

Random Walks, Constrained Multiple Hypothesis Testing and Image Enhancement*

Noura Azzabou^{1,2}, Nikos Paragios¹, and Frederic Guichard²

¹ MAS, Ecole Centrale de Paris,
Grande Voie des Vignes, Chatenay-Malabry, France
noura.azzabou@certis.enpc.fr, nikos.paragios@ecp.fr
<http://www.mas.ecp.fr>

² DxOLabs, 3, Rue Nationale, 92100 Boulogne, France
{nazzabou, fguichard}@dxo.com
<http://www.dxo.com>

Abstract. Image restoration is a keen problem of low level vision. In this paper, we propose a novel - assumption-free on the noise model - technique based on random walks for image enhancement. Our method explores multiple neighbors sets (or hypotheses) that can be used for pixel denoising, through a particle filtering approach. This technique associates weights for each hypotheses according to its relevance and its contribution in the denoising process. Towards accounting for the image structure, we introduce perturbations based on local statistical properties of the image. In other words, particle evolution are controlled by the image structure leading to a filtering window adapted to the image content. Promising experimental results demonstrate the potential of such an approach.

1 Introduction

In spite of the progress made in the field of image denoising, it is still an open issue. In fact, natural images contain various types of information such as texture, small details, noise, fine structure and homogeneous regions. Such conditions make image filtering a crucial and challenging task. Ideally, a denoising technique must preserve all image element except noise.

Prior art in image denoising consists of methods of various complexity. Local filter operators, image decomposition in orthogonal spaces, partial differential equations as well as complex mathematical models with certain assumptions on the noise model have been considered. The sigma filter method [15], the bilateral [22] filter, morphological operators [25] and the mean shift algorithm [5] are efficient local approaches to image denoising. The first two approaches compute a weighted average over the pixel neighborhood where weights reflect the spatial distance between pixels and also the difference between their intensities. Such methods account to a minimal extend for the image structure and introduce strong bias in process through the selection of the filter bandwidth.

Image decomposition in orthogonal spaces like wavelets [17], splines, fourier descriptors and harmonic maps is an alternative to local filtering. Images are represented through

* This work has been carried out during the affiliation of Ms. Azzabou and Prof. Paragios with CERTIS, Ecole Nationale des Ponts et Chaussees.

a class of invertible transformations based on an orthogonal basis. Filtering consists of modifying the coefficients of the transformation space where often the most important ones are eliminated. Reconstruction of the image using the new set of coefficients leads to natural denoising. In their origin such methods failed to preserve boundaries a limitation that has been addressed through more careful selection of the orthogonal basis driven from the image structure [6, 13]. Such techniques have good performance when dealing with edges but they fail to preserve small details and texture.

Partial differential equations[1], higher order nonlinear operators [2], and functional optimization [19, 21, 23] have been also considered to address image denoising. The anisotropic diffusion [20] was a first attempt to incorporate image structure in the denoising process. Despite numerous advantages, theoretical justification [3] and numerous provisions of such a method one can claim that it remains myopic and cannot deal with image textures. The Mumford-Shah framework [19], the total variation minimization [21], the Beltrami flow [12], and other cost functionals of higher order [2] make the assumption that the image consists of a noise-free smooth component and the oscillatory pattern which corresponds to the random noise. Within such a concept constraints at limited scale are also introduced and image is reconstructed through the lowest potential of a cost function, that is often recovered in an iterative fashion through the calculus of variations. In the most general case such cost functions are not convex and therefore the obtained solution could correspond to a local minimum. Such methods are also myopic and still fail to account for texture patterns despite recent advances [24].

In order to account for image structure [18] an effort to understand the behavior of natural images when seen through a set of orientation and scale selective band-pass operators was made [14, 16]. Central assumption on this effort was that images exhibit differentially Laplacian statistics [16]. Such information is critical to an image denoising approach since it suggests the optimal way to regularize the problem and design the most efficient algorithm. Despite promising results, such simplistic modeling often fails to capture dependencies in a larger scale as well as account for the presence of repetitive patterns like texture.

