

Volterra Filtering of Noisy Images of Curves

Jonas August*

Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Abstract. How should one filter very noisy images of curves? While blurring with a Gaussian reduces noise, it also reduces contour contrast. Both non-homogeneous and anisotropic diffusion smooth images while preserving contours, but these methods assume a single local orientation and therefore they can merge or distort nearby or crossing contours. To avoid these difficulties, we view curve enhancement as a statistical estimation problem in the three-dimensional (x, y, θ) -space of positions and directions, where our prior is a probabilistic model of an ideal edge/line map known as the curve indicator random field (CIRF). Technically, this random field is a superposition of local times of Markov processes that model the individual curves; intuitively, it is an idealized artist's sketch, where the value of the field is the amount of ink deposited by the artist's pen. After reviewing the CIRF framework and our earlier formulas for the CIRF cumulants, we compute the minimum mean squared error (MMSE) estimate of the CIRF embedded in large amounts of Gaussian white noise. The derivation involves a perturbation expansion in an infinite noise limit, and results in linear, quadratic, and cubic (Volterra) CIRF filters for enhancing images of contours. The self-avoidingness of smooth curves in (x, y, θ) simplified our analysis and the resulting algorithms, which run in $O(n \log n)$ time, where n is the size of the input. This suggests that the Gestalt principle of good continuation may not only express the likely smoothness of contours, but it may have a computational basis as well.

1 Introduction

Is noise a help or a hindrance? Despite intriguing phenomena such as stochastic resonance, noise is largely viewed a necessary evil for applied science. Filtering images of contours, the problem studied here, is especially susceptible to noise because spatial differentiation is often involved in obtaining initial measurements (Fig. 1, top). Large noise levels, however, can actually simplify the filtering problem analytically, making difficult problems (asymptotically) manageable [22]. For example, here we derive nonlinear contour enhancement filters that operate in a neighborhood of an infinite noise limit.

This work is part of a larger project to understand visual contour computations as the Bayesian inference of a random field [2]. Beginning with a random

* I am grateful to Steven Zucker for his deep insights and his caring and patient advice in supporting this work. I thank Patrick Huggins, Athinodoros Georgiades, and Ohad Ben-Shahar for helpful discussions. Funding was provided by AFOSR and the Heinz Endowment. Please direct correspondence to jonas@cs.cmu.edu.

process model of individual curves, we have developed a prior for *images of* these curves, which we call the *curve indicator random field* (CIRF), and we have obtained the complete CIRF statistics. The main contributions of this paper are three new linear and nonlinear Volterra filters for computing the posterior mean of the CIRF for noisy contour images.

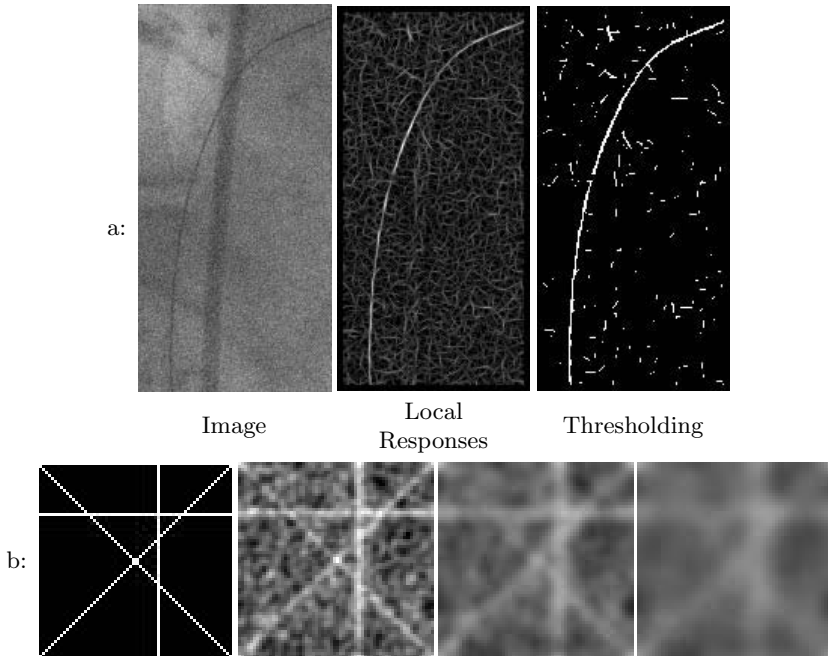


Fig. 1. (a) (Left) A guide wire imaged fluoroscopically during surgery. Finding the wire in noise is important [11] because surgery can take hours, all the while exposing the patient to radiation. Reducing exposures increases the noise level, however, and makes *local* filtering perform poorly (e.g., logical/linear positive-contrast line operator response tuned to thin curve, center, and its thresholding, right). In addition, local operators often employ numerical differentiation, which increase noise. (b) Crossing contours (left) in noise (center left) are distorted by Gaussian blurring (center right and right with blur sizes 5 and 10, respectively), revealing loss of distinction among overlapping contours, as well as reduced contrast. In this paper we exploit a model of random *images of* contours to derive nonlinear filters for overcoming noise without merging or attenuating contours.

