Image Denoising by Statistical Area Thresholding

D. COUPIER, A. DESOLNEUX* AND B. YCART
MAP5, UMR CNRS 8145, Université René Descartes, Paris
coupier@math-info.univ-paris5.fr
desolneux@math-info.univ-paris5.fr
ycart@math-info.univ-paris5.fr

Abstract. Area openings and closings are morphological filters which efficiently suppress impulse noise from an image, by removing small connected components of level sets. The problem of an objective choice of threshold for the area remains open. Here, a mathematical model for random images will be considered. Under this model, a Poisson approximation for the probability of appearance of any local pattern can be computed. In particular, the probability of observing a component with size larger than k in pure impulse noise has an explicit form. This permits the definition of a statistical test on the significance of connected components, thus providing an explicit formula for the area threshold of the denoising filter, as a function of the impulse noise probability parameter. Finally, using threshold decomposition, a denoising algorithm for grey level images is proposed.

Keywords: image denoising, mathematical morphology, area opening and closing, random image, threshold function, Poisson approximation, lattice animals

1. Introduction

The general problem of image denoising consists of deciding what is the "signal" and should be kept, and what is the noise, and must be removed. Many different criteria can be used to detect the noise-induced structures. For example, the oscillations due to an additive gaussian noise can be measured in terms of the wavelet coefficients. The noise may then be removed by a thresholding in the wavelet domain. Donoho and Johnstone [6] gave an explicit way to choose the threshold as a function of the variance of the noise. Their claim is that "denoising, with high probability, rejects pure noise completely". The underlying idea is that in pure noise, all the structures that actually belong to the image could not appear; or else, the structures coming from the image itself can be defined as those "objects" which would have a very small probability of appearing in a pure noise. This idea was implemented in [4, 5]

for the detection of alignments and meaningful level lines in an image.

Here, we shall focus on the size of connected components of the level sets of the image. Removing small components is a classical and efficient way of removing impulse noise from an image. This method, known as "the grain filter", was first introduced in the framework of Mathematical Morphology [16] by Vincent in [20] as morphological area openings and closings (see also [8, 21]). This filter is sometimes called the "extrema killer". It was then generalized by Masnou and Morel in [12], and by Monasse and Guichard in [14]. In [15] a similar filter was used, in the framework of gradient percolation, for recovering fuzzy images.

But the main question remains: how should the threshold for the area of the components that have to be kept, be chosen? A natural idea, imported from statistical inference, consists of fixing an a priori risk level ε (e.g. $\varepsilon = 0.001$), and deciding that anything that has probability lower than ε of occurring under a pure noise hypothesis cannot come from the noise and hence should be kept in the image. Thus for the threshold

area, one will choose the integer $s(n, p, \varepsilon)$, such that a connected component of size $k \geq s(n, p, \varepsilon)$ has a probability less than ε of appearing in a pure noise image with probability parameter p and size $n \times n$. Applying a grain filter with area threshold $s(n, p, \varepsilon)$ will ensure that, with probability larger than $1 - \varepsilon$, pure noise is eliminated. To implement this, one must be able to compute the probability of any connected component of size k appearing in a pure noise image. An exact computation is not feasible. However an approximation can be given if the image is large: our main theoretical result (Theorem 2.4) gives a Poisson approximation for the probability of occurrence for any image property which is local in the sense that its definition involves only a fixed number of connected pixels.

Our plan is as follows. Section 2 is devoted to the probabilistic model of noise in binary images: all pixels are independent, black with probability p or white with probability $1 - p$. The Poisson approximation result will be stated (Theorem 2.4) and an outline of its proof will be given (technical details will be postponed to the Appendix). Section 3 is devoted to applications. We will first explain how Theorem 2.4, together with numerical combinatorial results on square lattice animals[1], can be used to obtain an explicit formula for the size threshold $s(n, p, \varepsilon)$. An example of denoising for a binary image will be given. Then we shall extend the method to grey level images through threshold decomposition: the binary image corresponding to each grey level is treated separately, then all denoised binary images are recombined. Some experiments and a discussion of the obtained results come last.

2. Probability of a Local Property

Our probabilistic model for random images is the following. Let n be a positive integer. Consider the *pixel set* $\Xi_n = \{1, \ldots, n\}^2$. A *binary image of size n* is a mapping from Ξ_n to $\{0, 1\}$ (black/white). Their set is denoted by E_n. It is endowed with the probability distribution $\mu_{n,p}$ defined by: each pixel is black with probability p or white with probability $1 - p$, and all the pixel colors are independent. A *random image of size n and probability parameter p*, denoted by $\mathcal{I}_{n,p}$, is a random element of E_n with probability distribution $\mu_{n,p}$.

The pixel set Ξ_n is embedded in $\mathbb{Z}^2$ and naturally endowed with a *graph structure*. We consider in this paper the case of 4-connectivity (2 horizontal and 2 vertical neighbors). For purely technical reasons, it will be convenient that all pixels have the same neighborhood:

this is why we impose periodic boundary conditions, deciding that $(1, j)$ is a neighbor of (n, j) and $(j, 1)$ of (j, n). Thus the graph is a 2-dimensional torus. As usual, the *graph distance d* is defined as the minimal length of a path between two pixels. We shall denote by $B(x, r)$ the ball of center x and radius r with respect to the distance d. It is defined by

$$B(x, r) = \{y \in \Xi_n; \ d(x, y) \leq r\}.$$

Notice that this ball $B(x, r)$ is diamond-shaped (it is a rhombus) and that for $r < n/2$, it contains $2r^2 + 2r + 1$ pixels (see Fig. 1). For the rest of this section, the radius r is a fixed integer, and the image size n is larger than $2r + 1$.

The image properties we are interested in are all local, in the sense that they can be described inside balls of a fixed radius. All balls are translations of each other. We shall choose a ball of radius r, say $B(0, r)$, and fix a translation τ_x, from $B(0, r)$ to $B(x, r)$ for all x. We call *pattern*, and denote by D, an image defined on $B(0, r)$, and determined by its set of black pixels, denoted by $\beta(D)$ (see Fig. 1 for an example of pattern). Of course, $B(0, r) \setminus \beta(D)$ is the set of white pixels. We shall denote by $b(D)$ the cardinality of $\beta(D)$ (number of black pixels in the pattern). We shall deal with rather small levels of noise, seen as relatively sparse black pixels on a white background. This is of course a mere convention: swapping black and white, together with p and $1 - p$ does not change the model. Thus, in what follows, we will always assume that $p \leq \frac{1}{2}$.

If D is a pattern on $B(0, r)$ and τ is a translation of pixels, we shall denote by $\tau(D)$ the pattern on $B(\tau(0), r)$, whose set of black pixels is $\tau(\beta(D))$. If $\tau(0) = x$, we denote by $D(x)$ the property: "the restriction of the image to $B(x, r)$ is $\tau(D)$". The property


Figure 1. Example of an image on the ball $B(0, r)$ with $r = 3$. This small image is also called a *pattern*. The number of black pixels of this pattern D is $b(D) = 2$.

we are actually interested in is

$$\tilde{D} = (\exists x \in \Xi_n, D(x)).$$

In other words $\tilde{D}$ means: "a copy of pattern D can be found somewhere in the image".