To conclude, traditional state-of-the art techniques are often based on restoring image values based on local smoothness constraints within fixed bandwidth windows where image structure is not considered. Consequently a common concern for such methods is how to choose the most appropriate bandwidth and the most suitable set of neighboring pixels to guide the reconstruction process. In this context, the present work proposes a denoising technique based on multiple hypotheses testing. To this end, the reconstruction process is guided from multiple random walks where we consider a number of possible neighboring sites in the image and through a multiple hypotheses testing, we track the most suitable ones. Furthermore, image structure at a variable local scale is considered through a learning stage that consists of recovering probabilistic densities capturing co-occurrences of visual appearances at scale spaces. Kernels of fixed bandwidth are used to approximate such individual complex models for the entire visual spectrum. Random perturbations according to these densities guide the "trajectories" of a discrete number of walkers, while a weighted integration of the intensity through the random walks leads to the image reconstruction. Such a method is presented in [Fig. (1)].

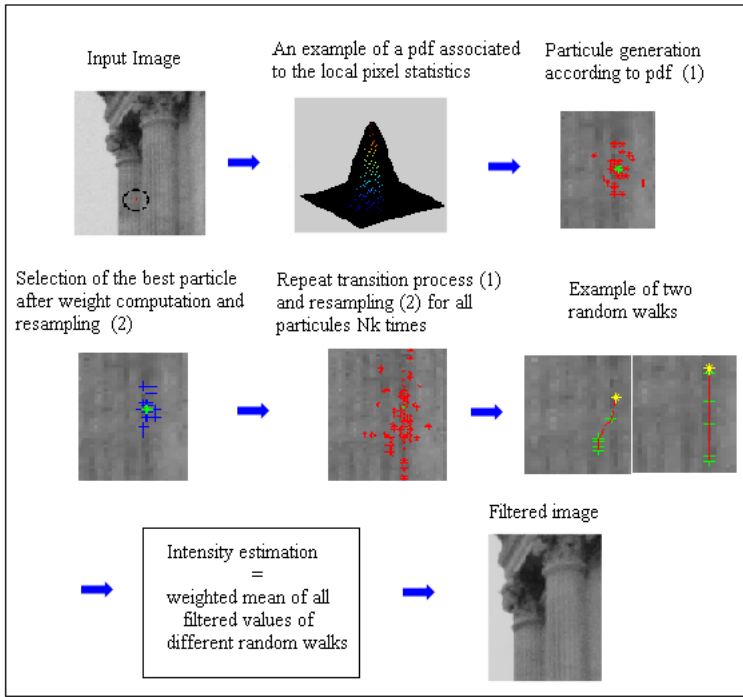


Fig. 1. Overview of Random Walks, Constrained Multiple hypotheses Testing and Image Enhancement

The reminder of this paper is organized in the following fashion; in section 2 we present density estimation of co-occurrence for encoding the structure present in the image. Random walks and particle filters are presented in section 3. While section 4 is devoted to the application of the particle filtering to denoising as well as some experimental results and comparisons with the state of the art methods . Finally, we conclude in section 5.

2 Statistics of Natural Images

Understanding visual content has been a constant effort in computer vision with applications to image segmentation, classification, retrieval and coding. Statistical modeling of images aims to recover contextual information at a primitive stage of visual processing chain. Co-occurrence matrices [11] have been a popular method to classification and segmentation of texture images.

Such a matrix is defined by a distance and an angle, and aim to capture spatial dependencies of intensities. The formal mathematical definition of an element (m, n) for a pair (d, θ) is the joint probability on the image that a m -valued pixel co-occurs with a n -valued pixel, with the two pixels are separated by a distance d and an angle θ :

$$C_{d,\theta}(m, n) = p_{(\mathbf{x},\mathbf{y}) \in \Omega}(m, n) (I(\mathbf{x}) = m, I(\mathbf{y}) = n, \mathbf{y} - \mathbf{x} = de^{i\theta})$$

with I being the observed image and Ω its domain. In the case of image denoising, the diagonal values of this matrix are of significant importance since implicitly they provide information on the geometric structure of the image. Inspired by such a concept, an intelligent denoising algorithm should be able to extract the most important correlations of local structure from the entire image domain, that is an ill-posed problem. Let us assume the absence of knowledge on the noise model. Then, in order to encode image structure we should seek for an estimate of the pdf.

