PARAMETRIC CONTOUR ESTIMATION BY SIMULATED ANNEALING

*Michael Jamieson, *Paul Fieguth, and †Leo J. Lee

*Dept. of Systems Design Engineering †Dept. of Electrical & Computer Engineering
University of Waterloo, Waterloo, ON, N2L 3G1, CANADA

ABSTRACT

Virtually all implementations of simulated annealing are simplified by assuming discrete unknowns, however *continuous*-parameter annealing has many potential applications to image processing. Widely-scattered problems such as formant tracking, boundary estimation and phase-unwrapping can all be approached as the annealed minimizations of continuous B-spline parameters.

The benefits of simulated annealing are well-known, including an insensitivity to initial conditions and the ability to solve problems with many local minima. Discrete-variable annealing has seen broad application, however continuous-variable annealing is limited by the computational challenge of Gibbs sampling. In this paper we develop efficient approaches to sampling, illustrated in the context of contour tracking in noisy images.

1. INTRODUCTION

The essential task in a wide variety of image-processing problems is estimating the shape and location of ordered irregular curves, or contours. Figure 1 presents three well-known examples of this type: the estimation of ordered formants from speech spectrograms, determining concentric classification boundaries in medical MRI images, and tracking phase-jumps in interferometric SAR images. Though the measurements and the underlying physics or mathematics for these examples are totally unrelated, in all three cases the image processing problem can be characterized similarly as the estimation of smooth, ordered (non-crossing) contours.

To be sure, there is a large literature of specialized methods for formant tracking, medical image segmentation, and phase unwrapping, however the mathematical characterizations of these problems, and others, share many features in common, so it is interesting to consider a unified approach.

Solving a contour estimation problem requires a parameterization of the contours and a performance metric which recognizes good solutions. B-splines [1] are a popular framework for representing smooth, continuous contours

The research support of the Natural Science and Engineering Research Council of Canada is acknowledged.

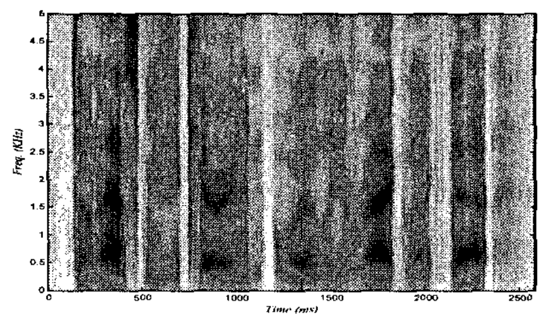

B-spline tracking of speech spectrogram formants.

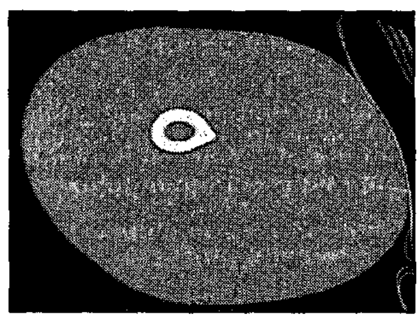

MRI thigh cross-section segmentation.

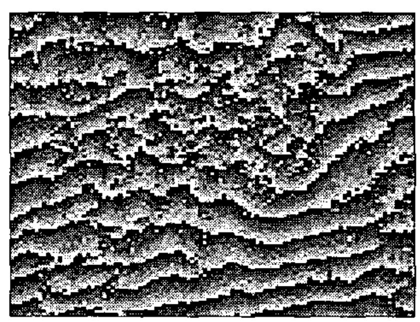

Interferometric SAR image phase contour unwrapping.

Fig. 1. Three examples from speech, medical imaging, and SAR interferometry which share a common underlying ordered-contour structure.

with relatively few parameters. Finding the 'best' set of ordered contours then implies solving an optimization problem in the B-spline parameters, the major obstacle to such optimization being the presence of local minima in the performance metric. Simulated annealing (SA) [2] is widely used to solve image-related optimization problems, due to its initialization insensitivity, avoidance of local minima, and accommodation of a wide variety of energy functions.

However B-spline parameters are *continuous*, whereas simulated annealing is nearly universally employed for *discrete*-valued parameters. A continuous-variable basis requires significant and fundamental changes to SA, in particular the efficient sampling from continuous conditional distributions, which is the key contribution of this paper. A discretization of the continous variable is, of course, a straightforward solution, however we shall see that much more efficient approaches exist.