The patterns D are the building blocks of all local properties. Indeed, there exists only a finite number of such patterns (precisely 2^{2r^2+2r+1}): let us denote their set by $\mathcal{D}$. Any *assertion* relative to the pixels in $B(0, r)$ will be called "local": it can be expressed in a unique way as a disjunction (logical "or", denoted by $\vee$) of distinct patterns.

The following definitions will be used in the counting of occurrences of a local property in an image.

Definition 2.1. Let ψ be a local assertion, relative to the pixels in $B(0, r)$.

1. The definition set of ψ, denoted by $\mathcal{D}(\psi)$, is the subset of $\mathcal{D}$ such that

$$\psi = \bigvee_{D \in \mathcal{D}(\psi)} D.$$

2. The black index $b(\psi)$ of ψ is the integer $b(\psi)$ defined by

$$b(\psi) = \min_{D \in \mathcal{D}(\psi)} \{b(D)\}.$$

3. A meaningful definition set of ψ, denoted by $\mathcal{D}_0(\psi)$, is a subset of $\mathcal{D}(\psi)$ such that

    (a) $\forall D \in \mathcal{D}_0(\psi), b(D) = b(\psi)$,
    (b) If τ is a translation, then $D, D' \in \mathcal{D}_0(\psi)$ and $\tau(\beta(D)) = \beta(D')$ imply $D = D'$,
    (c) $D \in \mathcal{D}(\psi)$ and $b(D) = b(\psi)$ imply $\exists \tau, \exists D' \in \mathcal{D}_0(\psi)$, s. t. $\tau(\beta(D)) = \beta(D')$.

    All meaningful definition sets have the same cardinality, which will be called the meaningful index of ψ, and denoted by $e(\psi)$.

The black index $b(\psi)$ is the minimal number of black pixels, in a pattern that satisfies ψ. One can see the meaningful index $e(\psi)$ as the maximal number of patterns with exactly $b(\psi)$ black pixels that satisfy ψ, up to possible translations. Both will be used to count occurrences of the *local property* based on ψ.

Example. Let us illustrate all these definitions by considering a simple example: the property "there exist two

connected black pixels". On the ball of radius $r = 1$, the definition set of this local assertion ψ is composed of all those patterns on $B(0, 1)$ whose center is black, and at least one of the 4 neighbors is also black (15 patterns). The black index $b(\psi)$ of ψ equals 2, and its meaningful index is $e(\psi) = 2$ (a possible meaningful definition set is made of the two patterns on $B(0, 1)$ such that the center and its right horizontal neighbor, resp. its top vertical neighbor, are the only black pixels in the ball $B(0, 1)$).

Definition 2.2. Let ψ be a local assertion, and $\psi(x)$ its localization on the ball centered at x:

$$\psi(x) = \bigvee_{D \in \mathcal{D}(\psi)} D(x).$$

We call local property based on ψ, and denote by $\tilde{\psi}$ the property

$$\tilde{\psi} = (\exists x, \psi(x)).$$

Our basic example of a local property $\tilde{\psi}$ is: "there exists a connected component of k black pixels". A connected component of size k is always included in a ball of radius $r \geq k/2$. The local assertion ψ is "there exists a connected component of size k in $B(0, r)$". The definition set is the set of all patterns on $B(0, r)$ having at least k connected black pixels. The black index is the minimal number of black pixels necessary for ψ to be satisfied (obviously k in our example). The meaningful index is the number of connected components of size k, up to translations (see Section 3).

For a fixed level p with $0 < p < 1$, if we let n tend to infinity, by the independence of pixels, it is easy to see that asymptotically any pattern will be present in a random image with a probability tending to 1 (see [3] for more precise results). Therefore the asymptotic probability for the random image $\mathcal{I}_{n,p}$ to satisfy $\tilde{\psi}$ is 1, whatever ψ. That asymptotic probability can be different from 1 only if $p = p(n)$ tends to 0 as n tends to infinity. Thus our images will have a relatively small proportion of black pixels.

A classical object of the theory of random graphs (see [1, 19] as general references), is the notion of *threshold function*. It describes the appearance of a given subgraph in a random graph. The notion of threshold function easily adapts to random images. Let $\mathcal{A}$ be an image property. The function $\theta(n)$ is called a threshold

function of $\mathcal{A}$ if for $p(n) \leq 1/2$ then

$$\lim_{n \to \infty} \frac{p(n)}{\theta(n)} = 0 \implies \lim_{n \to \infty} \mu_{n, p(n)}(\mathcal{A}) = 0 \,,$$

and

$$\lim_{n \to \infty} \frac{p(n)}{\theta(n)} = \infty \implies \lim_{n \to \infty} \mu_{n, p(n)}(\mathcal{A}) = 1 \,.$$

Notice that a threshold function is not unique. For instance if $\theta(n)$ is a threshold function for $\mathcal{A}$, then so is $c\theta(n)$ for any positive constant c. It is customary to ignore this and talk about "the" threshold function of $\mathcal{A}$. We then have the following lemma.

Lemma 2.3. *The threshold function of the local property $\tilde{\psi}$ is $n^{-\frac{2}{b(\psi)}}$.*

Proof: We shall just give here the main steps, since the detailed proof will appear in [3]. Let D be a pattern and let X_n denote the number of copies of D in the image. Then

$$\mu_{n, p(n)}(\tilde{D})$$
$$= \mu_{n, p(n)}(X_n > 0) \leq \mathbb{E}(X_n)$$
$$= n^2 p(n)^{b(D)} (1 - p(n))^{2r^2 + 2r + 1 - b(D)} \leq n^2 p(n)^{b(D)}.$$

On the other hand, let Y_n denote the number of copies of D occurring in balls $B(x, r)$ where both coordinates of x are multiples of $2r + 1$ (which implies that two such balls cannot meet). The number of such balls is n_r^2 where $n_r = \lfloor \frac{n}{2r+1} \rfloor$. Then

$$\mu_{n, p(n)}(X_n > 0)$$
$$\geq \mu_{n, p(n)}(Y_n > 0)$$
$$= 1 - \mu_{n, p(n)}(Y_n = 0)$$
$$= 1 - \left(1 - p(n)^{b(D)} (1 - p(n))^{2r^2 + 2r + 1 - b(D)}\right)^{n_r^2}$$
$$\geq 1 - \exp\left(-n_r^2 p(n)^{b(D)} \left(1 - p(n)\right)^{2r^2 + 2r + 1 - b(D)}\right)$$

Using these inequalities and the definition of a threshold function, we conclude that $\theta(n) = n^{-2/b(D)}$ is the threshold function of the property $\tilde{D}$. To conclude, one has to check that the threshold function of a disjunction of patterns is the smallest threshold function of these patterns. $\square$

Lemma 2.3 means that the appearance of a local property mainly depends on its black index: if $p(n)$

is small compared to $n^{-\frac{2}{b}}$, then the probability of any local property that needs b black pixels to be satisfied is small. If $p(n)$ is large compared to $n^{-\frac{2}{b}}$, then the probability is large. The particular case $b(\psi) = 0$ corresponds to the appearance of a white ball. If there exists α, with $0 \leq \alpha < 1$, such that for all n, we have $p(n) \leq \alpha$, then the probability for a white ball of being present in the random image always tends to 1 as n tends to infinity: there is no threshold function. From now on, we will always assume that the black index of ψ is positive.