$$p_f(d, \theta) = pdf(\{(d, \theta) \text{ where } \mathbf{x} \in \Omega, I(\mathbf{x}) = I(\mathbf{x} + de^{i\theta}) = f\})$$

To account for pixel values corrupted by noise, the constraint of exact matching could be relaxed, leading to:

$$p_{f,s}(d, \theta) = pdf(\{(d, \theta) \text{ where } \mathbf{x} \in \Omega, I(\mathbf{x}) = f, [\delta(\mathbf{x}, \mathbf{x} + de^{i\phi}) < \epsilon] \text{ and } [d < s]\})$$

where s is the scale considered for the pdf computation and $\delta(\cdot)$ is a metric that reflects similarity between to pixels in the image. This metric can be a simple distance such as the L^1 or the L^2 norm or more complex measures like correlation, histogram matching, mutual information, etc. In our experiments, we integrated the local variance into the pdf expression. In fact local variance (noted $\sigma(\cdot)$) is a simple primitive capable of describing texture at small scales. The new formulation of pdf is then as follows:

$$p_{f,\sigma,s}(d, \theta) = pdf(\{(d, \theta) \text{ where } \mathbf{x} \in \Omega, I(\mathbf{x}) = f, [\delta(\mathbf{x}, \mathbf{x} + de^{i\theta}) < \epsilon_1]; [\eta(\sigma(\mathbf{x}), \sigma(\mathbf{x} + de^{i\theta})) < \epsilon_2] \text{ and } [d < s]\})$$

As far as scale is concerned, different methods can be used to self-determine the scale like in the case of co-occurrence matrices. In the most general case we can assume scales of variable length that are self-adapted to the image structure. One can pre-estimate such pdf from the image using its empirical form.

However, $p_{f,s,\sigma}(d, \theta)$ aims to capture information of different structure, it describes spatial relation between similar patches in the image that may correspond to different population. To estimate $p_{f,s,\sigma}(d, \theta)$ non-parametric kernel-based density approximation strategies [26] like parzen windows were used.

Let $\{\mathbf{x}_i\}_{i=1}^M$ denote a random sample with probability density function p . The fixed bandwidth kernel density estimator consists of

$$\hat{p}(\mathbf{x}) = \frac{1}{M} \sum_{i=1}^M K_{\mathbf{H}}(\mathbf{x} - \mathbf{x}_i) = \frac{1}{M} \sum_{i=1}^M \frac{1}{\|\mathbf{H}\|^{1/2}} K(\mathbf{H}^{-1/2}(\mathbf{x} - \mathbf{x}_i))$$

where \mathbf{H} is a symmetric definite positive - often called a bandwidth matrix - that controls the width of the kernel around each sample point \mathbf{x}_i . Gaussian kernels are the most common selection of such an approach and that is what was considered in our case to approximate $p_{f,s,\sigma}(d, \theta)$. Once such pdf has been constructed from the image, we are able for a given image position \mathbf{x} and an observation ($f = I(\mathbf{x}), \sigma$) to generate a number of hypotheses for the most prominent position of the related image structure ([Fig. 2]).

One can now reformulate the problem of image denoising for a given pixel as a tracking problem in the image domain. Thus, given a starting position (pixel itself),




Fig. 2. Two pdf distribution $p_{f,\sigma}(d, \theta)$ for different values of f and σ (top ($f = 39, \sigma = 11.67$), bottom ($f = 96, \sigma = 3.55$)), and sample generation according to these pdf (red pixel) for two different positions

the objective is to consider a feature vector that upon its successful propagation along similar image structure, is able to remove and recover the original image value. To this end, we define

- a feature vector, that defines the current state of the reconstruction process s_t ,
- a measure of quality of a given hypothesis (feature vector) with respect to the image data.

with $[s_t = (\mathbf{x}_t, \hat{I}(\mathbf{x}))]$ being the state vector at a given time t . This state vector corresponds to the candidate site that can be used with in the filtering process and the reconstructed value induced by this site. The statistical interpretation of such an objective refers to the introduction of a probability density function (pdf) that uses previous states to predict possible new positions and image features to evaluate the new positions. The multiple hypotheses generation could be done in a number of fashions. Sequential Monte Carlo sampling is a well known technique that associates evolving densities to the different hypotheses, and maintains a number of them. Particle filters are popular techniques used to implement such a strategy.

