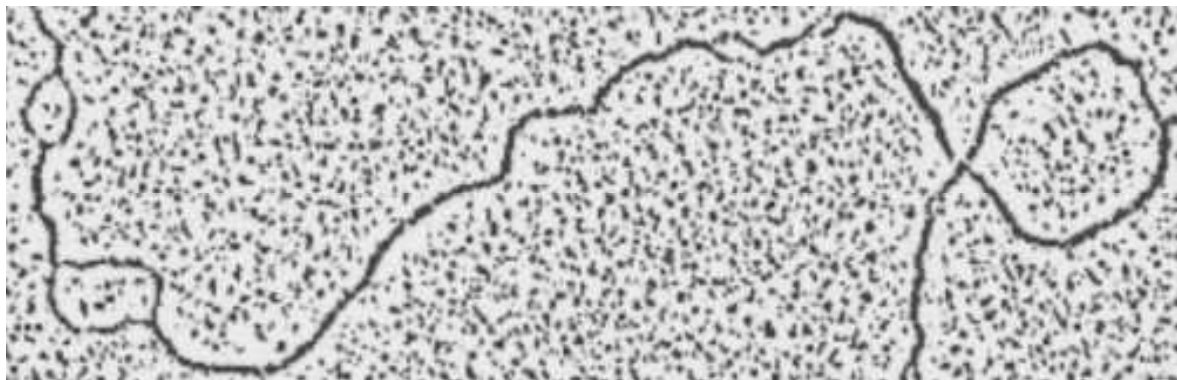
(a) Original image

(b) Closing by intersection of segments

(c) Path closing

(d) Area closing

Figure 8: *Example of path closing compared with closing with segments and area closing.*

References

1. M. Buckley and H. Talbot. Flexible linear openings and closings. In L. Vincent and D. Bloomberg, editors, *Mathematical Morphology and its application to image analysis*, Palo Alto, June 2000. ISMM, Kluwer.

2. H. J. A. M. Heijmans. *Morphological Image Operators*. Academic Press, Boston, 1994.

3. J. Serra. *Image Analysis and Mathematical Morphology*. Academic Press, London, 1982.

4. J. Serra, editor. *Image Analysis and Mathematical Morphology. II: Theoretical Advances*. Academic Press, London, 1988.

5. P. Soille. *Morphological Image Analysis, principles and applications*. Springer, 1999.

6. M Vandroogenboeck and H. Talbot. Fast computation of morphological operations with arbitrary structuring elements. *Pattern Recognition Letters*, 17:1451–1460, 1996.

7. L. Vincent. Grayscale area openings and closings, their efficient implementation and applications. In *Proceedings of the conference on mathematical morphology and its applications to signal processing*, pages 22–27, Barcelona, Spain, May 1993.