Stochastic Image Denoising

Francisco Estrada[1]
http://www.cs.utoronto.ca/~strider
David Fleet[1,2]
fleet@cs.utoronto.ca
Allan Jepson[2]
jepson@cs.utoronto.ca

[1] University of Toronto at Scarborough
1265 Military Trail, M1C 1A4
Toronto, On. CANADA
[2] University of Toronto
6 King's College Road, M5S 3G4
Toronto, On. CANADA

We present a novel algorithm for image denoising. Our algorithm is based on random walks over arbitrary neighbourhoods surrounding a given pixel. The size and shape of each neighbourhood are determined by the configuration and similarity of nearby pixels.

Assuming that pixels within the neighbourhood of x_0 are likely to have been generated by the same random process, we want the weights used to mix these pixels during denoising to depend on the similarity between them and x_0. At the same time, we require the random walk to follow a smooth path from x_0 to any other pixel in the neighbourhood, so the transition probabilities should also depend on the similarity between pairs of neighbouring pixels along any given path.

With this in mind, we define a random walk originating at pixel x_0 as an ordered sequence of pixels $T_{0,k} = \{x_0, x_1, \dots, x_k\}$ visited along the path from x_0 to x_k. Within this sequence, the probability of a transition between two consecutive pixels x_j and x_{j+1} is defined to be

$$p(x_{j+1}|x_j) = \frac{1}{K} e^{\left(\frac{-d(x_0,x_{j+1})^2}{2\sigma^2}\right)} e^{\left(\frac{-d(x_j,x_{j+1})^2}{2\sigma^2}\right)}, \tag{1}$$

where K is a normalization constant, $d(x_i, x_j)$ is a dissimilarity measure relating two image pixels, and σ is a scaling parameter. A sequence $T_{0,k+1}$ is generated from $T_{0,k}$ by sampling the neighbourhood of x_k and choosing a neighbour with probability $p(x_{k+1}|x_k)$. Under the first-order Markov assumption, the probability of a sequence starting at x_0 is given by

$$p(T_{1,k}|x_0) = \Pi_{j=1}^{k} p(x_j|x_{j-1}). \tag{2}$$

The first term in Eq. 1 accounts for the similarity between x_0 and neighbouring pixels, it prevents blurring across soft brightness or colour gradients. The second term imposes a preference for smooth transitions and prevents crossing over strong edges. Pixel transitions take place within standard 8-neighbourhoods. The distance function $d(x_i, x_j)$ is simply the Euclidean distance between the RGB values for pixels i and j, but we note that the same algorithm can be applied on any data for which a suitable neighbourhood structure and distance function can be determined.

Given m random walks each starting at x_0, and each with a length of k steps, the final denoised estimate $\vec{I}(x_0^*)$ is given by

$$\vec{I}(x_0^*) = \frac{1}{C} \sum_{i=1}^{m} \sum_{j=1}^{k} W_j^i \vec{I}(x_j), \tag{3}$$

where $\vec{I}(x_j)$ is the colour of a pixel x_j visited during the jth step of the random walk, W_j^i is the weight given to this pixel in the final estimate, proportional to the average log-transition probability along the path from x_0 to x_j, and C is a normalization constant.

We provide a thorough evaluation of denoising performance on images from the Berkeley Segmentation Database (BSD) and show that our method outperforms competing algorithms including the bilateral filter (BF) [3], the block matching (BM) algorithm [2], and the total variation (TV) method [1] in terms of Structured Image Similarity Index [4]. The table below shows the median SSIM index for the four algorithms as well as for the input noisy images for different amounts of noise.

	SD	BM	BF	TV	Unprocessed
Noise $\sigma = 5$	**.9728**	.8844	.8952	.8557	.9568
Noise $\sigma = 10$	**.9378**	.8819	.8961	.8270	.8306
Noise $\sigma = 15$	**.9004**	.8797	.8768	.7806	.7542

Our algorithm consistently outperforms competing methods, it yields high noise reduction with excellent preservation of detail as shown in Fig. 1. For noise that is correlated over small pixel neighbourhoods, a simple extension of our algorithm to perform multi-pass denoising yields significantly improved results as illustrated in Fig. 2.

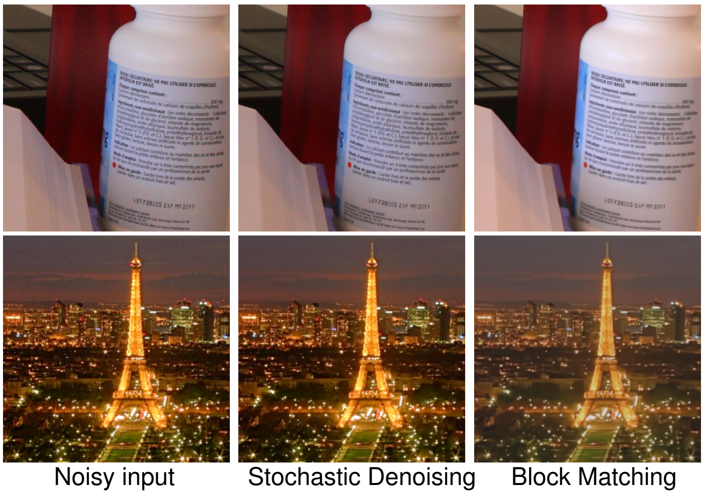

Figure 1: Denoising results on an input image with non-repetitive texture. The stochastic denoising result preserves very fine detail in the lettering on the bottle while removing noise. Conversely, the block matching algorithm over-smooths the textured region.

Figure 2: Single and multi-pass stochastic denoising results on a noisy image taken with a cell phone camera. The multi-pass denoising algorithm yields improved noise removal in this case, removing noise correlated over small pixel neighbourhoods while preserving sharp boundaries

[1] A. Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, 20:89–97, 2004.

[2] K. Davob, R. Foi, V. Katkovnik, and K. Egiazarian. Image denoising with block matching and 3d filtering. In *SPIE Electronic Imaging*, number 6064A-30, 2006.

[3] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *ICCV*, 1998.

[4] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004.