Lemma 2.3 suggests that the correct scaling for $p(n)$ when one studies a local property $\tilde{\psi}$ is $p(n) = cn^{-\frac{2}{b(\psi)}}$. Our main result shows that with this scaling, the probability of $\tilde{\psi}$ in a random image converges to a non trivial limit.

Theorem 2.4. *Let ψ be an assertion on $B(0, r)$, with black index $b(\psi)$ and meaningful index $e(\psi)$. Let $p(n) = cn^{-\frac{2}{b(\psi)}}$, where c is a positive constant. Then*

$$\lim_{n \to \infty} \mu_{n, p(n)}(\tilde{\psi}) = 1 - \exp(-e(\psi) c^{b(\psi)}) \,. \quad (1)$$

The reason why such a result is called a Poisson approximation becomes clear if one considers the property "there exists a black pixel". Let X_n be the total number of black pixels. Since all pixels are independent, the random variable X_n follows the binomial distribution with parameters n^2 and $p(n)$. In particular the probability that there exists a black pixel is

$$\mathbb{P}[X_n > 0] = 1 - (1 - p(n))^{n^2} \,.$$

Here the black index is 1 and the threshold function is n^{-2}. Take $p(n) = cn^{-2}$. Then the binomial distribution of X_n converges to the Poisson distribution with parameter c, and the probability that there exists a black pixel ($X_n > 0$) tends to $1 - \exp(-c)$.

The situation is not so simple as soon as the black index is larger than 1. Consider for instance again the local property $\tilde{\psi}$: "there exist two connected black pixels". We already saw that on the ball of radius $r = 1$, the definition set is composed of all those patterns on $B(0, 1)$ whose center is black, and at least one of the 4 neighbors is also black (15 patterns). Consider the number of occurrences of any of those patterns, somewhere in the random image. It is a sum of Bernoulli random variables. However they are not independent:

patterns on balls centered at two adjacent pixels have one pixel in common. The same can be said of any local property $\tilde{\psi}$: the number of occurrences of $\psi(x)$ can be viewed as a sum of (dependent) Bernoulli random variables. The sum of a large number of Bernoulli r.v.'s converges in distribution to a Poisson distribution, provided the dependencies between the variables are not too large. In the theory of random graphs, similar results are frequent (see e.g. [19] Lecture 1 p. 296, Lecture 2 p. 303 or Lecture 5 p. 314).

Proof of Theorem 2.4: There are several ways to prove a Poisson approximation result. We chose the famous "moment method" based on the following result ([1], Chapter 1, p. 25). $\square$

Lemma 2.5. *Let $(X_n)_{n\in\mathbb{N}^*}$ be a sequence of integer valued, nonnegative random variables and λ be a strictly positive real. For all $n, l \in \mathbb{N}^*$ define the quantity*

$$E_l(X_n) = \sum_{k \geq l} \mathbb{P}(X_n = k) \frac{k!}{(k-l)!}.$$

If, for all $l \in \mathbb{N}^*$, $\lim_{n\to\infty} E_l(X_n) = \lambda^l$ then (X_n) converges in distribution to the Poisson distribution with parameter λ.

In our case, X_n counts the number of occurrences in the random image of some patterns, to be precised later. The "moment" $E_l(X_n)$ is the expected number of ordered l-tuples of occurrences of those patterns.

Firstly, one should observe that patterns in the definition set of ψ cannot be all treated equally: since $p(n) = cn^{-\frac{2}{b(\psi)}}$, by Lemma 2.3 any pattern with more than $b(\psi)$ black pixels has a vanishing probability of being observed. Hence we can reduce the set of patterns to those having exactly $b(\psi)$ black pixels. In the example of two connected pixels with $r = 1$, $\mathcal{D}(\psi)$ has 15 different patterns, but only 4 of them have exactly 2 black pixels.

Now, one has to take care of multiple counts. Among the 4 patterns on $B(0, 1)$ that have 2 black pixels, 2 patterns have two horizontal black neighbors, and the 2 other patterns have two vertical black neighbors. Assume the image has only one occurrence of two horizontal black neighbors. If we examine all possible pixels x, we will find two adjacent centers for which $\psi(x)$ is satisfied. In order to obviate this problem, we need to count patterns *up to possible translations*. We say

that two patterns with black index $b(\psi)$ are equivalent if their sets of black pixels are translations of each other. The number of equivalence classes is the meaningful index $e(\psi)$ of Definition 2.1. (In the example of two connected black pixels, there are two equivalence classes: horizontal or vertical neighbors).

We choose a meaningful set, i.e. we fix a pattern for each equivalence class:

$$\mathcal{D}_0(\psi) = \{\bar{D}_1, \ldots, \bar{D}_{e(\psi)}\}.$$

The counting variable X_n to which Lemma 2.5 will be applied is the total number of occurrences of one of the patterns $\bar{D}_1, \ldots, \bar{D}_{e(\psi)}$, in the random image $\mathcal{I}_{n,p(n)}$:

$$X_n = \sum_{x \in \Xi_n} \sum_{i=1}^{e(\psi)} \mathbb{I}_{\bar{D}_i(x)}\big(\mathcal{I}_{n,p(n)}\big),$$

where $\mathbb{I}$ denotes the indicator function of an event. The expectation of X_n is

$$\mathbb{E}(X_n) = n^2\, e(\psi)\, (p(n))^{b(\psi)}(1 - p(n))^{2r^2+2r+1-b(\psi)}.$$

As n tends to infinity, it tends to $e(\psi)\, c^{b(\psi)}$, which is the parameter of the Poisson approximation in formula (1). In order to apply Lemma 2.5 to X_n, one has to check that the hypothesis holds.

Lemma 2.6.

$$\forall l \in \mathbb{N}^*, \qquad \lim_{n\to\infty} E_l(X_n) = \big(e(\psi)c^{b(\psi)}\big)^l.$$

The proof of Lemma 2.6 is rather technical and will be given in the Appendix.

