A Probabilistic Multi-Scale Model For Contour Completion Based On Image Statistics

Xiaofeng Ren and Jitendra Malik

Computer Science Division
University of California at Berkeley, Berkeley, CA 94720
xren@cs.berkeley.edu , malik@cs.berkeley.edu

Abstract

We derive a probabilistic multi-scale model for contour completion based on image statistics. The boundaries of human segmented images are used as "ground truth". A probabilistic formulation of contours demands a prior model and a measurement model. From the image statistics of boundary contours, we derive both the prior model of contour shape and the local likelihood model of image measurements. We observe multi-scale phenomena in the data, and accordingly propose a higher-order Markov model over scales for the contour continuity prior. Various image cues derived from orientation energy are evaluated and incorporated into the measurement model. Based on these models, we have designed a multi-scale algorithm for contour completion, which exploits both contour continuity and texture. Experimental results are shown on a wide range of images.

1 Introduction

Traditionally there are two approaches to grouping: region-based methods and contour-based methods. Region-based approaches, such as the Normalized Cut framework [19], have been popular recently. Region-based methods seem to be a natural way to approach the grouping problem, because (1) regions arise from objects, which are natural entities in grouping; (2) many important cues, such as texture and color, are region-based; (3) region properties are more robust to noise and clutter.

Nevertheless, contours, even viewed as boundaries between regions, are themselves very important. In many cases boundary contour is the most informative cue in grouping as well as in shape analysis. The intervening contour approach [9] has provided a framework to incorporate contour cues into a region-based framework. However, how to reliably extract contour information, despite years of research, is largely an open problem. Contour extraction is hard, mainly for the following reasons:

1. texture: natural scenes are often highly textured. Contour-based approaches often have difficulty dealing with textured regions and find a lot of false positives, largely because they do not have an inherent concept of texture.
2. low contrast: contrast varies a lot in natural scenes. For example, camouflage of animals. In many cases, contours are perceptually salient only because they form a consistent group.

The problem of contour completion has been studied extensively [4, 14, 17, 13, 20, 5, 18, 2]. Most of these approaches are two-stage: an early stage of detection, where hard decisions are made locally and prematurely; and a later stage of linking or grouping. This two-stage paradigm ignores the crucial fact that contour elements are not independent. A pixel is an edge if and only if there is a contour passing through it. The probability of a pixel being an edge is the posterior probability that there exists a contour passing through this pixel given the image.

Based on this observation, we propose a multi-scale Bayesian approach to contour completion and the classical problem of contour extraction. Two questions we need to answer in a Bayesian framework: a prior model of contour shape; and a local model of contour measurements, how the image arises from contours. These two questions we answer by empirical measurements of contours in natural images. There have been many recent studies on the statistics of natural images [22, 8, 7, 11, 3] . In our work we use the database of human segmented images reported in [11]. The contours in these segmentations are explored to understand natural scenes and to motivate our contour completion algorithm. Driven by this empirical analysis in Section 3, *higher order Markov models* are proposed as the prior model for contour shape. This is a significant distinction in our work from the related approaches such as Mumford [13], Williams and Jacobs [21], who used a first-order Markov model for contour shape. We also make use of the database of human segmented images to arrive at a measurement model, incorporating various local cues such as orientation energy and textureness. Based on the Markov assumption, we use dynamic programming to efficiently compute the posterior probability. The multi-scale contour completion algorithm is presented in Section 5. Experimental results are shown in Section 6.

2 Bayesian Contour Completion

In this section we give a formal analysis of our intuition in the introduction and motivate our work on images statistics. The key is that how likely a pixel is an edge is quantified by the posterior probability that there exists a contour passing through it:

Consider a pixel p in a given image I. Let M_p denote the measurement, or a feature vector, at the pixel p, and M the collection $\{M_p\}$ for all p. Let $\mathbf{b}_p \in \{0, 1\}$ be the binary random variable which denotes the existence of a boundary contour at pixel p. What we want to compute is the posterior distribution:

$$P(\mathbf{b}_p|M) = \frac{P(M, \mathbf{b}_p)}{P(M)} \propto P(M, \mathbf{b}_p) \tag{1}$$

The posterior probability of the non-contour case is determined by our background model. If we make the simplifying assumption that $\mathbf{b}_p = 0$ does not constrain the existence of contours at other pixels, we have

$$\begin{aligned} P(\mathbf{b}_p = 0|M) \propto & P(M|\mathbf{b}_p = 0)P(\mathbf{b}_p = 0) \\ = & P(\mathbf{b}_p = 0)P(M_p|\mathbf{b}_p = 0)P(M_{I\setminus\{p\}}) \end{aligned} \tag{2}$$

By the law of large numbers, the marginal probability M_Λ can be well approximated by an exponential function $P(M_\Lambda) \approx r^{|\Lambda|}$, where $|\Lambda|$ is the number of pixels in the subset Λ, and r is the expected likelihood.