We begin by developing continuous-parameter models for image processing, followed by an overview of simulated annealing in Section 3. Section 4 describes the main contribution, the continuous-parameter conditional sampling, followed by experimental results.

2. CONTOUR MODELS

Let us start by briefly describing the B-spline optimization problem for contour estimation, unifying the disparate image-processing problems of Figure 1 and formulating them as a continuous-parameter optimization problem, suitable for solving via simulated annealing in Section 3.

We require a performance metric ('energy' function) having two additive components: a prior model, capturing the intrinsic properties of the contours themselves, and a measurement model, defining the relationship between the contours and the image.

Consider the format tracking problem [3]; formant trajectories are ordered continuous functions of time

$$f_n(t) = \sum_{j=-1}^{J+1} c_n(j) \cdot \mathcal{B}^3\left(\frac{t}{\tau} - j\right), \qquad (1)$$

where $\mathcal{B}^3$ is the cubic B-spline generator [1], $c_n(j)$ is the j^{th} control point for the n^{th} formant, τ is the point spacing, and f_n is the n^{th} interpolated trajectory.

We want the contours to follow the dark spectrogram bands, so the external energy H_{ext} rewards "darkness" along the contour trajectories:

$$H_{ext} = -k_1 \sum_{n=1}^{N} \sum_{t=1}^{T} S\left[t, f_n(t)\right]. \qquad (2)$$

H_{ext} alone corresponds to a model-free approach to segmentation or edge detection, and is inadequate because of poor image quality and the multiplicity of possible tracks.

The internal energy function H_{int} specifies the prior or contextual preferences of trajectory shape, to guide preference *away* from unreasonable contours:

$$H_{int} = k_2 H_{smooth} + k_3 H_{range} + k_4 H_{cross}, \qquad (3)$$

a smoothness constraint, an expected frequency range, and a non-crossing condition; together these terms assert a model of formant behavior. Figure 2 illustrates the relationship between contour parameters and energy: the external energy identifies multiple favourable trajectories, however the internal energy constrains the choice to the correct path.

In the MRI application, the concentric classification boundaries can be expressed as closed 1-D radial splines. External energy can be calculated based on edges and the expected gray levels; the internal energy can assert smoothness and an expected thickness range for each tissue layer.

Phase-discontinuity lines in interferometric SAR images result from wrapped phase measurements. The external energy measures local phase continuity after correction; the internal energy function reflects preferences for smoothness and local parallelism.

3. SIMULATED ANNEALING

We wish to find the optimal contour configuration: the global minimum of the overall B-spline energy function. Though the use of SA is well established in hard optimization problems with discrete (or discretized) unknowns, the potential applications of continuous-parameter annealing are relatively unexplored. Unlike existing continuous-parameter annealing methods [4], our approach is developed from the discrete Gibbs Sampler [5]. The set of unknowns (spline control points) has a Gibbs distribution,

$$p(c|T) = \frac{1}{Z} e^{-H(c)/T}. \qquad (4)$$

At the limit $T \to 0$, the distribution becomes impulsive at the *global* energy minimum. In principle then, global optimization reduces to drawing a sample directly from the huge joint distribution $p(c|0)$ at zero-temperature, a computationally impossible operation.

The key to SA is that it is *much* easier to sample from the conditional distribution of a *single* parameter $x = c_i$, as in Figure 2(c). SA then proceeds by selecting T large, so that the parameters are only loosely constrained, then lowering T and more tightly constraining the contours. If the temperature reduction occurs slowly enough the process converges in probability to the global energy minimum [5].

The *key* challenge, and the contribution of this paper, is *how* to perform this conditional sampling.