3 Bayesian Tracking, Particle Filters and Multiple Hypotheses Testing

The Bayesian tracking problem can be simply formulated as the computation of the pdf relative to the present state s_t of a system, based on observations $z_{1:t}$ from time 1 to

time t : $p(s_t|z_{1:t})$. Assuming that one can have access to the prior pdf $p(s_{t-1}|z_{1:t-1})$, the posteriori pdf $p(s_t|z_{1:t})$ can be computed using the Bayes' rule:

$$p(s_t|z_{1:t}) = \frac{p(z_t|s_t)p(s_t|z_{1:t-1})}{p(z_t|z_{1:t-1})},$$

where the prior pdf is computed via the Chapman-Kolmogorov equation

$$p(s_t|z_{1:t-1}) = \int p(s_t|s_{1:t-1})p(s_{t-1}|z_{1:t-1})ds_{t-1},$$

and

$$p(z_t|z_{1:t-1}) = \int p(z_t|s_t)p(s_t|z_{1:t-1})ds_t$$

The recursive computation of the prior and the posterior pdf leads to the exact computation of the posterior density. Nevertheless, in practical cases, it is impossible to compute explicitly the posterior pdf $p(s_t|z_{1:t})$, and therefore an approximation method is to be introduced.

Particle filters, which are sequential Monte-Carlo techniques, estimate the Bayesian posterior probability density function (pdf) with a set of samples. Sequential Monte-Carlo methods have been first introduced in [9, 27]. For a more complete review of particle filters, one can refer to [10, 7].

Particle filtering methods approximate the posterior pdf by M random state sample $\{s_t^m, m = 1..M\}$ associated to M weights $\{w_t^m, m = 1..M\}$, such that

$$p(s_t|z_{1:t}) \approx \sum_{m=1}^M w_t^m \delta(s_t - s_t^m).$$

Thus, each weight w_t^m reflects the importance of the sample s_t^m in the pdf.

The samples s_t^m are drawn using the principle of Importance Density [8], of pdf $q(s_t|s_{1:t}^m, z_t)$, and it is shown that their weights w_t^m are updated according to

$$w_t^m \propto w_{t-1}^m \frac{p(z_t|s_t^m)p(s_t^m|s_{t-1}^m)}{q(s_t^m|s_{t-1}^m, z_t)}. \tag{1}$$

This equation shows that particle weights are updated using two mainly informations : the observation pdf which reflects the likelihood of seeing an observation z_t knowing the state s_t and the transition model which control the evolution of a particle state. The *sampling importance resampling* algorithm (SIR) consists in choosing the prior density $p(s_t|s_{t-1})$ as importance density $q(s_t|s_{1:t}^m, z_t)$. Doing so, equation (1) becomes simply

$$w_t^m \propto w_{t-1}^m p(z_t|s_t^m), \tag{2}$$

To sum up particle filtering consists of three main steps:

- particle drawing according the transition law $p(s_t^m|s_{t-1}^m)$
- computation of the likelihood of observations generated by the particle $p(z_t^m|s_t^m)$
- weight updating according to $w_t^m \propto w_{t-1}^m p(z_t|s_t^m)$




Fig. 3. Two different particle’s random walks(in red and yellow) starting from the same origin pixel (in magenta) where origin pixel is on the border (left image) or in an homogeneous region (right image)

After several steps a degeneracy issue occurs, such that all weights but few become null. In order to keep as many samples as possible with respectful weights, a resampling is necessary. Different resampling processes exist. The SIR algorithm consists in selecting the most probable samples in a random way, potentially selecting several times the same sample. An example of propagation of multiple hypotheses is shown in [Fig. (3)] for two different situations according to the position of the origin pixel.