In the contour case $P(\mathbf{b}_p = 1)$, we can not make the independence assumption. Indeed, $\mathbf{b}_p = 1$ if and only if there is a boundary contour passing through the pixel p. Let $\mathcal{C} = \{\gamma |\ \gamma$ is a curve passing through $p\ \}$ be the collection of all such contours. Since one and only one such contour passes through p in the image I, we have

$$P(M, \mathbf{b}_p = 1) = \sum_{\gamma \in \mathcal{C}} P(M, \gamma)$$

For each such curve γ,

$$P(M, \gamma) = P(M|\gamma)P(\gamma)$$

What we need in this probabilistic formulation is:

1. $P(\gamma)$, the prior model on contour shape;
2. $P(M|\gamma)$, the model of local image measurements conditioned on the presence or the absence of a contour.

3 Modeling Contour Shape

The key fact we need to keep in mind when studying contours, or any object in natural images, is that these objects are multi-scale. This multi-scale phenomenon mainly arises from two sources:

1. Objects in the natural world are themselves multi-scale. For example, a object has parts: the parts in the figure of a person include nose, head, torso, arm, hand, finger, etc., all of which are different in scale.
2. Despite the possible bias introduced by the observer, objects are usually viewed from an arbitrary distance and angle.

3.1 Scale Invariance

Scale invariance in natural images has been reported by various authors. An in-depth study of scale invariance in boundary contours is beyond the scope of this paper. We study one phenomenon here: contours consist of segments, which roughly correspond to the parts of objects or the scale of local details. We consider the decomposition of contours at extremal points, i.e., the points whose curvatures are locally maximal. Figure 1(a) shows some examples of this decomposition. The distribution of the length of the resulted contour segments reveals to us properties of the underlying mechanism which generates the contours. For example, if a first-order Markov model were accurate, then this distribution of segment length, or the time to wait until a high curvature event occurs, is exponential.

Figure 1 (b) shows the distribution of contour segment length. Ignoring the range where the length is small, in which the decompositions are not reliably, we observe from this distribution the following power law:

$$\text{frequency} \propto \frac{1}{(\text{contour length})^{1.994}}$$

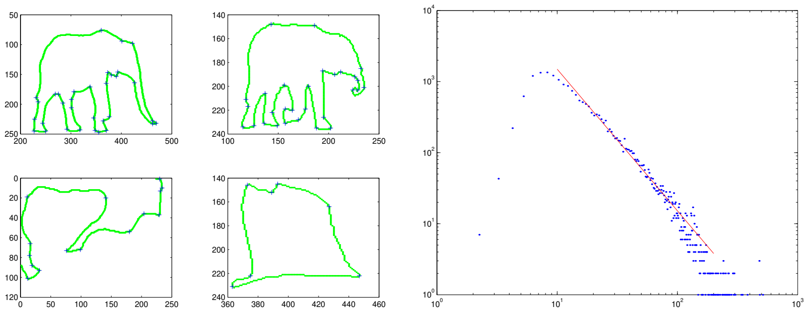

(a) (b)

Fig. 1. Examples of the contour decomposition. High curvature points, where the contours are segmented, are shown in the plots. The decompositions are consistent with human perception.

This power law is consistent with the distribution of region area [11], and with the intuition that natural objects are self-similar. We can formalize this intuition and give a simple explanation: suppose we have a contour of fixed length l_0 at a base scale. When viewed at scale s (say $s = d$ where d is the distance to the observer), its apparent length is l_0/s. At the same time, because the field of view is s^2 times larger, the probability of observing this contour increases by s^2. Hence the frequency of observing this contour of an apparent length l is proportional to $s^2 = (l_0^2)/(l^2)$. I.e., this contour induces a whole distribution:

$$f(l) \propto \int_s s^2 \delta_{l_0}(ls)ds = \frac{l_0^2}{l^2}$$

We take the expectation w.r.t. l_0 to obtain the overall distribution. This does not change the structure of the power law:

$$\bar{f}(l) = E[f(l)] \propto \frac{E[l_0^2]}{l^2}$$

it still decreases quadratically with contour length l.