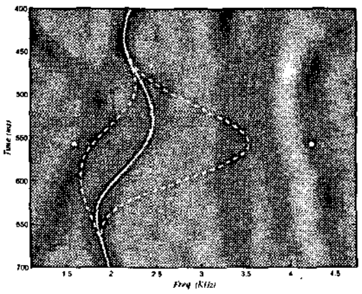

(a)

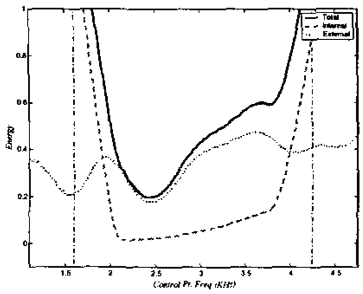

(b)

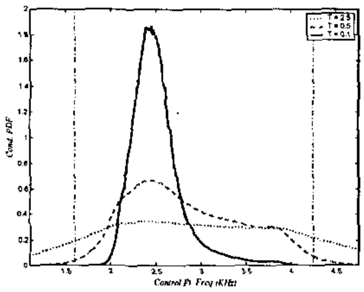

(c)

Fig. 2. (a) B-spline curves from three values for the central control point. (b) Energy as a function of the control point value. (c) The resulting conditional PDF at different temperatures; at low temperatures the distribution narrows.

4. CONDITIONAL SAMPLING

For conditional sampling in the discrete domain, it is simple to find the energy associated with each possible value of a single parameter x and to choose a new setting according to (4). However when x is continuous and the analytical form of $p(x)$ unknown, the only recourse is to make a finite number of $h(x)$ evaluations, form an estimated PDF $\hat{p}(x)$, and sample from it. Since each energy measurement has an associated computational cost, possibly involving a wide variety of image processing operations through H_{ext}, an accurate $\hat{p}(x)$ must be constructed from the smallest possible set of measurements $\{h(x_i)\}$.

Uniformly-spaced x_i is by no means optimal: at low temperatures the distribution is extremely narrow, and most measurements will be wasted at points where $p(x)$ is vanishingly small. By using past $h(x)$ measurements as a guide, an adaptive strategy can achieve a much higher accuracy.

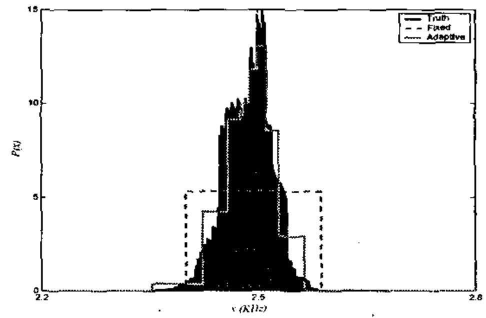

Fig. 3. The central portion of a PDF (filled) is more closely modelled by a 16-sample adaptive estimate (solid) than a 16 uniform sample estimate (dashed).

We begin with a heuristic 'Max-Bracketing' (MB) algorithm for choosing the x_i. If we assume the x_i to be sorted, the block with the largest area,

$$\arg_i \max_i \frac{1}{2}(x_{i+1} - x_{i-1})e^{-H(x_i)/T} \tag{5}$$

representing the greatest integrated uncertainty, is broken by placing new measurements at its left and right edges. Figure 3 illustrates the improvement offered by MB.

While MB is an effective heuristic, more principled approaches are possible. Asserting a prior model for $h(x)$ variation defines the distribution of possible PDFs given $\{h(x_i)\}$ (Figure 4). That is, $p(x|\{h(x_i)\})$ is a random function, the statistics of which can be invaluable in adaptively selecting the next sampling location $\tilde{x}$. For instance, $\tilde{x}$ might subdivide the inter-measurement segment,

$$\arg_i \max_i \int_{x_i}^{x_{i+1}} \overline{p}(x|\{h(x_i)\}) \tag{6}$$

with the greatest mean contribution to the PDF (MPC). Alternatively, $\tilde{x}$ could be placed on the location of greatest uncertainty $var\left[p(\tilde{x}|\{h(x_i)\})\right]$. Such adaptive sampling techniques outperform the regular grid approach at low temperatures, as shown in Figure 5.

Although the computational cost of evaluating the random function statistics is prohibitive for all but the most computationally demanding H_{ext}, the excellent performance motivates a search for equivalent heuristic or analytical approaches. Certainly efficient and robust adaptive sampling techniques are critical to the utility of any continuous Gibbs annealing algorithm.