4 Random Walks and Image Denoising

We now consider the application of such a schema to image denoising. Thus, given an origin pixel (\mathbf{x}) reconstruction is equivalent to recover a number of "random" positions ($\mathbf{x} = \mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_\tau$) with similar properties to \mathbf{x} to reconstruct the corrupted origin value ($I(\mathbf{x})$). The set of the obtained trajectories of each particle and their corresponding weights will represent the "filtering window". To this end, the use of "constrained" multiple hypotheses will be considered. This approach requires the definition of a perturbation model as well as a likelihood measure that reflects the contribution of a trajectory to the denoising process.

4.1 Likelihood Measure

Measuring similarities between image patches has been a well studied problem in computer vision. Within the proposed approach, filtering is done in a progressive fashion and therefore a need exists to measure the contribution of a new element in the filtering process. Parallel to that, each particle corresponds to a random walk where a certain number of pixels have been selected and contribute to the denoising process. Therefore, we define two metrics, one that accounts for the quality of potential additions and one for the intra-variability of the trajectories.

- The L^2 error-norm between local neighborhoods centered at the current position \mathbf{x}_t and at the origin pixel \mathbf{x} .

$$D_{sim}(t) = \frac{1}{(2W + 1)^2} \sum_{\mathbf{v} \in [-W, W] \times [-W, W]} |I(\mathbf{x} + \mathbf{v}) - I((\mathbf{x}_t + \mathbf{v}))|^2$$

where W is the bandwidth which must be carefully selected to get a reliable measure of similarity while being computationally efficient.

- In order to account for the intra-variability of the trajectories, we consider the variance, centered at the origin value,

$$D_{intra}(t) = \frac{1}{t} \sum_{\tau=0}^t (I(\mathbf{x}_\tau) - I(\mathbf{x}))^2$$

that measures the "uniformity" of the trajectory and could also be determined within a larger neighborhood (not at the pixel level). This terms insures edges and fine structure enhancement since random walks with small intra-variability are favored.

These two metrics are considered within an exponential function to determine the importance of a new sample given the prior state of the walk.

$$w_t = e^{-\left(\frac{D_{sim}(t)}{2\sigma_g^2} + \frac{D_{intra}(t)}{2\sigma_v^2}\right)} \tag{3}$$

The next step consists of defining an appropriate strategy for samples perturbation. This step is based on the statistical model for image structure introduced in the second section. The distribution $p_{f,\sigma}$ determines the transition model between particle at position \mathbf{x}_t and \mathbf{x}_{t+1} . In fact similar patches have similar values of local mean and variance and our transition model guide particle to those patches, since displacements that guarantee this similarity are favored.

4.2 Implementation and Validation

In this section we will be concerned about the application of the particle filtering process to denoising. To this end, for each pixel \mathbf{x} of the image, we generate N number of particles by applying N perturbations to the initial position \mathbf{x} . Then, each particle is propagated using a perturbation driven from the conditional distribution of the image statistics described in section (2). The process is repeated for (T) iterations. In each step of the process, we associate to each random walk a weight according to the likelihood measure defined in expression (3). We define then the walk value $\hat{I}_t^m(\mathbf{x})$ as the average value along the walk. It corresponds to the value used to reconstruct the original pixel according the "random walk" m :

$$\hat{I}_t^m(\mathbf{x}) = \frac{\sum_{\tau=0}^t D_{sim}(\tau) I(\mathbf{x}_\tau^m)}{\sum_{\tau=0}^t D_{sim}(\tau)}$$

Linear combination of the hypotheses weights and the corresponding denoised values is used to produce the current state of the process:

$$\hat{I}_t(\mathbf{x}) = \sum_{m=0}^N w_t^m \hat{I}_t^m(\mathbf{x})$$

In order to avoid degeneration of samples, as well as use with maximum efficiency all hypotheses, a frequent resampling process is used. In practice we use (N=30) particles,