We can compare this inverse square law with the predictive first-order Markov model as used by Mumford [13], Williams and Jacobs [21]. In their work, curvature is assumed to be white noise; hence the tangent direction is a Brownian motion. That would imply that the frequency a contour appears decreases exponentially with its length. Exponential models are common; they have the memoryless property and are easy to work with. We would like to assume that. However, *This is not true empirically.*

The next question we want to explore is *self-similarity*. We use region area as an indication for the scale of the object. We study whether the distribution of the contour segment length is the same for different ranges of region area. Our results, which we omit here, show that the distributions are almost identical for groups of different region sizes. This result justifies the intuition that objects themselves are multi-scale and self-similar in nature. *It suggests that any algorithm for contour completion should be intrinsically multi-scale.*

3.2 Higher-Order Markov Models

We have seen in the previous section that the Markov assumption is not accurate. In this section we extend the first-order Markov model to high-order ones, and measure the information these higher-order Markov models convey.

Let $C(\cdot)$ be the representation of the curve. This could be an intrinsic representation, e.g., curvature parameterized by curve length. Or, in the context of contour completion, to make the computation easier we represent the curve by its tangent directions parameterized by time. We adopt the random process view of contour generation, in which we predict $C(t+1)$ based on the information we have up to time t.

We extend the basic Markov model over scale. Let $C^{(0)} = C$ be a curve γ at the base scale. $C^{(1)}, \cdots, C^{(k-1)}$ are the scaled versions of the original curve γ. Define the k-th Markov model over scales to be:

$$P(C^{(0)}(t+1)|\gamma) = P(C^{(0)}(t+1)|C^{(0)}(t), C^{(1)}(t), \cdots, C^{(k-1)}(t)) \tag{3}$$

Figure 2 shows the information gain when we extend the Markov model over scales. We observe that there is a substantial gain from combining orientation information at coarser scales. The use of higher-order Markov models is empirically justified by cross-validation (Figure 2 (b)). The intuition is that scaled curves, at coarser scales, combine information in a neighborhood. For example, if $C^{(1)}(t)$ is to the left of $C^{(0)}(t)$, it makes a left turn at t and therefore $C^{(0)}(t+1)$ is more likely to turn to the left. This intuition has been confirmed empirically from samples of conditional distributions, which we omit here. Since long-range dependencies in contours is likely caused by interactions at coarser scales, these models over scales are a natural choice for local contour modeling. This is not intended to be a global model of contours, since in these models there is no notion of topological constraints, such as closure and no self-intersection; nevertheless, it is sufficient for our purpose of contour completion.

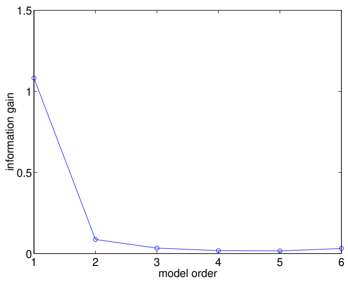

(a)

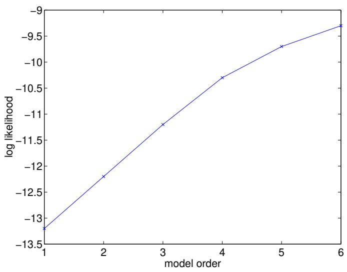

(b)

Fig. 2. (a) Information gain as the order of our Markov models increase. The model of order 1 corresponds to the traditional Markov model. (b) 10-fold cross validation of higher-order Markov models. The log-likelihood is normalized by the length of contours.

4 Local Measurement Model

We have seen the use of higher order Markov models to represent the prior distribution of contour shape. In this section we turn to the images themselves. We study how the contours locally give rise to image measurements. We formalize the problem as computing the local posterior distribution of contour elements. In Section 4.1 we derive the posterior model of local contour detection from image statistics of both the orientation energy and textureness. In Section 4.2 we derive a probabilistic model of positional uncertainty to go beyond the assumption of conditional independence. And in Section 4.3 we derive in a similar fashion a model of local orientation uncertainty.