Now Lemma 2.5 implies that X_n converges in distribution to the Poisson distribution with parameter $e(\psi)c^{b(\psi)}$. Therefore $\mu_{n,p(n)}(X_n > 0)$ tends to $1 - \exp(-e(\psi)c^{b(\psi)})$. It is clear that $X_n > 0$ implies that $\mathcal{I}_{n,p(n)}$ satisfies $\tilde{\psi}$. Hence $\mu_{n,p(n)}(X_n > 0) \leq \mu_{n,p(n)}(\tilde{\psi})$. Moreover, the event $(\tilde{\psi} \setminus (X_n > 0)) = (\tilde{\psi} \cap (X_n = 0))$ implies the appearance of a pattern with at least $b(\psi) + 1$ black pixels in a ball of radius r, and by Lemma 2.3, its probability tends to 0 as n tends to infinity. Therefore,

$$\lim_{n\to\infty} \mu_{n,p(n)}(X_n > 0) = \lim_{n\to\infty} \mu_{n,p(n)}(\tilde{\psi})$$
$$= 1 - \exp\big(-e(\psi)c^{b(\psi)}\big).$$

It should be noticed that the asymptotics of X_n does not depend on the choice of the meaningful definition set $\{\bar{D}_1, \ldots, \bar{D}_{e(\psi)}\}$. It does not depend either on the radius r of the ball. Consider for instance the property $\tilde{\psi}$ "the image contains two horizontally connected black pixels". Its definition set for the ball $B(0, r)$ has $r^2 2^{2r^2+2r}$ elements. Among these, only $2r^2$ have exactly 2 black pixels, and there is only one equivalence class up to translations, whatever r. Therefore r is a phantom parameter, as should be expected. It serves only to ensure that properties remain local.

3. Application to Image Denoising

In the previous section, we computed the asymptotic probability of appearance of any local property in a random binary image. This provides the basis of a statistical test to decide whether an observed pattern in an image may be due to noise or not, and this test can be applied for image denoising. In this section, all the considered images will be corrupted by the same kind of noise, namely impulse noise. This type of noise models for example the fact that some (unknown) part of the data is lost. We will assume that the probability parameter of the noise is known. We will first start with the denoising of binary images, and then extend it to grey level images using their threshold decomposition.

3.1. Binary Images

Let I_0 be the original (non degraded) binary image of size $n \times n$. This original image I_0 is then corrupted by impulse noise, which has a probability parameter p in the white components and a probability parameter q in the black ones (see Fig. 3 for an example). We shall see in the next section why it is important to allow black and white pixels to be destroyed with a different probability. Thus the noisy image I is given by

$$\forall x, \ I(x) = I_0(x) \cdot (1 - \zeta_p(x))$$
$$+ (1 - I_0(x)) \cdot \zeta_q(x), \qquad (2)$$

where the $\zeta_p(x)$'s (resp. $\zeta_q(x)$'s) are independent Bernoulli random variables with parameter p (resp. q). In other words, we have the following conditional probabilities

$$\mathbb{P}(I(x) = 0 \mid I_0(x) = 1) = p \quad \text{and}$$
$$\mathbb{P}(I(x) = 1 \mid I_0(x) = 0) = q.$$

As can be seen in Fig. 3, the impulse noise creates small black and white connected components. These small components will be removed using a statistical decision based on their size ("size", in this paper, always means "area"). We are first interested in the black connected components (with respect to 4-connectivity). The results of the previous section give us the threshold function and also the probability of appearance of such components. More precisely, the threshold function for a given (fixed) black component of size k is $\theta(n) = n^{-2/k}$ and its asymptotic appearance probability in a $n \times n$ image of noise with probability parameter $p(n) = c\theta(n)$, as n goes to infinity, is equal to

$$1 - e^{-c^k}.$$

Now, if we are interested in the appearance of a component of size k (i.e. any of them, not only a given one), Theorem 2.4 claims that the asymptotic (for large n) probability of appearance is

$$1 - e^{-a_k c^k},$$

where a_k is the number of 4-connected components one can make with exactly k pixels (up to translations). Writing $c^k = n^2 p^k$, we thus have an approximation for the probability of appearance of a component of size k in the $n \times n$ image, with a proportion p of black pixels. We denote by $\text{PA}(n, k, p)$ this approximation :

$$\text{PA}(n, k, p) = 1 - e^{-n^2 a_k p^k}.$$

The 4-connected components are known in the combinatorics literature as "square lattice animals" or "polyominoes". Counting these objects is a difficult combinatorial problem and there is no general expression for a_k. However, some asymptotic results are known: a concatenation argument [10] shows that there exists a constant a, called *growth constant*, such that:

$$\lim_{k \to \infty} (a_k)^{\frac{1}{k}} = \sup_{k \geq 1} (a_k)^{\frac{1}{k}} = a.$$

The exact value of a is unknown. Numerical estimates give $a \simeq 4.06$ and the best published rigorous bounds for it are $3.9 < a < 4.65$ (see [2, 9, 11]). But thanks to some numerical studies[2], the values of the sequence $(a_k)_{k \geq 1}$ are known up to $k = 47$, which will be enough in practice for denoising applications. The first terms are: $a_1 = 1, a_2 = 2, a_3 = 6, a_4 = 19$, etc. Furthermore,

Figure 2. The size threshold $s(n, p, \varepsilon)$ as a function of the noise probability parameter $p \in [0, 0.18]$, for $n = 256$ and $\varepsilon = 10^{-1}$, 10^{-2} and 10^{-3}.

numerical computations show that for $p \leq p_{\max} \simeq 0.2$, one has $a_{k+1} p \leq a_k$ for $k \in [1, 47]$, which ensures that $\text{PA}(n, k + 1, p) \leq \text{PA}(n, k, p)$. This means that the probability of appearance of an animal is a decreasing function of its size. This is rather reasonable: for fixed values of p and n, it would not make much sense to keep a connected component of size k and to remove one of size $k' > k$.

Let us fix a (small) positive real ε which will be our risk probability, in the sense of statistical testing. If the size k of a connected component observed in a noisy image I is such that $\text{PA}(n, k, p) \leq \varepsilon$, then we will consider that it comes from the original image I_0, and keep it. If $\text{PA}(n, k, p) > \varepsilon$, it will be regarded as noise and removed. Thus the size threshold for the components we keep is defined by:

$$s(n, p, \varepsilon) = \inf\left\{k \,;\, \text{PA}(n, k, p) = 1 - e^{-n^2 a_k p^k} \leq \varepsilon\right\}. \tag{3}$$

A component with size $k \geq s(n, p, \varepsilon)$ has a very low probability (less than ε) of appearing in a pure noise image. In Fig. 2(b), we plot the size threshold $s(n, p, \varepsilon)$ as a function of the noise probability parameter p, for a fixed value of $n = 256$ and three different values of ε: 10^{-1}, 10^{-2} and 10^{-3}.

The algorithm for the binary image denoising can be decomposed in four steps:

1. Extract all the 4-connected black components of the noisy image I.
2. Remove the ones which have a size less than $s(n, p, \varepsilon)$ (i.e. change their pixels into white). Obtain a new binary image.
3. Extract all the white 4-connected components of this new image.
4. Remove the ones which have a size less than $s(n, q, \varepsilon)$ (i.e. change their pixels into black), to obtain the final denoised image denoted by $\tilde{I} = TI$.

