In *Third IEEE-CS Workshop on Perceptual Organization in Computer Vision* at *International Conference of Computer Vision*, Vancouver, July 2001.

A Field Model for Contour Organization and Partial Differential Equations

Jonas August and Steven W. Zucker   Yale University*

From its earliest roots in Gestalt psychology, thinking about perceptual organization has been predominantly discrete in nature. For example, in one standard grouping demonstration, a subject is presented with a stimulus consisting of perhaps a dozen dots, and the theoretical task is to explain why the subject chose one particular partitioning of the dots among the finite set of possibilities. In illusory contour perception as well, the stimulus is a number of inducing endpoints, where one seek to understand the (perceptual) selection of one of a finite set of arrangements. The "units" of perceptual organization here (dots, endpoints) are discrete, and typically few in number. Artificial intelligence has reinforced this view, where the units of perceptual organization become the "atoms" of LISP programs and symbol manipulation. Modern graph theoretic models of perceptual organization in computer vision have entrenched this discrete thinking, even to the extent that the formal task of perceptual organization has become identified with the grouping of a finite set of (fixed) entities.

Unfortunately, natural images are ambiguous, and not only in terms of the groupings of discrete units: even the existence or absence of these units is uncertain. In contour perception and enhancement, which we consider here, this ambiguity implies that the true space of "units" is the uncountably infinite set of all possible curve groups; explicitly enumerating them is unthinkable. To us this suggests that a continuum or field model is more appropriate, and therefore we work to bring perceptual organization closer to signal processing. Here we outline our framework for contour organization and then describe a rapid but accurate method for solving a partial differential equation used in contour filtering.

*Contact via e-mail: jonas.august@yale.edu, zucker-steven@cs.yale.edu

Framework. Our random field inference framework for contour organization begins with a Markov process model for each contour. Mumford, Williams and co-workers imagined a particle at R_t (in the space of positions (x, y) and directions θ) 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 $Q := \frac{\sigma^2}{2}\frac{\partial^2}{\partial \theta^2} - \cos\theta \frac{\partial}{\partial x} - \sin\theta \frac{\partial}{\partial y} - \lambda^{-1}$ is the generator of directed Markov process R_t, σ bounds the orientation perturbations and λ is the mean of the contour length T. In our framework the Markov process R_t models all image contours, some observed without corruption of any sort (e.g., no noise or blur), some poorly observed (e.g., medical images), and some invisible (e.g., occluded contours, Mumford's original application). Indeed, the particular (stationary) Markov process contour model is unspecified at the level of our framework; more exotic processes, which include scale [Williams, Thornber] or curvature κ [August, Zucker], can be used as well. At this level of generality, the Markov process R_t takes on values ("states") i in state space $\mathcal{I}$, e.g. $i = (x, y, \theta)$ or $i = (x, y, \theta, \kappa)$. A natural image will have an unknown (e.g., Poisson) number $\mathcal{N}$ of contours, $R^{(1)}_{t_1}, \dots, R^{(\mathcal{N})}_{t_\mathcal{N}}$, which are 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 $U = U_i$ of assertions of local contour existence:

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

$$U_i := \sum_{n=1}^{\mathcal{N}} \int_0^{T_n} 1\!\!1\{R^{(n)}_{t_n} = i\} dt_n, \qquad i \in \mathcal{I}.$$

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. 1). Unlike for Markov random fields where computationally intensive Monte-Carlo procedures predominate and exact results are rare, for the CIRF we have analytically derived the following result (subject to a uniformity condition on contour endpoints):

Proposition 1 *The k-th (joint) cumulant of the CIRF at sites $i_1, \dots, i_k$ is proportional to $\sum g_{j_1 j_2} \cdots g_{j_{k-1} j_k}$, where g_{ij} is the expected amount of time the process R_t spent in state j given that it started in i, or $G = (g_{ij}) = -Q^{-1}$, and the sum is over all permutations $j_1, \dots, j_k$ of $i_1, \dots, i_k$.*

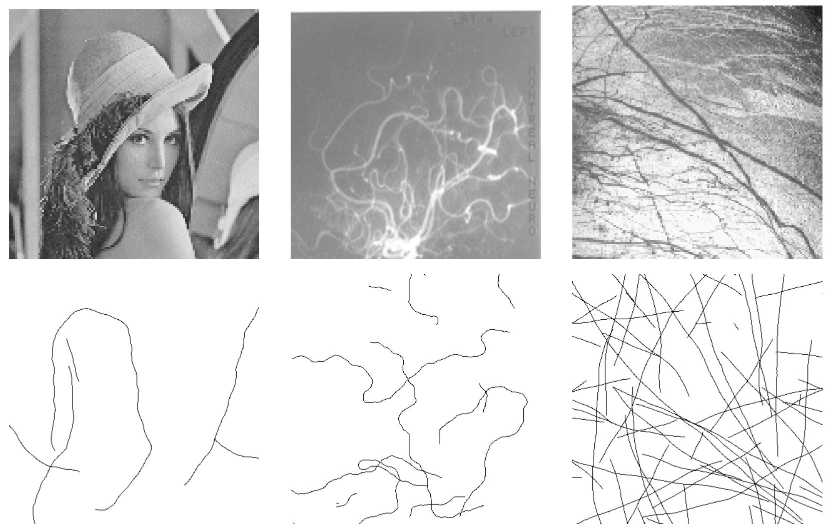

Fig. 1: Observe the similarity of natural images (top: "Lenna," angiogram, ice cracks on Jupiter moon Europa) to random samples of the curve indicator random field (CIRF) for a directed Markov process in (x, y, θ) (bottom: various parameter settings of CIRF). The CIRF acts as a prior for contour organization in our framework.