4.1 Detection Uncertainty

The use of orientation energy (e.g., [12, 16]) has been an essential part of computer vision. The orientation energy at an angle θ is defined as:

$$OE_\theta^2 = (I * f_{1,\theta})^2 + (I * f_{2,\theta})^2 \tag{4}$$

where $f_{1,\theta}$ and $f_{2,\theta}$ are the second Gaussian derivative and its Hilbert pair oriented at the angle θ. It has roots in biological vision, and it has proven to be extremely successful in both contour and texture analysis. However, people have also realized the difficulties which come with the success. One major problem is the wide existence of textured regions, which typically have high responses to traditional contour-based techniques. Recently people have started to look into the interaction of texture cues with contour cues (e.g., Malik et. al. [10]). They introduced the notion of $p_{texture}$, based on the χ^2 distance of texton histograms (see detailed explanations in their paper):

$$p_{texture} = 1 - \frac{1}{exp[-(\chi_{LR}^2 - \tau)/\beta]} \tag{5}$$

They target at suppressing OE responses in homogeneous texture regions, by multiplying p_{con}, a non-linear transform of OE, by $p_{texture}$. In this paper we extend the notion of $p_{texture}$ to be orientation dependent; i.e., χ_{LR}^2 is computed for each angle θ. Let $p_{con,\theta} = 1 - exp(-OE_\theta/\sigma_{IC})$ as in [10] . We gate $p_{con,\theta}$ with $p_{texture,\theta}$ separately in each angle θ

$$p_{b,\theta} = p_{con,\theta} * (1 - p_{texture,\theta}) \tag{6}$$

After the gating, we take the maximum $p_b^* = \max p_{b,\theta}$. Let θ^* denote the angle where $p_{b,\theta}$ achieves the maximum, i.e., $p_{b,\theta^*} = p_b^*$.

In Figure 3, we show the power of this cue combination. At the back of the zebra, the orientation energy is strong in the vertical direction. This is not the true direction of the boundary contour (in the scale where the human segmentations are done). However, $p_{texture}$ is low only at the horizontal direction. This enables us to locally detect the maximum p_b^* in the correct direction.

Figure 4 shows the statistics of $p_{con,\theta}$ and $p_{texture,\theta}$. The likelihood ratio in Part (a) is used in our contour completion algorithm. Part (c) and (d) justifies the multiplicative form of our definition of $p_{b,\theta}$. We also empirically measure the information gain of

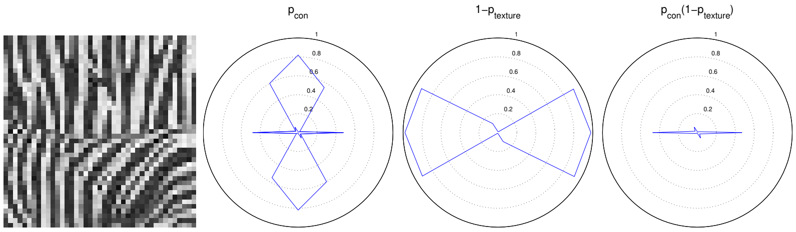

(a) (b) (c) (d)

Fig. 3. Illustration of $p_{texture,\theta}$. (a) a patch from a zebra image; we consider the pixel at the center. (b) the distribution of raw $p_{con,\theta}$ (not gated with $p_{texture}$), maximal in the vertical direction. (c) $1 - p_{texture,\theta}$, maximal in the horizontal direction. (d) the product $p_{b,\theta} = p_{con,\theta}(1 - p_{texture,\theta})$, sharply peaked in the horizontal direction.

the cues from histograms of the joint distribution. In our experiments, the marginal entropy of $\mathbf{b}_p$, *edge* vs. *non-edge*, is 0.2492 (in bits); the information gain of $p_{con,\theta}$ is 0.0425; and the information gain of $p_{texture,\theta}$ over $p_{con,\theta}$ is 0.0091. These information-theoretic measures illustrate the relative importance of these local cues.

4.2 Positional Uncertainty

So far in our local model, the posterior probability p_b^* is estimated independently at each pixel. We have ignored the important relationship: the correlation between neighboring pixels. Traditionally this is done by non-maximum suppression. However, due to image noise we can never be certain about the contour localization. There is always a positional uncertainty. To be consistent with our philosophy of avoiding local hard decisions, we again make use of image statistics to derive a probabilistic model of positional uncertainty, which serves as a soft non-maximum suppression.