To summarize, this denoising filter T can be written as:

$$T = T_{s(n,q,\varepsilon)}^+ \circ T_{s(n,p,\varepsilon)}^-,$$

where T_s^+ (resp. T_s^-) is the morphological area opening (resp. closing) of size s defined by L. Vincent in [20]. See Fig. 3 for an example of the obtained result and for a comparison with the results obtained with a more standard binary filter (namely the median filter).

Before explaining how this method will be extended to grey level images, let us make a few general comments.

- The method is valid when p is not too large, since we need $a_k p^k$ to be small. In practice, we are limited to $p \leq p_{\max} \simeq 0.2$.
- The dependence on ε is low since it is in fact a $\log(\varepsilon)$-dependence. Indeed, $1 - e^{-n^2 a_k p^k}$ is approximately

Figure 3. First row: on the left, the original binary image I_0 of size 256×256; on the right, the corrupted image I. White pixels have been changed with probability $p = 0.1$ and black ones with probability $q = 0.2$. Second row: on the left, the result of the denoising algorithm with $\varepsilon = 10^{-2}$ when removing first black components and then white ones (i.e. applying $T^+_{s(n,q,\varepsilon)} \circ T^-_{s(n,p,\varepsilon)}$); on the right, denoising by first removing white components and then black ones (i.e. applying $T^-_{s(n,p,\varepsilon)} \circ T^+_{s(n,q,\varepsilon)}$). The two images are not the same, illustrating the fact that the two operators T^-_s and T^+_s do not commute. However in both cases, small black and white components due to noise have been removed. Only the boundaries of the remaining ones are different. Third row: results obtained when applying median filtering with a disk of radius $r = 2$ (on the left) and a disk of radius $r = 5$ (on the right). The value of the parameter $r = 2$ seems too small since some noise is still present in the black components. Now, for the value $r = 5$, one can notice that the black corners have been eroded. The reason for this is that, in the noisy image, the probability parameters p and q (used to corrupt respectively the white and black pixels) were such that $q > p$.

equal to $n^2 a_k p^k$ when the value of this expression is small. If we replace a_k by a^k, the threshold for the minimal size of the components we keep is approximately given by

$$s(n, p, \varepsilon) \simeq \frac{\log \varepsilon - 2 \log n}{\log a + \log p} .$$

- The boundaries of the remaining components are not smoothed. This comes from the fact that when some noise is at the boundary of a component, it becomes part of it. In order to remove it, one would need other *a priori* knowledge of the original image (such as smooth or straight boundaries as in the case of Fig. 3). It is actually a general problem of image denoising:

one has to define some *a priori* model for the image. Here the underlying model is that the original binary image is made of "large" (as compared to the noise) black and white connected components.

- The two filters T_s^- and T_s^+ do not commute (see Fig. 3). This was already noted by Vincent in [20]. One solution he proposed is to use them in alternating sequential filters [17, 18] with increasing sizes of area. This may not be a real issue, since he also noticed that $T_s^- \circ T_s^+$ and $T_s^+ \circ T_s^-$ are visually extremely close (this will be even more true for grey level images). Another solution, proposed by Masnou and Morel in [12], and then formalized by Monasse [13], is to process simultaneously upper and lower level sets. This grain filter denoted G_t (where t is the area threshold) is done by a pruning of the tree of all level sets, built thanks to the inclusion principle (this algorithm which is very fast is called the Fast Level Set Transform [14]).

  The fact that the foreground and the background of a binary image are treated in a complementary way is a general problem in Mathematical Morphology. Many operators are not self-dual, and they often occur pairwise: like dilation/erosion and opening/closing for example. It is worth mentioning that in [7], H. Heijmans describes a general method to construct morphological operators which are self-dual.

- It is generally considered that, for consistency reasons, using the 4-connectivity on the black (or white) pixels should be followed by using the 8-connectivity for the complementary set. From a theoretical point of view, the method we proposed can be extended to 8-connectivity in a straightforward way. In order to apply the method, one would have to count the number of 8-connected components of size k, which are not available in the literature, whereas the (a_k)'s are known up to $k = 47$. Consequently, for our application to image denoising, we decided to treat the foreground and the background in the same way, with 4-connectivity.

- If the original image I_0 is all white ($I_0 \equiv 0$), and if it is corrupted by some noise with probability parameter p as described by Eq. (2), we obtain an image I which is "pure noise". The probability that it contains a connected component with size larger than $s(n, p, \varepsilon)$ is (by definition of $s(n, p, \varepsilon)$ and thanks to Theorem 2.4) less than ε. Thus,

$$\mathbb{P}\big(T_{s(n,p,\varepsilon)}^- I = I_0\big) \geq 1 - \varepsilon,$$

which means that, with probability larger than $1 - \varepsilon$, pure noise is completely removed.

Thus ε represents the "significance level" of our statistical method for denoising: the probability of not removing a component coming from the noise is less than ε. In practice, we generally take $\varepsilon = 10^{-3}$ or 10^{-2} (the results are visually the same). This parameter ε is completely independent of the image (which is not the case of the size n of the image, or the probability parameter p of the impulse noise): ε has to be fixed by the user in the same way as the risk level in statistical hypothesis testing.

- For reasons of simplicity, the denoising filter T was described in the framework of an image of size $n \times n$. The method extends straightforwardly to an image of size $m \times n$, where $m \neq n$, by simply changing the size threshold $s(n, p, \varepsilon)$ into $\tilde{s}(m, n, p, \varepsilon)$ defined by

$$\tilde{s}(m, n, p, \varepsilon) = \inf\big\{k \,;\, \mathrm{PA}(\sqrt{nm}, k, p)$$
$$= 1 - e^{-mna_k p^k} \leq \varepsilon\big\}.$$

3.2. *Grey Level Images*

Let u be a grey level image, of size $n \times n$ and grey level values in the range $[0, 255]$. Assume that this image is corrupted by impulse noise with probability parameter p. This means that the observed noisy image v may be written in the form:

$$\forall x, \ v(x) = (1 - \zeta_p(x)) \cdot u(x) + \zeta_p(x) \cdot \nu(x), \quad (4)$$

where the $\zeta_p(x)$'s are independent Bernoulli random variables with parameter p and the $\nu(x)$'s are i.i.d.r.v.'s, uniformly distributed on $[0, 255]$.