Given the apparently slight distinction between a set of random curves and an image of them (the CIRF), it is not surprising that the CIRF has been overlooked as an object of study in vision. But there are significant advantages to making the CIRF a focus for understanding visual contours. First, it provides an exact notion of an ideal edge or line map (Fig. 1, top), satisfying the two desiderata of being (a) nonzero-valued along the true contours, and (b) zero-valued else-

where. The CIRF therefore provides a basis for formalizing saliency [24,13] as the problem of estimating the CIRF. Second, unlike dictionary-based relaxation labeling [8] and line processes in Markov random fields [7,14], the CIRF does not require an explicit and arbitrary enumeration and weighting of local field configurations to specify its joint statistical structure. Instead, by formally defining the field as a function of the curves depicted by it, the statistical properties of the field become a derivable consequence of those of the contours. Third, being a field, the CIRF makes it possible to formalize what is even meant by observing contours under various forms of corruption. Without the CIRF, a concept of a noisy contour might mean a smooth curve made rough; with the CIRF, a noisy *image of* a contour is simply the pointwise addition of one field (the CIRF) with another (white noise), for example. Fourth, the filters we derive provide a different notion of blur scale than convolution by a Gaussian (Fig. 1, bottom), or linear scale-space [12]. In particular, smoothing will take place along the (putative) contours, and unlike nonhomogeneous [19] and anisotropic diffusion [26], it allows for multiple, nearby and crossing curves. Finally, the CIRF provides a *local* representation of contour intersections; in contrast, parameterized curves require a strictly global computation to determine intersections. As ink builds up where contours cross in an artist’s sketch, so too does the (local) value of the CIRF increase at intersections.

There are two major sources of inspiration for this project. The first is the set of relaxation labeling [21] approaches to contour computation, which explicitly emphasize the need to include contextual interaction in inference, in contrast to purely local techniques [4] that detect edges largely independently. These approaches formalize contour inference as a labeling problem using a notion of label compatibility derived from good continuation in orientation [28] or curvature as well [18,10]. The second source of inspiration is the work on curves of least energy [9] or elastica [16] and its probabilistic expression as stochastic completion fields [27]. The CIRF was introduced to unify these works in a probabilistic framework.

2 A Prior on Images of Curves

Our random field inference framework for contour enhancement begins with a Markov process model for each contour. Mumford [16], Williams and co-workers [27] imagined a particle at R_t (in the space of positions (x, y) and directions θ in $\mathbb{R}^2 \times \mathbb{S}$) whose direction is slightly perturbed at each time instant t before taking its next step forward. The particle’s probability density $p(x, y, \theta, t)$ diffuses according to $\frac{\partial p}{\partial t} = Qp$, where $Q := \frac{\sigma_\kappa^2}{2} \frac{\partial^2}{\partial \theta^2} - \cos \theta \frac{\partial}{\partial x} - \sin \theta \frac{\partial}{\partial y} - \lambda^{-1}$ is the generator of Markov process R_t (the *direction process*), σ_κ bounds the direction perturbations and λ is the mean of the (exponentially-distributed) contour length T . In our framework the Markov process R_t models all image contours, irrespective of how they are observed. Indeed, the particular (stationary) Markov process contour model is unspecified in our framework; more exotic processes, which include scale [25] or curvature κ [3] can be used as well. At this level of

generality, the Markov process R_t takes on values (*states* or *sites*) i in a finite state space \mathcal{I} , e.g., $i = (x, y, \theta)$ (where space is discretized). The Markov process is summarized by its *Green's operator* $G = (g_{ij})$, for $i, j \in \mathcal{I}$, where g_{ij} is the average amount of time the process spends in site j given that it started in site i (Fig. 3). A natural image will have an unknown number \mathcal{N} (assumed to be Poisson-distributed with mean $\bar{\mathcal{N}}$) of contours, $R_{t_1}^{(1)}, \dots, R_{t_{\mathcal{N}}}^{(\mathcal{N})}$, which are independent and identically distributed (i.i.d.) with lengths $T_1, \dots, T_{\mathcal{N}}$.

So far we have only described the individual contours, but we know of them only through a (spatially distributed) field of measurements $M = (M_i)$ from an image (e.g., orientation-selective edge operator responses), which are corrupted due to noise and blur. To cope with such ambiguity, we suggest that contour organization be formalized as the estimation of an ideal field U of assertions of local contour existence.

Definition 1. *The curve indicator random field (CIRF) $U = (U_i)$ is:*

$$U_i := \sum_{n=1}^{\mathcal{N}} \int_0^{T_n} \mathbb{1}_{\{R_{t_n}^{(n)} = i\}} dt_n, \quad \forall i \in \mathcal{I}.$$

Here, the integral over t_n is called the *local time* at i for process $R_{t_n}^{(n)}$; the different processes (curves) are combined via superposition. In words, U_i is the (random) amount of time that particles spent in state i , and so non-zero values of U_i indicate that a contour passed through i (Fig. 2). We suggest that the CIRF be used a prior on images of curves not only because it captures an intuitive notion of the “perfect” edge map, but also because its can be tractably and fully characterized. In particular, the following is a formula for the CIRF cumulants¹—a complete set of *joint* statistics of this prior—and provides, in effect, an exact description of contextual interactions in an image of curves.

Proposition 1. *Suppose the Markov process R_t is spatially homogeneous and the endpoint distributions are spatially uniform. Let the contour density be $\eta := \bar{\mathcal{N}}\lambda|\mathcal{I}|^{-1}$. The k -th (joint) cumulant of the CIRF at sites i_1, \dots, i_k is $\text{cum}_{i_1, \dots, i_k} := \text{cum}\{U_{i_1}, \dots, U_{i_k}\} = \eta \sum g_{j_1 j_2} \cdots g_{j_{k-1} j_k}$, where the sum is over all permutations j_1, \dots, j_k of i_1, \dots, i_k .*

The proof requires the Feynman-Kac formula, a broadly-applied result in quantum physics (see [2] for details). Since CIRF cumulants of order greater than two are (generally) nonzero, the CIRF is therefore a non-Gaussian random field [17]. The cumulants are the coefficients of a Taylor series expansion of the *cumulant generating functional* $\text{CGF}_U(c) := \ln \text{MGF}_U(c)$ of U , where $\text{MGF}_U(c) := \mathbb{E} \exp \langle c, U \rangle$ is the *moment generating functional* of U at c and $\langle a, b \rangle$ denotes the *inner product* between vectors $a, b \in \mathbb{R}^{|\mathcal{I}|}$. We now turn the main question for this paper: how can the CIRF be used in a filter?