We base our analysis on the quadratic model originally proposed in [15]. Empirically we have found it to perform well on real images. Their approach is to fit a parabola in a local neighborhood, and use the information from the parabola fitting, such as d, the distance to the center of the parabola, and c, the curvature of the parabola. Details are omitted here. In Figure 5 (a)-(c) we show the statistics of the various outputs from this parabola fitting. In this study of non-maximum suppression, we compare the statistics of edge pixels to near-edge pixels (i.e., pixels that are within a distance of 4 from a human marked contour). The results are qualitatively different from what we have seen in the previous section, where we compare edge pixels to non-edge pixels. p_b^* alone is not a good cue for non-maximum suppression. Instead, both the distance d and the curvature c are informative. For edge pixels, the curvature tends to be negative, and the distance to the center of the parabola tends to be small. In this work we choose d, the distance , for our model of positional uncertainty. We observe the following relationship

$$P(\mathbf{b}_p|M_p, d_p) = P(\mathbf{b}_p|M_p)\frac{P(d_p|M_p, \mathbf{b}_p)}{P(d_p|M_p)}$$

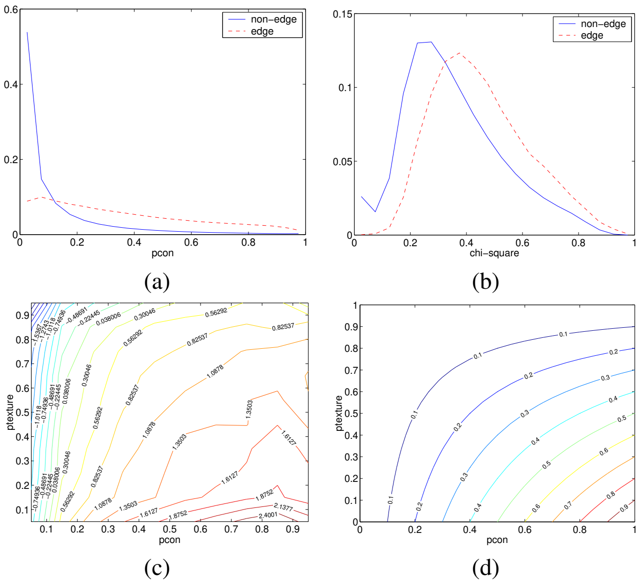

Fig. 4. Image statistics of orientation energy ($p_{con,\theta}$) and textureness ($p_{texture,\theta}$). (a) the marginal distributions of $p_{con,\theta}$, both the edge case and non-edge case. (b) the marginal distributions of the χ^2 distance used in Equation eqrefeq:ptexture to produce of $p_{texture,\theta}$. (c) the contour plot of the likelihood ratio $\frac{P(\cdot|edge)}{P(\cdot|non-edge)}$. (d) the contour plot of the multiplicative cue combination (Equation (6)).

Accordingly we collect the statistics for the ratio $P(d_p|p_b^*, \mathbf{b}_p = 1)/P(d_p|p_b^*, \mathbf{b}_p = 0)$, as shown in Figure 5 (d). Motivated by this statistics , we choose the following model to update p_b^*:

$$p_b^* \leftarrow p_b^* \left[(\alpha_1 + \alpha_2 d)^{-\beta_1 - \beta_2 p_b^*} \right] \tag{7}$$

In this parametric model, when d is large the multiplicative factor is small, and p_b^* is suppressed. This suppression is more significant when p_b^* is large.

4.3 Orientation Uncertainty

In the previous sections we have studied the posterior probability of $\mathbf{b}_p$, i.e., the existence of a boundary contour at a pixel, as a function of $p_{con,\theta}$, $p_{texture,\theta}$, and d (the distance to the center of a local parabolic fit). However, due to image noise and further complications such as junctions, the true contour might not be oriented at θ^*, the maximum $p_{b,\theta}$ orientation. We represent the orientation uncertainty with a distribution $\{p_\theta\}$ over θ at each pixel. We further assume that this distribution is Gaussian, and θ^* is an unbiased estimate of the peak. We use two features to estimate the variance σ_θ^2: p_b^* and $p_b^\perp$. $p_b^\perp$ is defined as follows: if θ^* is the angle of the maximum $p_{b,\theta}$, we choose the angle $\theta^\perp$ perpendicular to θ^* and define $p_b^\perp = p_{b,\theta^\perp}$. Figure 6 (a) shows the statistics of the variance of local orientation for each $(p_b^*, p_b^\perp)$ pair. Motivated by this statistics,