For each level $\lambda \in [0, 255]$, we can consider the thresholded images $u_\lambda = \mathbb{I}_{u \geq \lambda}$ and $v_\lambda = \mathbb{I}_{v \geq \lambda}$. The grey level images may then simply be recovered by $u = \sum_\lambda u_\lambda$ and $v = \sum_\lambda v_\lambda$. The binary noisy image v_λ is a corrupted version of the binary image u_λ; they are related by

$$\mathbb{P}(v_\lambda(x) = 0 \mid u_\lambda(x) = 1) = p \times \frac{\lambda}{256} \quad \text{and}$$
$$\mathbb{P}(v_\lambda(x) = 1 \mid u_\lambda(x) = 0) = p \times \left(1 - \frac{\lambda}{256}\right).$$

We are thus back in the framework described for binary images with parameters $p_\lambda = p\lambda/256$ and $q_\lambda = p(1 - \lambda/256)$. The image v_λ can be denoised following the

Figure 4. Left: image v obtained with impulse noise with probability parameter $p = 0.15$ on the Lena image. Middle: thresholded image v_λ for the grey level $\lambda = 150$. Right: denoised image $\tilde{v}$ obtained by the noise-adapted grain filter T with $\varepsilon = 10^{-3}$.

method described in the previous subsection. Finally, we reconstruct a grey level image by simply adding the binary ones: $\tilde{v} = \sum_\lambda \tilde{v}_\lambda$. This can be summarized by the formula

$$\tilde{v} = Tv = \sum_{\lambda=0}^{255} T^+_{s(n,q_\lambda,\varepsilon)} \circ T^-_{s(n,p_\lambda,\varepsilon)}(v_\lambda), \quad \text{where}$$

$$p_\lambda = p\frac{\lambda}{256} \quad \text{and} \quad q_\lambda = p\left(1 - \frac{\lambda}{256}\right). \tag{5}$$

Figures 4 and 5 give two examples of results obtained by this filtering.

One natural question that can be asked is whether the filter T defined by formula (5) is a morphological filter (see [16, 18] and references therein for the definition and properties of morphological filters). Unfortunately, the answer is negative. For two grey levels $\lambda \geq \lambda'$, one has $v_\lambda \leq v_{\lambda'}$, and for a fixed area threshold t one would have $T^-_t(v_\lambda) \leq T^-_t(v_{\lambda'})$ (because area openings and closings are morphological operators). Now, the two thresholds $s(n, p_\lambda, \varepsilon)$ and $s(n, p_{\lambda'}, \varepsilon)$ can be different, i.e. $s(n, p_\lambda, \varepsilon) > s(n, p_{\lambda'}, \varepsilon)$ and thus it is not necessarily true that $T^-_{s(n,p_\lambda,\varepsilon)}(v_\lambda) \leq T^-_{s(n,p_{\lambda'},\varepsilon)}(v_{\lambda'})$. This happens when v_λ and $v_{\lambda'}$ both contain the same small black connected component of size k such that $s(n, p_\lambda, \varepsilon) > k > s(n, p_{\lambda'}, \varepsilon)$. However, in the experimental results, we noticed that this rarely happens: for most values of λ, one has $\tilde{v}_\lambda \leq \tilde{v}_{\lambda-1}$.

In order to illustrate the interest of adapting the area threshold to each grey level, we treated the same image using our method, then using a fixed area threshold (for this we used the algorithm developed by Monasse in [13]). The results are those of Figs. 5 and 6. Figure 6

shows the result of the usual grain filter, denoted by G_t, for two different values of the area threshold: $t = 10$ and $t = 20$. One can notice that the parameter value $t = 10$ seems too low since there is still some remaining noise (for example on the coat of the cameraman). On the other hand the value $t = 20$ seems too large, since some of the original structures have disappeared (for instance the white parabola at the top of the building) and still too low (there is some remaining noise on the coat). These results have to be compared with the one of Fig. 5(c). This last figure shows that thanks to the adapted area threshold $s(n, p_\lambda, \varepsilon)$ a small white component can be kept and at the same time, a larger grey component removed. These results also illustrate what we have proposed in this paper, namely an adapted and automatic way to choose the right parameter for the area openings and closings.

3.3. Extension to Other Noise Models

In the previous subsection, we have explained how the theoretical results of Section 2 can be used to denoise an image degraded by impulse noise. Now, even if the proposed denoising procedure corresponds to an impulse noise model, it is interesting to see how it works in the presence of white noise. An example of the obtained results is shown on Fig. 7: we used again the cameraman image, which is here degraded by white noise with standard deviation $\sigma = 15$. It is then denoised using the filter defined by Eq. (5) with parameter value $p = 0.2$ and $\varepsilon = 10^{-2}$ (on the middle image) and $p = 0.1$ and $\varepsilon = 10^{-2}$ (on the right image). The main question is here: how to choose the value of the parameter p used

Figure 5.   From left to right, top to bottom: (a) the original cameraman image u (size 256×256); (b) degraded image v, with impulse noise probability parameter $p = 0.2$; (c) filtered image $\tilde{v}$, obtained with $\varepsilon = 10^{-3}$; (d) image of the difference $u - \tilde{v}$. It shows that most of the noise has been removed, except at the boundaries of the objects and also in the grass texture.

Figure 6.   Result of the filtering of the noisy image v with the usual grain filter G_t with area threshold $t = 10$ on the left and $t = 20$ on the right.

in the filter, in relation to the standard deviation σ of the white noise ?

In the case of impulse noise, we were able to relate the size threshold s of the area openings and closings to the impulse noise probability parameter p and to the grey level λ. The main result was then: if we take $u = 0$ in Eq. (4), then the degraded image v is pure impulse noise, and after filtering (Eq. (5)), we have, by definition of the threshold $s(n, p, \varepsilon)$, that $Tv = u$ with probability larger than $1 - \varepsilon$.

Figure 7. From left to right: (a) the cameraman image degraded by white noise with standard deviation $\sigma = 15$; (b) denoising of the previous image by the filter defined by Eq. (5) with parameter values $p = 0.2$ and $\varepsilon = 10^{-2}$; (c) denoising by the same filter with parameter values $p = 0.1$ and $\varepsilon = 10^{-2}$.

Now, if we want to obtain in the same way a denoising filter for white noise, we have first to be able to relate the size threshold s used for the area openings and closings to the standard deviation σ of the white noise and to the grey level λ of the thresholded image. In order to do this, let us consider a pure white noise image w: all the $w(x)$'s are independent identically distributed random variables with distribution $\mathcal{N}(0, \sigma^2)$ (gaussian with mean 0 and variance σ^2). For $\lambda \in \mathbb{R}$, let us consider the thresholded image $w_\lambda = \mathbb{I}_{w \geq \lambda}$. We then have

$$\mathbb{P}(w_\lambda(x) = 1) = \int_\lambda^{+\infty} \frac{1}{\sigma\sqrt{2\pi}} e^{-x^2/2\sigma^2} \, dx.$$

This last term, denoted by $\tilde{p}_{\lambda,\sigma}$ should be the analogue of the probability parameter p_λ defined in the case of impulse noise (Eq. (5) in the previous subsection). Now the main difference here is that $\tilde{p}_{\lambda,\sigma}$ is not necessarily small (it goes to 1 as λ goes to $-\infty$), and the thresholded image w_λ may contain arbitrarily large connected components. Thus a filter like the area opening or closing will never be able to remove all the noise. The problem here is that the type of filter we have considered is not adapted to white noise.