¹ Cumulants are a function of the moments of a distribution, but can be easier to work with. The lowest-order cumulants are the mean and covariance.




Fig. 2. Modeling images of curves with the CIRF prior. The contours in these images (left: angiogram (top) and surface of a Jupiter moon (bottom)) are primarily characterized by their number (\mathcal{N}), smoothness (σ_κ^{-1}), and extent (λ). Given only noisy local measurements $m = (m_i), i = (x, y, \theta)$, (center: logical/linear line operator responses [10]), we seek to enhance contours using a prior on images of curves. Several random samples of the CIRF prior in $\mathbb{R}^2 \times \mathbb{S}$ (right: various parameter settings) resemble (up to \mathcal{N} , σ_κ , and λ) the qualitative structure of the local responses, but without the corruption. These random samples of the prior are intended only to reflect the “contourness” of the local responses, not the particular contours, which are modeled by the posterior. At each (x, y) , the maximum over θ was taken for both local responses and the CIRF. The goal is to bias the enhancement method toward contour images (as opposed to pebble or sand texture images, for example).

3 Estimating Noisy CIRFs

3.1 The Likelihood

The likelihood represents how our ideal contour image U relates to the measurements M , characterizing the corruptions in local edge and line operator responses. We focus on noise contamination in the measurements because (1) edge/line detection requires differentiation, which amplifies noise and (2) the local signal-to-noise ratio (SNR) drops for fixed sensor noise, when the local contrast (i.e., signal) is low. For concreteness and simplicity, our measurement model is $M = U + N$, where N is white Gaussian noise N with zero-mean and variance σ_N^2 , and the likelihood is

$$p(m|u) = (\sqrt{2\pi}\sigma_N)^{-|I|} \exp\left(-\frac{1}{2}\sigma_N^{-2}\|m - u\|^2\right). \quad (1)$$

By a change in variables in m , other observation models with i.i.d. noise can be included, including those estimated from histograms of edge/line operator responses both on and off the actual contours [6,2].

3.2 The Posterior

Using Bayes' rule, the posterior distribution $P(du|m)$ is proportional to $p(m|u)$ times $P(du)$, where du is an infinitesimal region around realization u . The CIRF prior is written $P(du) := \mathbb{P}\{U \in du\}$ because a density for U exists only in a generalized sense². Expanding the norm in the likelihood (1), the posterior therefore becomes

$$P_\epsilon(du) := Z_\epsilon^{-1} \exp\left(\epsilon \left\langle m - \frac{1}{2}u, u \right\rangle\right) P(du) = P(du|m),$$

where the subscript $\epsilon := \sigma_N^{-2}$ (the *inverse noise variance*) indicates conditioning on the measurements m , and the *normalizing constant* is

$$Z_\epsilon := \int \exp\left(\epsilon \left\langle m - \frac{1}{2}u, u \right\rangle\right) P(du).$$

Although we do not have an explicit expression for the prior $P(du)$ (and so neither for the posterior), in §4 we shall find that we can indirectly use the prior through its cumulants in order to obtain Volterra filters.

3.3 Bayes Estimation

Filtering here means the estimation of the (unknown) CIRF U given a noisy realization m . In Bayes decision theory, estimation is formalized by specifying the posterior and a loss function $\text{loss}(U, u)$ that penalizes estimate u when the true unknown is U ; the Bayes estimate is then $\tilde{u} := \arg \min_u \mathbb{E}[\text{loss}(U, u)|m]$. Maximum a posteriori (MAP) estimation corresponds to the 0/1 loss function, which penalizes all errors equally; this loss is zero only when the estimate is perfect (at *all* sites i), and 1 otherwise. Despite the popularity of the MAP estimator, we believe an additive loss function such as the squared error $\|U - u\|^2$ is more appropriate for high-noise filtering because it penalizes small errors at a few sites much less than large errors at many sites, unlike the 0/1 loss function (MAP). Therefore we define our filtering task to be the computation of the minimum mean squared error (MMSE) estimate

$$\tilde{u} = \arg \min_u \mathbb{E}_\epsilon \|U - u\|^2, \tag{2}$$

where ϵ again denotes the conditioning on m . By a standard argument [5], the MMSE estimate \tilde{u} is equal to the *posterior mean* $\mathbb{E}_\epsilon U$ of the CIRF U .

² Specifically, at each site i there is positive mass for $u_i = 0$.

Since the posterior mean is generally difficult to compute, we seek simplifications. A typical approach would be to approximate the CIRF prior with a Gaussian, giving rise to MMSE linear filtering [2]. We avoid this method, however, because it ignores all the CIRF cumulants beyond the mean and covariance. Instead, we focus on an aspect of our problem that we have not yet considered: *extremely noisy conditions*. In this way we can simplify our filtering problem tremendously, while exploiting the higher-order statistical structure of the CIRF.

4 Volterra Filtering at Large Noise Levels

Volterra filters constitute one class of nonlinear filters that subsumes linear filters [23]. By definition, the output of a k -th order Volterra filter is a k -th degree polynomial³ function of the input signal. Analogous to the connection between MMSE linear filters and second-order statistics, MMSE Volterra filters are related to higher-order statistics. In particular, a MMSE k -th order Volterra filter generally requires up to the $(2k)$ -th order moments of the input and unknown [1, 20]. Unfortunately, the computational complexity of simply applying the general k -th order Volterra is $O(n^{k+1})$, where $n = |\mathcal{I}|$; solving for the MMSE filter coefficients is much more costly in general.