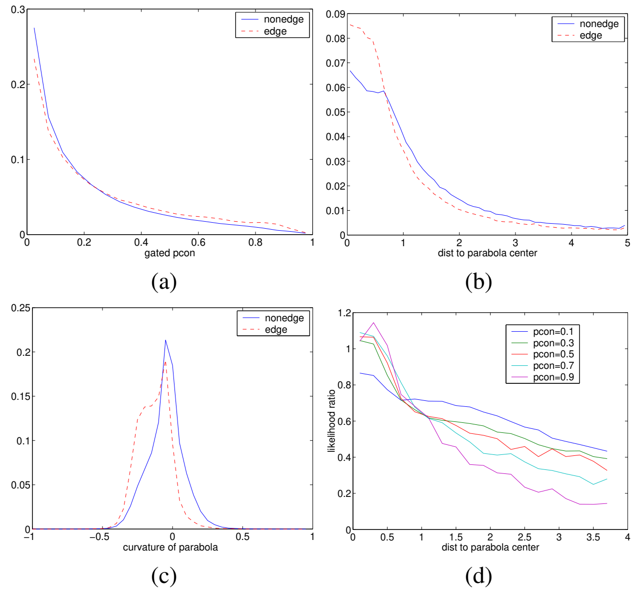

Fig. 5. Image statistics for probabilistic non-maximum suppression. (a) the marginal distributions of $Pcon$, edge vs. non-edge. Note here we are comparing edge pixels with near-edge pixels, hence the orientation energy alone gives us little information. (b) the marginal distributions of the distance d in parabola fitting. c) the marginal distributions of the curvature c in parabola fitting. (d) the distributions of the likelihood ratio, as a function of the distance d, conditioned on p_b^*.

we use a simple parametric model as shown in Figure 6 (b) :

$$\sigma_\theta^2 = \exp(-\beta \frac{p_b^*}{p_b^\perp + \alpha}) \tag{8}$$

When p_b^* is large, the uncertainty is low; similarly when $p_b^\perp$ is large, the uncertainty is high. This simple parametric model has the desired properties; clearly other models could be used as well. Once we obtain an estimate of σ_θ^2, we distribute the probability mass p_b^* over θ according to the uncertainty.

There are three important cases here: (1) when the contrast is high, we have a reliable estimate of local orientation; (2) when the contrast is low, local orientation can be arbitrary; (3) at a junction, the $p_{b,\theta}$ profile is complicated. Building a local junction detector is extremely difficult. Our approach is to make a soft decision locally, and let contour completion to find the most probable orientation. This frees us from searching for a precise solution for the distribution of orientation $\{p_\theta\}$.

5 Multi-Scale Contour Completion

To motivate the multi-scale model from a practical viewpoint, Figure 7 shows the output of a single-scale contour completion algorithm on the same image scaled to different

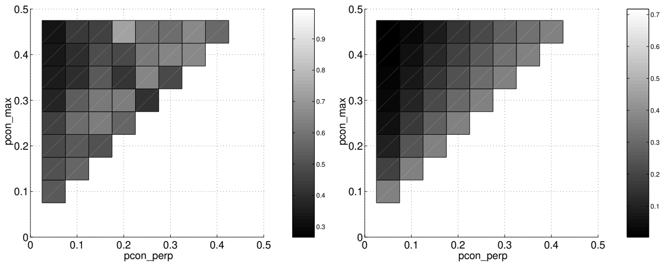

(a) (b)

Fig. 6. The uncertainty of local orientation σ_θ^2 as a function of p_b^* and $p_b^\perp$. (a) the statistics from real images. (b) fitting the parametric model (Equation (8)).

sizes. The results are dramatically different. At finer scales, the algorithm is unable to complete contours over a long distance in the image. At coarser scales, the algorithm is unable to properly detect and enhance curved contours. For any model based on the Markov assumption and the underlying preference for straight lines, no matter how the parameters are tuned, it only works at a certain scale, therefore unable to handle the wide range of scales existing in natural images.

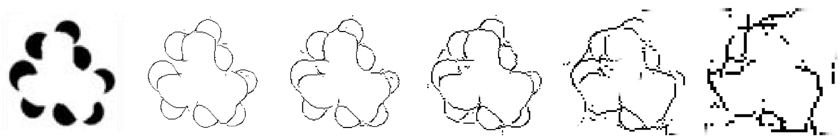

Fig. 7. A synthetic image and the outputs of a single-scale contour completion algorithm on this image at different sizes. The figures correspond to completion over increasingly coarse scales.

We will start with a single-scale version of our probabilistic contour completion algorithm and then extend it to multi-scale. The multi-scale property comes in two ways: (a) contour propagation is done at each scale and the results are combined to produce the final estimate; (b) prior models at finer scales are conditioned on the results from coarser scales.