Generally, when using a probabilistic approach for filtering, one needs a model for the image and one for the noise. Here, we do not need a model for the image, since we only use an "a contrario" hypothesis. It means that we only need to know that "the image is not noise" in the sense that large connected components, which have a very small probability of appearing in impulse noise, necessarily belong to the image. This approach does not work in the case of white noise since the size of the connected components of level

sets is not a good way to discriminate white noise from the image (both contain large components). Nevertheless, if we are able to find some characteristic geometric features (as the size of connected components in the case of impulse noise), the proposed approach could be extended to white noise or to other models of noise.

4.   Conclusion

We have introduced a mathematical model for random images, in which we were able to compute the probability of appearance of any "local pattern" (Theorem 2.4). This was then used to give an explicit formula for the size threshold $s(n, p, \varepsilon)$, such that the probability of appearance of a component of size $k \geq s(n, p, \varepsilon)$ in a $n \times n$ image of pure noise with probability parameter p is less than ε. Using this value of $s(n, p, \varepsilon)$ for the area openings and closings defined by Vincent will ensure that, with probability larger than $1 - \varepsilon$, pure noise is completely removed. This denoising process was then extended to grey level images using their threshold decomposition. There, the proposed area threshold depends on both the probability parameter p of the impulse noise and the grey level λ of the level set.

Now, some questions remain, that have not been addressed in this paper: if the probability parameter p of the impulse noise is unknown, what is the best way to estimate it ? For a binary pure noise image, the best estimate of p is simply the ratio of the number of black pixels to the area of the image. Then, by analogy, a first answer for binary images (like the chessboard for example) is to compute the relative number of black

pixels outside a dilation of the "large" black components. Thus, for a grey level image, it is possible to use the threshold decomposition to obtain initial estimates of $p_\lambda = p \times \lambda/256$ and then to estimate p using, for example, a linear regression. Now, it is not clear that this estimate will be a good one since natural grey level images often contain textures creating small components which over-estimate p. In order to obtain a reliable estimate of p, it would be necessary to use also some information extracted from the statistical moments (like the covariance, three-point probability, etc...) measured on the image.

Appendix A

Proof of Lemma 2.6: Fix $l \in \mathbb{N}^*$. Recall that X_n counts the number of occurrences of the meaningful patterns $\bar{D}_1, \ldots, \bar{D}_{e(\psi)}$ in the random image $\mathcal{I}_{n,p(n)}$ where $p(n) = cn^{-\frac{2}{b(\psi)}}$. We are interested in:

$$E_l(X_n) = \sum_{k \geq l} \mathbb{P}(X_n = k) \frac{k!}{(k-l)!} .$$

We need to prove that $E_l(X_n)$ tends to $(e(\psi)c^{b(\psi)})^l$ as n tends to infinity.

One can see $E_l(X_n)$ as the average number of ordered l-tuples of copies of the patterns $\bar{D}_1, \ldots, \bar{D}_{e(\psi)}$ in $\mathcal{I}_{n,p(n)}$. Thus, we can write:

$$E_l(X_n)$$
$$= \mathbb{E}\left(\sum_{\substack{x_1, \ldots, x_l \\ x_i \neq x_j}} \sum_{1 \leq j_1, \ldots, j_l \leq e(\psi)} \mathbb{I}_{\bar{D}_{j_1}(x_1) \wedge \cdots \wedge \bar{D}_{j_l}(x_l)}\left(\mathcal{I}_{n,p(n)}\right) \right)$$
$$= \sum_{s=1}^{l} \sum_{\substack{(x_1, \ldots, x_l) \\ \in \mathcal{C}(s)}} \sum_{\substack{1 \leq j_1, \ldots, j_l \\ \leq e(\psi)}} \mu_{n,p(n)}\left(\bar{D}_{j_1}(x_1) \wedge \cdots \wedge \bar{D}_{j_l}(x_l) \right),$$

where, for $s = 1, \ldots, l$, $\mathcal{C}(s)$ represents the set of l-tuples $(x_1, \ldots, x_l)$ of pixels in Ξ_n such that the set $\{B(x_1, r), \ldots, B(x_l, r)\}$ is composed of s equivalence classes for the 4-connectivity relation.

The term corresponding to $s = l$ in the last sum will be denoted by $E'_l(X_n)$ and the rest by $E''_l(X_n)$. The quantity $E'_l(X_n)$ can be seen as the average number of ordered l-tuples of copies of $\bar{D}_1, \ldots, \bar{D}_{e(\psi)}$, on non-overlapping balls. We will first show

that:

$$\lim_{n \to \infty} E'_l(X_n) = \left(e(\psi)c^{b(\psi)} \right)^l . \qquad (6)$$

Then we will prove that $E''_l(X_n)$ tends to 0 as n tends to infinity.

We want to choose l pixels $x_1, \ldots, x_l$ such that the balls of radius r centered on those pixels are two by two disjoint. For the first pixel x_1, there are n^2 possibilities. Let $2 \leq j \leq l$ and suppose pixels $x_1, \ldots, x_{j-1}$ have been chosen. For the j-th choice, the set of all pixels x such that $d(x, x_k) \leq 2r$ for some $1 \leq k \leq j-1$, must be avoided. The cardinality of this set is bounded by $(j-1) \times (8r^2 + 4r + 1)$ whatever $x_1, \ldots, x_{j-1}$. This bound does not depend on n. So, asymptotically the number of choices for the j-th element is n^2, and consequently the cardinality of $\mathcal{C}(l)$ is equivalent to n^{2l}. On the other hand, if two balls $B(x, r)$ and $B(x', r)$ are disjoint, then for all $1 \leq j, j' \leq e(\psi)$, the random variables $\mathbb{I}_{\bar{D}_j(x)}$ and $\mathbb{I}_{\bar{D}_{j'}(x')}$ are independent. Therefore, we obtain the first limit (relation (6)):

$$E'_l(X_n) \sim n^{2l} \left(e(\psi)p(n)^{b(\psi)}(1 - p(n))^{2r^2+2r+1-b(\psi)} \right)^l$$
$$\sim \left(e(\psi)c^{b(\psi)} \right)^l.$$

The factor $e(\psi)^l$ comes from the choice of the $e(\psi)$ patterns $\bar{D}_1, \ldots, \bar{D}_{e(\psi)}$ for the l chosen balls.