Fig. 4. Original image (a) and method noise of Total Variation (b), Anisotropic Diffusion (c), Bilateral Filter (d), Random Walks (e), Non Local Mean (f)

Table 1. PSNR values for denoised image (The PSNR of the image corrupted by gaussian noise of $\text{std}=20$ is equal to 22.15)

	TV	AD	Bilateral	NLmean	R.Walks
Barbara	26.18	26.45	26.75	29.46	28.77
Boat	27.72	28.06	27.82	28.70	28.53
FingerPrint	26.08	24.81	24.12	26.22	26.00
House	28.43	29.41	29.18	31.47	30.87
Lena	28.45	29.27	29.28	31.18	30.32

with ($T=4$) pixels contributing to each walk. To illustrate the random walks filtering an overview of the hole process is presented in [Fig. (1)].

Towards objective validation of the method, we used natural images corrupted by a synthetic gaussian noise ($\sigma_n=20$). We compared our approach to well known filtering techniques such as the bilateral filter [22], the Non Local Mean [4] approach, the total variation [21] and the anisotropic filtering [20] using an edge stopping function of the type $(1 + |\nabla I|^2 / K^2)^{-1}$. The parameters of different methods were tuned to get a good balance between texture preserving and noise suppression. As for qualitative criteria, we used the "method noise" criterion, which corresponds to the difference between the noisy image and the filtered one [Fig. (4)]. Ideally, the "method noise" should be free of all image information and must look like random noise. In case of the total variation, the anisotropic diffusion and the bilateral filter [Figs. (4.b), (4.c) and (4.d)]




Fig. 5. Results of random walks filtering on natural images corrupted by additive gaussian noise; (i) Input image, (ii) Denoised Image

the method noise contains much more image details than our technique and the Non local mean. They fail to preserve the image texture because they are not based on patch comparison. Our Method has similar performance to the Non local mean filtering in terms of detail preservation. This is not surprising since particle weight is dependent on patch similarity. Thus, in homogeneous zones particles evolve in an isotropic fashion which is equivalent with considering a window with a fixed bandwidth. In textured regions, particles transitions rely on image structure where trajectories describing these structure are favored.

As far as quantitative validation is concerned we used the Peak Signal to Noise Ration criterion defined by

$$PSNR = 10 \log_{10} \frac{255^2}{MSE} \quad MSE = \frac{1}{\|\Omega\|} \sum_{\mathbf{x} \in \Omega} (I_0(\mathbf{x}) - \hat{I}(\mathbf{x}))^2$$

where I_0 is the noise free ideal image and \hat{I} its estimation by the denoising process.

In table 1, we report experimental validation results for our method on image used ofently to evaluate performance of image denoising algorithms. Parameters of each method were selected to reach its best PSNR value. For our technique the following set of parameter ($N=30, T=4, \sigma_v=\sigma_g=15, W=4$) provides the best results. We can see that quantitative validation is in harmony with the qualitative one. The NLmean achieves the best result since it scans large width window (15×15 in our tests) to find similar patches. Our approach aims to introduce a technique where similar patch selection is

guided by image structure. However, in the case of strong noise structures are not easily tractable which limits the performance of the method when compared to the Non local mean.

5 Conclusions

In this paper we have proposed a novel technique for image filtering. The core contribution of our method is the selection of an appropriate walk in the image domain towards optimal denoising. Such concept was implemented in a non-exclusive fashion where multiple hypotheses are maintained. The use of monte-carlo sampling and particle filters were considered to inherit such a property in the method. Furthermore, inspired by co-occurrence matrices we have modeled global image structure towards optimizing the selection of trajectories of the multiple hypotheses concept. To further adapt the method to the image structure such modeling was updated on line using local structure. Promising experimental results demonstrate the potentials of our approach.

Computational complexity is a major limitation of the method. The use of smaller number of hypotheses could substantially decrease the execution time. Improving the learning stage and guiding the particles to the most appropriate directions is a short term research objective. To this end, we would like to provide techniques capable of selecting the scales of each operator. Furthermore, we would like to consider kernels of variable bandwidth when recovering the non-parametric form of the learned distribution that are more efficient to capture image structure. More long term research objectives refer a better propagation of information within trajectories. Particle filters is a fairly simple approach that mostly propagates the mean value and the weights. The propagation of distributions can better capture the importance of the trajectories as well as the effect of new additions. In addition to that, geometric constraints on the "walks" could also improve the performance of the method in particular when texture is not present.