5.1 Single-Scale Version

We represent a curve γ by $\{q(t), T(t)\}$, where $q(t)$ denote the locations and $T(t)$ the tangent directions of γ at $q(t)$. Let $P_{\text{length}}(n)$ be the distribution of contour lengths and $|\gamma|$ the length of the curve γ.The prior probability of γ can be written as

$$P(\gamma) = P_{\text{length}}(|\gamma|)P(T(\cdot))P(\gamma(\cdot)|T(\cdot)) \tag{9}$$

The conditional probability

$$P(M|\gamma) \approx P(M_{I\setminus\{\gamma\}}) \prod_t P(M_\gamma|T(t)) \tag{10}$$

We approximate $P(M_{I\setminus\{\gamma\}})$ by $r^{|I|-|\gamma|}$, where r is the expected likelihood in Section 2. Apply the Markov assumption here; let

$$\begin{aligned} g_t^+ &= P_{location}(q(t+1)|q(t), T(t)) \\ g_t^- &= P_{location}(q(t-1)|q(t), T(t)) \\ h_t^+ &= \quad P_{tangent}(T(t+1)|T(t)) \\ h_t^- &= \quad P_{tangent}(T(t-1)|T(t)) \\ m_t &= P_{measurement}(M_{q(t)}|T(t))/r \end{aligned}$$

Then we have the likelihood ratio:

$$\begin{aligned} L(M,\gamma) =& \frac{P(M,\gamma)}{P(M, \mathbf{b}_p = 0)} \propto P(\gamma)\frac{P(M_\gamma|\gamma)}{P(M_\gamma|\mathbf{b}_{\gamma(\cdot)} = 0)} \\ \approx& P_{\text{length}}(|\gamma|)P(T(0))\frac{P(M_p|T(0))}{P(M_p|\mathbf{b}_p = 0)} \\ &\times \prod_{t<0} g_{t+1}^- h_{t+1}^- m_t \prod_{t>0} g_{t-1}^+ h_{t-1}^+ m_t \end{aligned} \tag{11}$$

This is a Hidden Markov Model with $q(t)$ and $T(t)$ as hidden variables. Because of the one-dimensional nature of contours, we can apply dynamical programming to solve this computational problem, which is essentially the same as the alpha-beta algorithm or the stochastic completion approach in [21]. The details are omitted here. Let γ^+ denote the partial contour $\{\gamma(t); t > 0\}$, and γ^- for $\{\gamma(t); t < 0\}$. Let $\phi(p, \theta, k)$ be the sum of the messages arriving at (p, θ) at step k. We maintain:

$$\begin{aligned} \phi(p, -\theta, k) &\propto P(M, \exists\gamma \text{ through } (p,\theta) \text{ s.t. } |\gamma^+| = k) \\ \phi(p, \theta, k) &\propto P(M, \exists\gamma \text{ through } (p,\theta) \text{ s.t. } |\gamma^-| = k) \end{aligned}$$

$\phi(p, \theta, k)$ are recursively computed using propagation and diffusion as in [21, 20]. Given ϕ, we can calculate the likelihood ratio $L_p = \frac{P(M, \mathbf{b}_p = 1)}{P(M, \mathbf{b}_p = 0)}$ as:

$$\begin{aligned} L_p \propto& \sum_\theta P(T(0) = \theta)P(M_p|T(0) = \theta) \\ &\times \sum_{n>1}\left(P_{\text{length}}(n)\sum_{k=1}^{n-2}\phi(p,\theta,k)\phi(p,-\theta,n-k-1)\right) \end{aligned} \tag{12}$$

With this probabilistic interpretation, we can incorporate contour cues, avoid making premature hard decisions, and readily extend the model to multi-scale.

5.2 Multi-Scale Version

At the coarsest scale, the single-scale version of the contour completion algorithm is applied. Suppose we have obtained the posterior distributions $P(\mathbf{b}_p^{(s)} = 0|M^{(s)})$ and $P(T_p^{(s)}|M^{(s)})$ for scales $s = 1, \cdots, k-1$. (note $\sum_\theta P(T_p^{(s)} = \theta|M^{(s)}) = P(\mathbf{b}_p^{(s)} = 1|M^{(s)})$.)

The prior distributions in Equation (9) is now all conditioned on the coarser scales. We see that the message passing algorithm remains the same, except that the prior of tangent directions $P_{\text{tangent}}(\theta'|\theta)$ now becomes

$$P^{(k)}_{\text{tangent}}(\theta'|\theta) = E_{\{T^{(s)}_q, s<k\}}\left[P_{\text{tangent}}(\theta'|\theta, \{T^{(s)}_q, s < k\}\right] \tag{13}$$

We refer to this multi-scale processing as *multi-scale conditioning*. In practice, we simplify the computation by conditioning $P^{(k)}_{\text{tangent}}$ at scale k only on $T_q^{(k-1)}$ at scale $k-1$, and use the maximum-probability direction $T_q^{*(s)}$ to replace the expectation over all the possible directions $\{T_q^{(s)}\}$. Some results obtained by the multi-scale algorithm are shown in the next section. Finally, to combine the results from individual scales, we have

$$P(M, T_p = \theta) = \sum_s P_{\text{scale}}(s) P^{(s)}(M, T_p = \theta)$$

where we use the self-similar $1/s^2$ distribution in Section 3.1 for P_{scale}.