Here we derive linear, quadratic and cubic Volterra filters for approximating the posterior mean of the CIRF corrupted by large amounts of white Gaussian noise. Our approach is based on the following observation: in such noisy conditions, the inverse noise variance ϵ will be small. This suggests that we consider the infinite noise limit where ϵ approaches zero; in particular, we perform a Taylor series expansion of the posterior mean $\mathbb{E}_\epsilon U$ around $\epsilon = 0$. We then exploit simplifying aspects of the CIRF, especially the self-avoiding nature of the direction process, to *obtain Volterra filters that run in $O(n \log n)$ time*.

Now we proceed with the derivation. Observe that the posterior mean can be computed by differentiating the logarithm of the normalizing constant, or

$$\frac{\partial \ln Z_\epsilon}{\partial(\epsilon m_i)} = Z_\epsilon^{-1} \int \frac{\partial \exp(\epsilon \langle m - \frac{1}{2}u, u \rangle)}{\partial(\epsilon m_i)} P(du) = \int u_i P_\epsilon(du) = \mathbb{E}_\epsilon U_i. \quad (3)$$

Therefore to get a Taylor series for the posterior mean, we need only get a Taylor series for $\ln Z_\epsilon$ and then differentiate with respect to ϵm_i .

Where do we get a Taylor series for $\ln Z_\epsilon$? Observe that the normalizing constant Z_ϵ is just the moment generating functional of the random vector $W = \langle m - \frac{1}{2}U, U \rangle$ in disguise, or, equivalently, W has the cumulant generating functional $\text{CGF}_W(\epsilon)$ such that

$$\ln Z_\epsilon = \ln \int \exp(\epsilon w) P(du) = \ln \text{MGF}_W(\epsilon) = \text{CGF}_W(\epsilon).$$

³ Note the distinction between the (nonlinear) polynomial filters here and those linear filters whose *kernel* is based on some (e.g., Hermite) polynomial.

But the coefficients in a Taylor series expansion of a cumulant generating functional are by definition the cumulants, and therefore we need to obtain the cumulants of W , a (quadratic) polynomial in U . We now only sketch the rest of our derivation, as the full version takes 20 pages [2].

Since the cumulants of W are not directly known, we computed them in terms of McCullagh's so-called generalized cumulants (an algebraic structure that eases computations of polynomials of moments and cumulants) of U [15]. We then differentiated with respect to ϵm_i to get the first few terms in the Taylor expansion of $\mathbb{E}_\epsilon U_i$ as a function of the generalized cumulants of U .

Besides spatial homogeneity, two properties of the CIRF played a significant role in simplifying our calculations. First, the *contour density* η is *small* ($0 \leq \eta \ll 1$) in a large number of images, as contours constitute a small fraction of the total image (which is why pen ink does not run out quickly). In the continuum, an extreme form of this claim certainly holds: the bounding contours in an image of smooth surfaces in an occlusion relationship constitute a set of measure zero (with respect to the standard area measure in the plane). This low density is even more appropriate for the CIRF based on the direction process, where the curves live in a three-dimensional space. The low density of the CIRF was crucial for simplifying McCullagh's formula for expressing generalized cumulants in terms of a sum of products of (ordinary) cumulants: we neglect all terms which are a product of two or more CIRF cumulants (having magnitude $O(\eta^2)$ by Prop. 1). Thus the low density of the CIRF allowed us to express the posterior mean as a weighted $O(\eta)$ -sum of cumulants of U , each of the form $\sum_{i,j,k,l,\dots} \text{cum}_{r,i,j,k,\dots} m_i m_j \cdots \delta_{kl} \cdots$.

The second CIRF property that simplified our derivation is the *approximate self-avoidingness of the direction process*. If our Markov processes were perfectly straight, they would never return to their starting points, even if infinitely long. We observe in Fig. 3 that the direction process rarely returns to its starting position. Many permutations in cumulants of the CIRF (Prop. 1) could therefore be neglected.

To support the error-prone task of explicitly summing over the permutations in the CIRF cumulants in the $O(\eta)$ -sum for the posterior mean, we constructed a diagram corresponding to each permutation (Fig. 4). Nodes indicate sites, arrows indicate ordering in the permutation, and multiplicative constants such as η are temporarily suppressed. In Fig. 4a, permutation (r, i) of term $\sum_i \text{cum}_{r,i} m_i$ is $\sum_i g_{ri} m_i = (Gm)_r$, or " G acts on input m , and the result is evaluated at readout r " (G is understood for arrows away from r , G^* is understood for arrows toward r , and " X " represents an input m at a node). Using the diagram we see that the reverse permutation (i, r) gives rise to $G^* m$ (evaluated at r). In Fig. 4b, permutation (j, i, r) of term $\sum_{i,j} \text{cum}_{r,i,j} m_i m_j$ is $G^* \text{diag}(m) G^* m$, where $\text{diag}(m)$ is a diagonal matrix with m along the diagonal; the reverse permutation (r, i, j) is $G \text{diag}(m) G m$ using the reversed diagram. In Fig. 4c, permutation (r, i, j, k) of term $\sum_{i,j,k} \text{cum}_{r,i,j,k} m_i \delta_{jk}$ has a loop, which represents contour intersection (at node $j = k$) and arises from the Kronecker delta δ_{jk} , which in turn comes from the quadratic term in W . Terms that have larger loops can require taking