As the cumulants define the moments, and the moments characterize the distribution, this results provides complete knowledge of the prior distribution $\mathbb{P}(U)$. Given the measurement field M, our goal is to compute that approximation $\tilde{u}$ of the CIRF U that minimizes the mean square error (MMSE):

$$\tilde{u} := \arg\min_{u} \mathbb{E}_M ||u - U||^2 = \mathbb{E}_M U,$$

where $\mathbb{E}_M$ is the expectation operator conditioned on M. Assuming that the likelihood $\mathbb{P}(M|U)$ is local and conditionally independent, and that U is approximately a binary field, we have derived the following nonlinear filter:

Result 1 *A MMSE estimate of the CIRF U given measurements M is:*

$$\mathbb{E}_M U_i \approx f_i b_i, \quad \textit{such that} \quad (Q + \operatorname{diag} c) f + \gamma = 0, \quad (Q^* + \operatorname{diag} c) b + \gamma = 0,$$

where $f = (f_i)$, $b = (b_i)$, $\operatorname{diag} c$ is a diagonal operator, c is a vector dependent on M, γ is a constant field over $\mathcal{I}$, and Q^* is the transpose of Q.

Observe that solving for f and b requires solving two linear partial differential equations (PDEs). Nonetheless, this filter *is* nonlinear, and in two ways. First, the output at i is the product $f_i b_i$, analogous to the source/sink product in Williams' and Jacobs' stochastic completion fields (SCFs). Second, the input comes in as a perturbation of an operator Q or Q^*, and so, for example, $f = (I - G \operatorname{diag} c)^{-1} G\gamma = \sum_{k=0}^{\infty} (G \operatorname{diag} c)^k G\gamma$, since the Green's operator $G = -Q^{-1}$. If we interpret SCFs with field-like inputs c (outside of the original

Fig. 2: Spatial PDE $Qh + w = 0$ was solved for a slightly blurred impulse h, with parameters $\sigma = 1/24, \lambda = 100, \Delta x = \Delta y = 1$, and at discrete directions 4 (top) and 0 (bottom) of 44. Depicted is the integral over θ. Our direct method is isotropic (right), producing a more focused result at discrete direction 4 than does the first-order method (left).

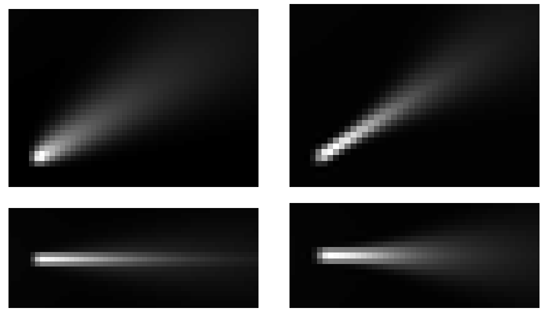

discrete context of illusory contour endpoints of Mumford, Williams and co-workers), then our filter produces the SCF if we stop this sum at $k = 1$. The other terms, polynomial in c at all orders k, make this a (nonlinear) Volterra filter.

A direct method for applying Green's operator. Observe that the above infinite sum for f applying the Green's operator to a field w to compute a field $h = Gw$. For the directed Markov process in (x, y, θ)-space, Williams and co-workers applied Green's operator by (numerically) solving the spatio-temporal diffusion equation $\frac{\partial p}{\partial t} = Qp$ and then (numerically) integrating over time. Here, by applying the Laplace transform (w.r.t time) analytically, and then (numerically) solving the simpler *spatial* (elliptic) PDE $Qh + w = 0$, we eliminate the time dimension, and can solve this PDE directly. In particular, by assuming periodic boundary conditions we can take the discrete Fourier transform (DFT) to obtain a tridiagonal system (solvable in linear time):

$$\pi \left(\frac{im}{M\Delta x} + \frac{n}{N\Delta y} \right) H_{m,n,f-1} + \left(\frac{\sigma^2}{2} f^2 + \lambda^{-1} \right) H_{m,n,f} + \pi \left(\frac{im}{M\Delta x} - \frac{n}{N\Delta y} \right) H_{m,n,f+1} = W_{m,n,f},$$

where the DFTs of h, w, resp., are $H_{m,n,f}$ and $W_{m,n,f}$ at frequencies (m, n, f) out of (M, N, F) total frequencies. This method is more accurate (essentially double precision) than using first-order differencing (Fig. 2), and much faster than the recent isotropic wavelet technique of Zweck and Williams. We are currently using this direct method to implement our nonlinear filter for contour organization, and results on natural images will be reported in [August, *Ph.D. Thesis*, 2001].