There remains to prove that $E''_l(X_n)$ tends to 0 as n tends to infinity. The intuition is that if two patterns occur in overlapping balls, then locally more than $b(\psi)$ black pixels are present in a ball of radius $2r$. This has vanishing probability, by Lemma 2.3.
Let $1 \leq s \leq l-1$ and $(x_1, \ldots, x_l)$ be an element of $\mathcal{C}(s)$. Let $C_1, \ldots, C_s$ represent the connected components of the set $\cup_{k=1}^{l} B(x_k, r)$. Then by independence between them (they concern disjoint pixel sets):

$$\mu_{n,p(n)}\left(\bar{D}_{j_1}(x_1) \wedge \cdots \wedge \bar{D}_{j_l}(x_l) \right)$$
$$= \prod_{m=1}^{s} \mu_{n,p(n)}\left(\bigwedge_{k; B(x_k,r) \in C_m} \bar{D}_{j_k}(x_k) \right).$$

As a consequence of $s \leq l-1$, there exists at least one connected component, say C_1, having at least two elements. Since the black pixel sets of two different patterns of $\mathcal{D}_0(\psi)$ cannot be translated of each other, there must be at least $b(\psi)+1$ black pixels in C_1. Thus

we have

$$\mu_{n,p(n)}\left(\bigwedge_{k;B(x_{k,})\in C_1}\bar{D}_{j_k}(x_k)\right)\leq p(n)^{b(\psi)+1}.$$

For the other connected components, we simply bound

$$\mu_{n,p(n)}\left(\bigwedge_{k;B(x_k,r)\in C_m}\bar{D}_{j_k}(x_k)\right)\leq\mu_{n,p(n)}\left(\bar{D}_{j_{k_m}}(x_{k_m})\right)$$
$$\leq p(n)^{b(\psi)},$$

for any index k_m such that $B(x_{k_m},r)\in C_m$. Therefore, we obtain the following result:

$$\mu_{n,p(n)}\left(\bar{D}_{i_1}(x_1)\wedge\cdots\wedge\bar{D}_{i_l}(x_l)\right)\leq p(n)^{sb(\psi)+1}.$$

Finally, the set $\mathcal{C}(s)$ only has $O(n^{2s})$ elements and the number of ways to choose l elements among $\bar{D}_1,\ldots,\bar{D}_{e(\psi)}$ does not depend on n. Consequently, the desired result follows:

$$E_l''(X_n)\leq\sum_{s=1}^{l-1}O\left(n^{2s}\times n^{-\frac{2(sb(\psi)+1)}{b(\psi)}}\right)=\sum_{s=1}^{l-1}O\left(n^{-\frac{2}{b(\psi)}}\right)$$
$$=o(1)\,.\qquad\square$$

Acknowledgments

Notes

1. Square lattice animals or "polyominoes" are simply defined as connected clusters of squares in the plane (for example, the "Tetris" game uses all lattice animals of size 4).
2. for up-to-date information on the topic, see the web-site of the "On-line Encyclopedia of Integer Sequences", `http://www.research.att.com/~njas/sequences/` and references therein.

References

1. B. Bollobás, *Random Graphs*, Academic Press: London, 1985.
2. A.R. Conway and A.J. Guttmann, "On two-dimensional percolation," *J. Phys. A: Math. Gen.*, Vol. 28, pp. 891–904, 1995.
3. D. Coupier, A. Desolneux, and B. Ycart, A zero-one law for first-order logic on random images, in *Proceedings Math-Info'04*, Drmota et al. (Ed.), to appear 2004.
4. A. Desolneux, L. Moisan, and J.-M. Morel, "Meaningful alignments," *Int. J. Computer Vision*, Vol. 40, No. 1, pp. 7–23, 2000.
5. A. Desolneux, L. Moisan, and J.-M. Morel, "Edge detection by Helmholtz principle," *Journal of Mathematical Imaging and Vision*, Vol. 14, No. 3, pp. 271–284, 2001.
6. D.L. Donoho and I.M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *Journal of Amer. Stat. Assoc.*, Vol. 90, pp. 1200–1224, 1995.
7. H.J.A.M. Heijmans, "Self-dual morphological operators and filters," *Journal of Mathematical Imaging and Vision*, Vol. 6, No. 1, pp. 15–36, 1996.
8. H.J.A.M. Heijmans, "Connected morphological operators for binary images," *Computer Vision and Image Understanding*, Vol. 73, No. 1, pp. 99–120, 1999.
9. I. Jensen and A.J. Guttmann, "Statistics of lattice animals (polyominoes) and polygons," *J. Phys. A: Math. Gen.*, Vol. 33, pp. L257–L263, 2000.
10. D.A. Klarner, "Cell growth problems," *Canad. J. Math.*, Vol. 19, pp. 851–863, 1967.
11. D.A. Klarner and R.L. Rivest, "A procedure for improving the upper bound for the number of *n*-ominoes," *Canad. J. Math.*, Vol. 25, pp. 585–602, 1973.
12. S. Masnou and J.-M. Morel, "Image restoration involving connectedness," in *Proc. DIP'97, Vienna, Austria*, Vol. 3346. SPIE, 1998.
13. P. Monasse, "Contrast invariant representation of digital images and application to registration," PhD thesis, University Paris 9, 2000.
14. P. Monasse and F. Guichard, "Fast computation of a contrast invariant image representation," *IEEE Transactions on Image Processing*, Vol. 9, No. 5, pp. 860–872, 2000.
15. B. Sapoval and M. Rosso, "Gradient percolation and fractal frontiers in image processing," *Fractals*, Vol. 3, pp. 23–31, 1995.
16. J. Serra, *Image Analysis and Mathematical Morphology, Part. I*, Academic Press: New York, 1982.
17. J. Serra, *Image Analysis and Mathematical Morphology, Part. II: Theoretical Advances*, Academic Press: London, 1988.
18. J. Serra and L. Vincent, "An overview of morphological filtering," *Circuits, Systems and Signal Processing*, Vol. 11, No. 1, pp. 47–108, 1992.
19. J. Spencer, "Nine lectures on random graphs," in *Ecole d'été de probabilité de Saint-Flour XXI*, P. Bernard (Ed.), Vol. 1541 of *L.N. in Mathematics*, Springer-Verlag, New York, 1991, pp. 293–343.
20. L. Vincent, "Morphological area opening and closing for grayscale images," in *Proc. NATO Shape in Picture Workshop, Driebergen,* Springer-Verlag: The Netherlands, 1992, pp. 197–208.
21. L. Vincent, "Grayscale area openings and closings: Their applications and efficient implementation," in *Proc. EURASIP Workshop on Mathematical Morphology and its Applications to Signal Processing, Barcelona, Spain*, 1993, pp. 22–27.

David Coupier is 25 years old, he is PhD student at the MAP5. He has studied mathematics at the University of Orsay—Paris XI. He is working on zero-one laws and Poisson approximations for random images. Web page: http://www.math-info.univ-paris5.fr/∼coupier/

Agnès Desolneux is 30 years old, she is CNRS researcher at the MAP5. She defended her PhD thesis in applied mathematics in 2000 under the direction of Jean-Michel Morel at the ENS Cachan. She is working on statistical methods in image analysis. Web page: http://www.math-info.univ-paris5.fr/∼desolneux/

Bernard Ycart is 45 years old, he is Professor of mathematics at the University Paris 5 and director of the MAP5 (FRE CNRS 2428). He is specialist of applied probabilities (Markov processes, stochastic algorithms). Web page: http://www.math-info.univ-paris5.fr/∼ycart/