References

1. L. Alvarez, F. Guichard, P.-L. Lions, and J.-M. Morel. Axioms and fundamental equations of image processing. *Archive for Rational Mechanics*, 123:199–257, 1993.
2. G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*. Springer-Verlag, 2001.
3. M. Black, G. Sapiro, D. Marimont, and D. Heeger. Robust anisotropic diffusion. *IEEE Transactions on Image Processing*, 7:421–432, 1998.
4. A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *CVPR*, pages 60–65, 2005.
5. Y. Cheng. Mean Shift, Mode Seeking, and Clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17:790–799, 1995.
6. M. Do and M. Vetterli. Pyramidal directional filter banks and curvelets. In *ICIP01*, pages III: 158–161, 2001.
7. A. Doucet, J. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.
8. Arnaud Doucet. On sequential simulation-based methods for bayesian filtering. Technical Report CUED/F-INFENG/TR. 310, Cambridge University Department of Engineering, 1998.

9. N. Gordon. Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation. *IEE Proceedings*, 140:107–113, 1993.
10. N. Gordon. A Tutorial on Particle Filters for On-line Non-linear/Non-Gaussian Bayesian Tracking. *IEEE Transactions on Signal Processing*, 50:174–188, 2002.
11. R.M. Haralick, K. Shanmugam, and I. Dinstein. Textural features for image classification. *SMC*, pages 610–621, 1973.
12. R. Kimmel, R. Malladi, and N. Sochen. Image processing via the beltrami operator. In *ACCV*, pages 574–581, 1998.
13. E. Le Pennec and S. Mallat. Sparse geometric image representations with bandelets. *IEEE Transactions on Image Processing*, pages 423–438, 2005.
14. A. Lee, K. Pedersen, and D. Mumford. The nonlinear statistics of high-contrast patches in natural images. *International Journal of Computer Vision*, pages 83 – 103, 2003.
15. S. Lee. Digital image smoothing and the sigma filter. *CVGIP*, 24(2):255–269, November 1983.
16. A. Levin, A. Zomet, and Y. Weiss. Learning to perceive transparency from the statistics of natural scenes. In *NIPS*, 2002.
17. S. Mallat. A theory for multiscale signal decomposition: The wavelet representation. *IEEE Transactions on Pattern and Machine Intelligence*, pages 674–693, 1989.
18. P. Monasse and Guichard. Fast computation of a contrast invariant image representation. *IEEE Transactions on Image Processing*, 9:860–872, 2000.
19. D. Mumford and J. Shah. Optimal Approximation by Piecewise Smooth Functions and Associated Variational Problems. *Communications on Pure and Applied Mathematics*, 42: 577–685, 1989.
20. P. Perona and J. Malik. Scale space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern and Machine Intelligence*, 12:629–639, 1990.
21. L. Rudin, S. Osher, and E. Fatemi. Nonlinear Total Variation Based Noise Removal. *Physica D*, 60:259–268, 1992.
22. Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *ICCV*, pages 839–846, 1998.
23. D. Tschumperle and R. Deriche. Vector-valued image regularization with pde's : A common framework for different applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 506–517, 2005.
24. L. Vese and S. Osher. Modeling Textures with Total Variation Minimization and Oscillating Patterns in Image Processing. *Journal of Scientific Computing*, pages 553–572.
25. L. Vincent. Morphological grayscale reconstruction in image analysis: Applications and efficient algorithms. *IEEE Transactions on Image Processing*, 2:176–201, 1993.
26. M. Wand and M. Jones. *Kernel Smoothing*. Chapman & Hall., 1995.
27. W. West. Modelling with mixtures. In J. Bernardo, J. Berger, A. Dawid, and A. Smith, editors, *Bayesian Statistics*. Clarendon Press, 1993.