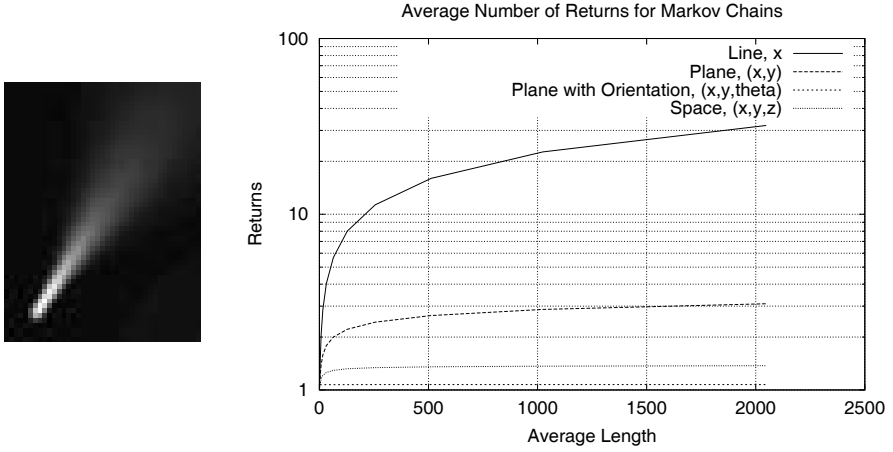


Fig. 3. (Left) The Green’s operator of the direction process (see §2); brightness indicates average time spent at a site, given that process started near lower left, traveling diagonally to the upper right (parameters: $\sigma_\kappa = 1/24$ with 44 directions; $\lambda = 100$ throughout this paper). (Right) The self-avoidingness of the direction process: a comparison of average number of returns to starting location for various Markov processes (chains), plotted as a function of total contour length. The average number of returns is at least one because the first visit is included. Unlike random walks on the line, in the plane, or even in space, the average number of returns of the direction process (“plane with orientation”) does *not* increase with contour length; the direction process is therefore (approximately) self-avoiding. Spatial stepsize (Δx , etc.) is one, except $\Delta\theta = 2\pi/(\text{number of directions})$; also, $\sigma_\kappa \approx \Delta\theta$ throughout this paper unless specified explicitly.

componentwise products of operators such as $G \odot G^*$, and are neglected due to the self-avoidingness of the direction process. See [2] for details.

These diagrams are loosely analogous to Feynman diagrams, which support perturbation calculations in statistical physics and quantum field theory. Independently of this work, Ruderman [22, pp. 29–33] used Feynman diagrams to calculate the posterior mean of an underlying signal observed in Gaussian white noise, where the perturbation is in terms of the non-Gaussian underlying signal. In contrast, the perturbation used here is due to the Gaussian noise; the underlying non-Gaussian signal (the CIRF) is modeled exactly.




Fig. 4. Diagrams dramatically simplified the cumulant permutation computations in the lengthy derivation of the Volterra CIRF filters. See text.

Using this diagram-based technique, we obtained linear, quadratic, and cubic Volterra filters which asymptotically approximate the posterior mean of the CIRF. Let $a \odot b$ denote the vector which is the componentwise product of vectors a and b , i.e., $(a \odot b)_i = a_i b_i$.

Result 1 (High-noise MMSE Volterra CIRF Filters) *Suppose that the curve indicator random field U (for approximately self-avoiding Markov processes) is corrupted with additive Gaussian white noise of variance $\sigma_N^2 =: \epsilon^{-1}$ to produce measurement vector m . Let $\zeta := g_{ii}\lambda$, where λ is the average curve length. Then the minimum mean squared error estimate of U given m has the following approximate asymptotic expansion as $\sigma_N \rightarrow \infty$ ($\epsilon \rightarrow 0$):*

Linear Filter:

$$\tilde{u}^{(1)} = \eta\{1 - 2\epsilon\zeta + \epsilon(Gm + G^*m)\} + O(\eta\epsilon^2 + \eta^2\epsilon) \quad (4)$$

Quadratic Filter:

$$\begin{aligned} \tilde{u}^{(2)} = \eta\{ & 1 - 2\epsilon\zeta + 3\epsilon^2\zeta^2 + \epsilon(1 - 2\epsilon\zeta)(Gm + G^*m) \\ & + \epsilon^2(G \operatorname{diag}(m) Gm + Gm \odot G^*m + G^* \operatorname{diag}(m) G^*m)\} \\ & + O(\eta\epsilon^3 + \eta^2\epsilon) \end{aligned} \quad (5)$$

Cubic Filter:

$$\begin{aligned} \tilde{u}^{(3)} = \eta\{ & 1 - 2\epsilon\zeta + 3\epsilon^2\zeta^2 - 4\epsilon^3\zeta^3 \\ & + \epsilon(1 - 2\epsilon\zeta + 3\epsilon^2\zeta^2)(Gm + G^*m) \\ & + \epsilon^2(1 - 2\epsilon\zeta)(G \operatorname{diag}(m) Gm + Gm \odot G^*m + G^* \operatorname{diag}(m) G^*m) \\ & + \epsilon^3(G \operatorname{diag}(m) G \operatorname{diag}(m) Gm + G \operatorname{diag}(m) Gm \odot G^*m \\ & + Gm \odot G^* \operatorname{diag}(m) G^*m + G^* \operatorname{diag}(m) G^* \operatorname{diag}(m) G^*m)\} \\ & + O(\eta\epsilon^4 + \eta^2\epsilon). \end{aligned} \quad (6)$$

These filters compute quickly: by implementing the operator G in the Fourier domain [2], its computational complexity is $O(n \log n)$ for the direction process, where $n = |\mathcal{I}|$. Since multiplying a diagonal matrix with a vector has $O(n)$ complexity, the net complexity for all of these filters is therefore $O(n \log n)$ for the direction process. This is far better than for the general k -th order Volterra filters, which have $O(n^{k+1})$ complexity.