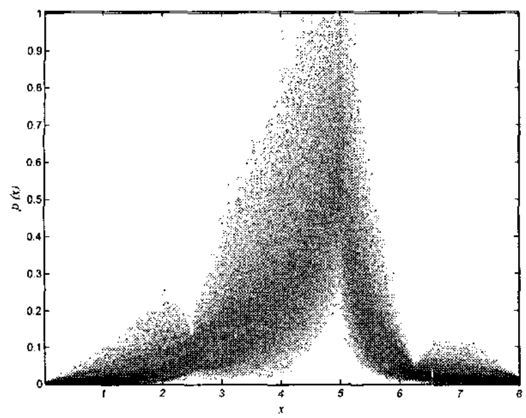

Fig. 4. Given measurements of energy at $x = 0, 2.5, 5.0, 6.2, 8$ the ensemble of probable PDFs. Determining the energy at x_i does not uniquely determine $p(x_i)$

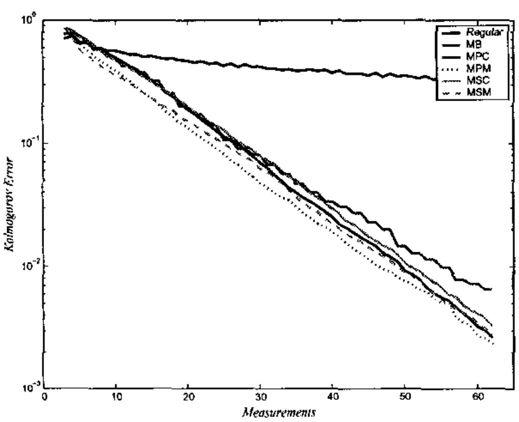

Fig. 5. Closeness to the ideal PDF as a function of the number of energy measurements. The adaptive approaches vastly outperform regular uniform sampling (black).

5. RESULTS & DISCUSSION

Both MB-adaptive and regular-grid versions of the proposed contour-estimation method were implemented for the formant tracking problem. Figure 6 illustrates the convergence of the estimated formant trajectories between intermediate and low temperatures, respectively. Using the regular grid, there was an average distance of 156 Hz between the annealed estimates and hand-labelled formant trajectories. The adaptive approach achieved an average error of 141 Hz in less than a quarter of the time. In both cases, most of the discrepancy occurs in low-confidence regions of the hand-labelled solution.

We see a growing interest in continuous-parameter Gibbs annealing for tackling a variety of image analysis

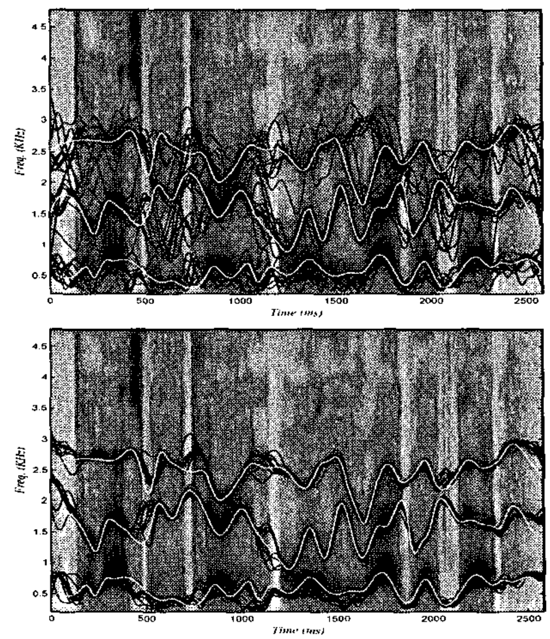

Fig. 6. Formant estimates (dark blue) at medium (top) and low (bottom) temperatures converge toward hand-labelled trajectories (light contours).

problems not amenable to other optimization techniques, where efficient and robust adaptive sampling techniques will clearly play a critical role.

6. REFERENCES

[1] P. Brigger, J. Hoeg, and M. Unser, "B-spline snakes: A flexible tool for parametric contour detection," *IEEE Trans. on Image Process.*, vol. 9, pp. 1484–1496, 2000.

[2] S. Kirkpatrick, S.D. Gellat, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 5, pp. 671–680, 1983.

[3] D. O'Shaughnessy, *Speech Communications: Human and Machine (2nd Edition)*, IEEE Press, 2000.

[4] D. Vanderbilt and S.G. Louie, "A Monte Carlo simulated annealing approach to optimization over continuous variables," *Journal of Computational Physics*, vol. 56, pp. 259–271, 1984.

[5] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. on PAMI*, vol. 6, pp. 721–741, 1984.