5 Results

Here we apply our high-noise MMSE Volterra CIRF filters to some noisy synthetic and real images. To test robustness, we evaluated responses to image of a horizontal straight line under varying noise variance σ_N^2 . For the sythetic images in this paper, we mapped from the planar image to a discrete $\mathbb{R}^2 \times \mathbb{S}$ by copying the image to each direction, so that our measurements $m = m(x, y, \theta)$ were constant in θ . Noise was quantified using the peak signal-to-noise ratio (SNR), or

$10 \log_{10}(\mathbf{I}_{\max} - \mathbf{I}_{\min})^2 / \sigma_N^2$ (in decibels), for an image \mathbf{I} having minimum value \mathbf{I}_{\min} and maximum \mathbf{I}_{\max} . The 0 dB result is shown in Fig. 5 (left), showing significant noise cleaning (at each (x, y) , maximum response over θ is shown). The error between noisy and noisy-free responses (Fig. 5, right) shows improvement for increasing order of nonlinearity.

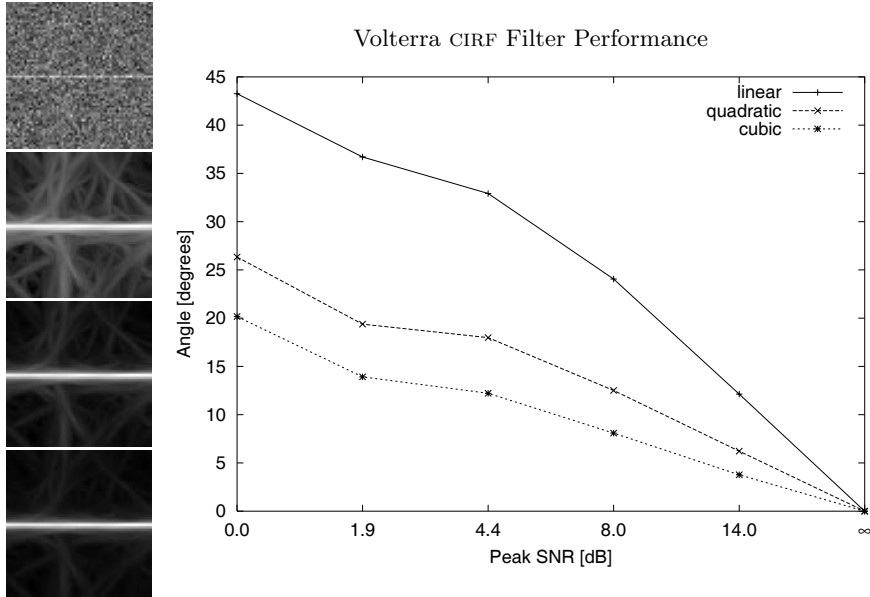


Fig. 5. Noisy image of a line (top left) is cleaned via linear, quadratic, and cubic CIRF filters (second from top to bottom, respectively; parameters: $\epsilon = 10, \zeta = 0, \sigma_\kappa = 1/10$, with 64 directions). (Right) Noise performance of the high-noise MMSE Volterra filters in $\mathbb{R}^2 \times \mathbb{S}$. Plotted is the angle between the mean-shifted noisy and noise-free filter responses for each filter type and at each noise level. Angle is measured between two vectors in $\mathbb{R}^{|\mathcal{I}|}$. Notice the advantage of the cubic filter over the others.

Next we tested the filter's response on both open and closed curves (Fig. 6). Observe that the response is well resolved at the contour crossings; even though we have no initial estimate of local direction. Fig. 7 is a more strenuous test at contour crossings. Unlike isotropic, nonhomogeneous, or even anisotropic diffusion, the Volterra filter performs smoothing along even overlapping contours, and therefore the preservation of junctions is possible. In addition, quite powerful orientation-selectivity is apparent (fixed θ slices of the response are shown in bottom rows). We studied the effect of parameter variations in Fig. 8. For most examples, we let $\zeta \geq 0$ be small (usually 0) and let $\epsilon > 0$ be about 10, although the exact settings did not appear to have a great effect.

To apply the Volterra filters to real images, we set the filter input m to be the (directional) output of a logical/linear operator (with parameter degree




Fig. 6. Filtering a noisy image of a circle with a diagonal line (top left, peak SNR=4.2 dB) with linear and cubic CIRF filters (responses at top center and right, resp.). The cubic response was thresholded and the resulting isosurface in $\mathbb{R}^2 \times \mathbb{S}$ is shown from different viewpoints (bottom, with $\theta \in [0, \pi)$). The filter is orientation selective and separates the contours in $\mathbb{R}^2 \times \mathbb{S}$. Parameters were $\epsilon = 10.5, \zeta = 0, \sigma_\kappa = 1/5$, with 32 directions.

$d = 0.1$ in Iverson’s implementation⁴, and other values at their default), instead of copying over direction. For the ship’s wake in Fig. 9 (top left), we zoom in the highlighted portion (top right) and then obtain logical/linear negative-contrast line responses and their thresholding (second row). Responses to high-noise Volterra filters ($\epsilon = 10.5, \zeta = 0, \sigma_\kappa = 1/15, \lambda = 100$, and 96 directions) is shown in the bottom three rows of Fig. 9. Observe how the many bogus responses in the thresholded local responses are not present in the quadratic and cubic CIRF filter responses. A prostate boundary is enhanced with the cubic CIRF filter in Fig. 10. Returning to our opening example of a noisy guide wire image (Fig. 1), we see that the CIRF Volterra filters reduce noise while enhancing the guide wire (Fig. 11).

6 Conclusion

In this paper we took a Bayesian estimation approach to contour enhancement at large noise levels, where we used an exact prior for images of curves, the CIRF.

⁴ Code and settings available at <http://www.ai.sri.com/~leei/loglin.html>.




Fig. 7. Crossing lines are teased apart with the cubic CIRF filters. From left to right in top row: original image, image corrupted with additive white Gaussian noise (peak SNR=8 dB), and cubic CIRF filter response (parameters were $\epsilon = 10.5, \zeta = 0, \sigma_\kappa = 1/10$, with 64 directions). Observe the noise reduction without merging the nearby junctions, unlike Gaussian blurring (Fig. 1). The direction-specific (i.e., without taking the max over direction) cubic responses (middle row, left to right: $\theta = 0^\circ, \dots, 135^\circ$) and the corresponding thresholded responses (bottom row) reveal that the initially strong responses at the inappropriate directions are chiseled away by the action of the filter.




Fig. 8. The effect of parameter variations. Since η amounts to a simple threshold, we only varied ϵ and ζ in cubic CIRF filter (6) applied to Fig. 6 (top left), as follows: $\zeta = 0$ (a,b,c), 1 (d,e,f), 50 (g,h,i); $\epsilon = 0.5$ (a,d,g), 10.5 (b,e,h), 50 (c,f,i). At large ζ , cubic responses (g,h,i) resembled linear response (Fig. 6, center); otherwise, the cubic CIRF filter was insensitive to parameter changes.

Exploiting both the high-noise limit and the approximate self-avoidingness of the direction process model of random, smooth curves, we obtained MMSE Volterra filters for contour enhancement that run in $O(n \log n)$ time, where n is the size of the input. These filters perform well even for crossing or nearby curves, thus this might be interpreted as an extension of anisotropic diffusion ideas where multiple local directions are allowed.

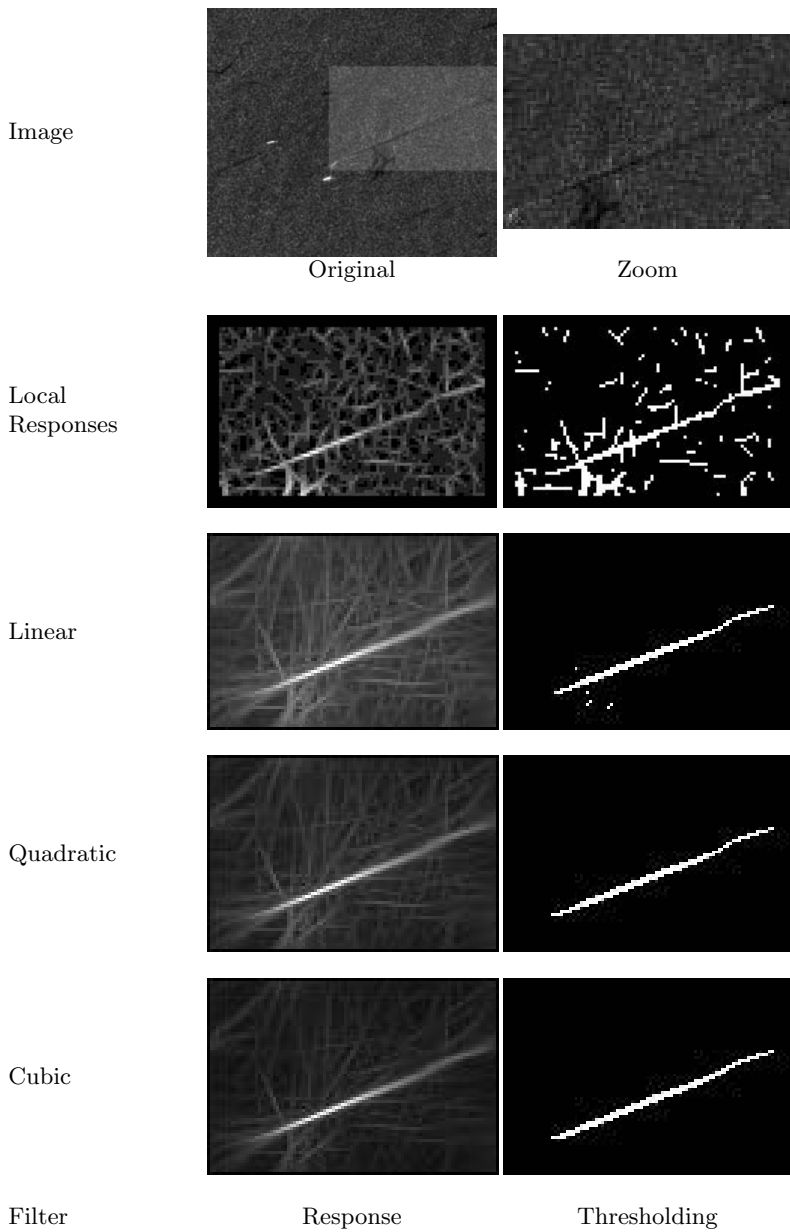


Fig. 9. Finding a ship's wake (see text for explanation).




Fig. 10. To find the boundary of a prostate imaged with ultrasound (left), we obtained blurred, directed local edge measurements (center, region around transducer and upper corners were cropped out), and then applied the cubic CIRF filter (right).

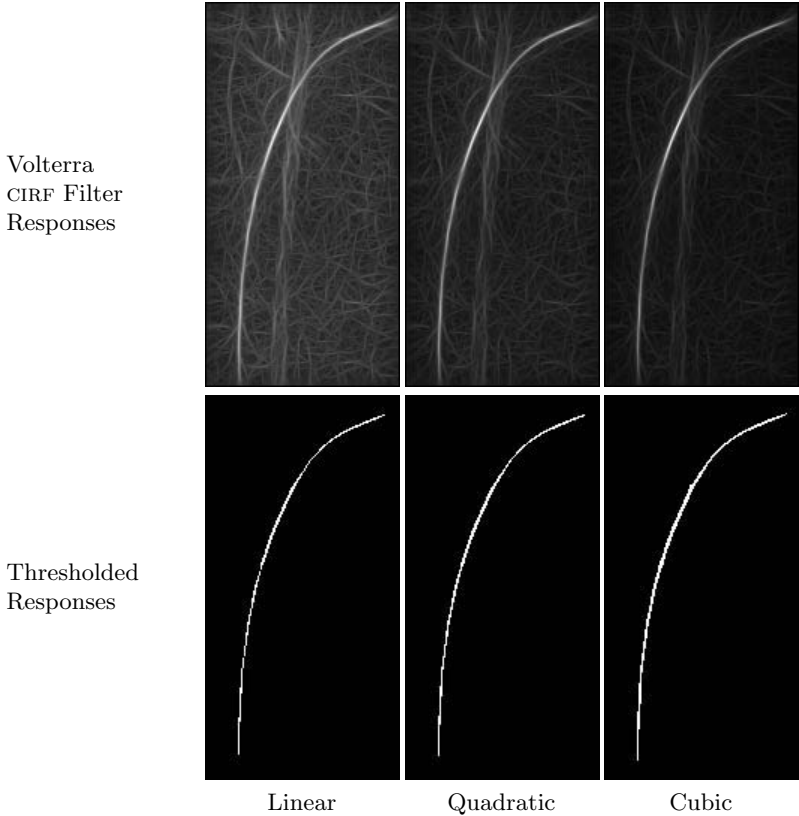


Fig. 11. Finding a guide wire in the image of Fig. 1. Responses to high-noise Volterra CIRF filters ($\epsilon = 10.5, \zeta = 0, \lambda = 100, \sigma_\kappa = 1/10$, with 64 directions). Observe how the cubic filter at $\sigma_\kappa = 1/10$ enhances the guide wire, by increasing, not reducing, contrast (cf. Fig. 1). To produce these two dimensional images from the actual discrete (x, y, θ) -space responses, the maximum over θ was taken.

References

1. P. Alper. A consideration of discrete volterra series. *IEEE Transactions on Automatic Control*, 10:322–327, 1965.
2. J. August. *The Curve Indicator Random Field*. PhD thesis, Yale University, 2001.
3. J. August and S. W. Zucker. A markov process using curvature for filtering curve images. In *Energy Minimization Methods for Computer Vision and Pattern Recognition*, 2001.
4. J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:679–698, 1986.
5. G. Casella and R. L. Berger. *Statistical Inference*. Duxbury, Belmont, CA, 1990.
6. D. Geman and B. Jedynak. An active testing model for tracking roads in satellite images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(1):1–14, 1996.
7. S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741, 1984.
8. E. R. Hancock and J. Kittler. Edge-labeling using dictionary-based relaxation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(2):165–181, 1990.
9. B. K. P. Horn. The Curve of Least Energy. *ACM Transactions on Mathematical Software*, 9:441–460, 1983.
10. L. A. Iverson. *Toward Discrete Geometric Models for Early Vision*. PhD thesis, McGill University, Montreal, 1994.
11. S. N. Kalitzin, B. M. ter Haar Romeny, and M. A. Viergever. Invertible orientation bundles on 2d scalar images. In *Proc. Scale-Space '97*, LICS, pages 77–88. Springer, 1997.
12. J. J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.
13. M. Lindenbaum and A. Berengolts. A probabilistic interpretation of the saliency network. In *European Conf. on Computer Vision*, 2000.
14. J. L. Marroquin. A markovian random field of piecewise straight lines. *Biological Cybernetics*, 61:457–465, 1989.
15. P. McCullagh. *Tensor Methods in Statistics*. Monographs in Statistics and Applied Probability. Chapman and Hall, Cambridge, England, 1987.
16. D. Mumford. *Algebraic Geometry and Its Applications*, chapter Elastica and Computer Vision, pages 491–506. Springer-Verlag, 1994.
17. C. L. Nikias and A. P. Petropulu. *Higher-Order Spectra Analysis: A Nonlinear Signal Processing Framework*. Prentice Hall, Englewood Cliffs, NJ, 1993.
18. P. Parent and S. W. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8):823–839, August 1989.
19. P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–39, jul 1990.
20. B. Picinbono and Y. Gu. Mean square estimation and projections. *Signal Processing*, 19:1–8, 1990.
21. A. Rosenfeld, R. A. Hummel, and S. W. Zucker. Scene labeling by relaxation operations. *IEEE Transactions on Systems, Man, and Cybernetics*, 6(6):420–433, 1976.

22. D. L. Ruderman. *Natural Ensembles and Sensory Signal Processing*. PhD thesis, U. C. Berkeley, 1993.
23. M. Schetzen. *The Volterra and Wiener Theories of Nonlinear Systems*. Wiley, New York, 1980.
24. A. Sha'ashua and S. Ullman. Structural saliency: The detection of globally salient structures using a locally connected network. In *Proceedings, Second International Conference on Computer Vision*, 1988.
25. K. Thornber and L. Williams. Analytic solution of stochastic completion fields. *Biological Cybernetics*, 75:141–151, 1996.
26. J. Weickert. *Anisotropic diffusion in image processing*. ECMI. Teubner, Stuttgart, 1998.
27. L. Williams and D. Jacobs. Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation*, 9(4):837–858, 1997.
28. S. W. Zucker, R. Hummel, and A. Rosenfeld. An application of relaxation labelling to line and curve enhancement. *IEEE Trans. Computers*, C-26:393–403, 922–929, 1977.