6 Experiments and Evaluation

Figure 8 shows our results on a few simple synthetic images with subjective contours. Figure 9 shows the use of multi-scale completion on a synthetic image. Figure 10 shows the use of multi-scale conditioning for real images. Figure 11 shows the completion results at various scales. And Figure 13 contains more completion results. Please see the interpretations therein.

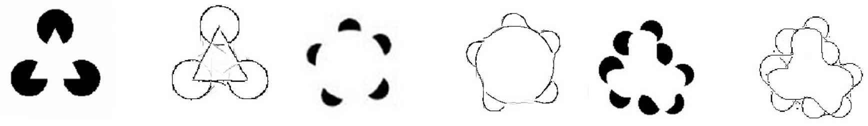

Fig. 8. Results on a few classical synthetic images.

We have quantitatively evaluated the performance of our algorithm on a large set of images. We again turn to the ground truth from our segmentation database. The problem of edge detection by itself is a classification problem. We define the following two performance measures: *hit rate*, the number of correctly labeled edge pixels divided by the total number of edge pixels; and *false rate*, the number of labeled pixels which are not true edges divided by the total number of labeled pixels. Figure 12 shows our results. The plots correspond to two different tolerance thresholds on localization error. We declare a hit when a labeled pixel is within a distance of n from a true contour. We can observe from the curves, that our algorithm consistently outperforms the raw OE cue. This advantage increases when the threshold n increases.

In summary, this paper has made the following contributions:

(a) (b) (c) (d)

Fig. 9. A synthetic example of the use of multi-scale conditioning. (a) the original input. (b) the completion at the coarse scale. (c) the completion at the fine scale without using the conditioning in Equation (13). (d) the completion with the conditioning.

(a) (b) (c) (d) (e)

Fig. 10. A further example of the use of multi-scale conditioning. It is a patch extracted from the penguin image in Figure 11. (a) the original image. (b) the completion at the coarse scale. (c) the completion at the fine scale without conditioning. (d) the completion with conditioning. (e) the posterior log-likelihood at the pixels along the central low-contrast contour. Both algorithms successfully complete the low-contrast contour in the middle of the patch. The signal is significantly enhanced by the use of multi-scale conditioning at low-contrast locations.

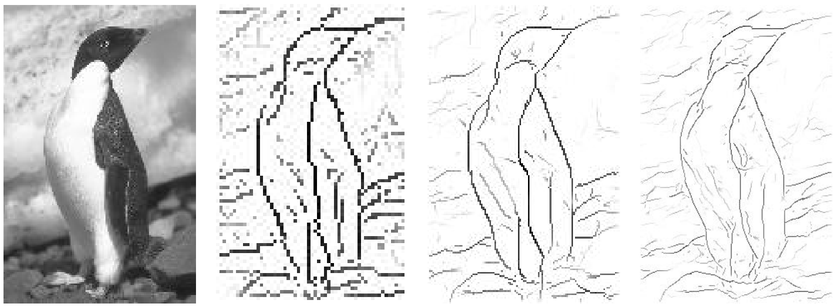

Fig. 11. A complete example of multi-scale contour completion. The completion is done at three distinct scales, from coarse to fine. We notice that the results are qualitatively different. At the coarse scale, large gaps of low contrast contours are easily completed, and noise are generally suppressed. But there are some details which we can only see at fine scales.

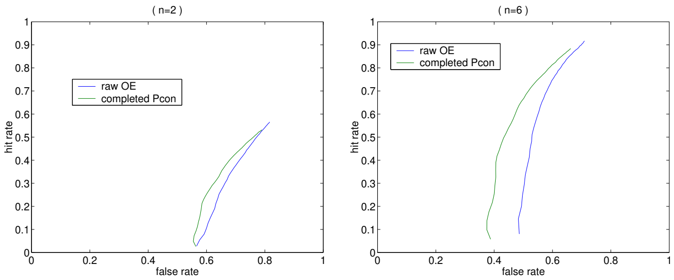

(a) (b)

Fig. 12. ROC-like curves for performance evaluation. (a) with $n = 2$ (see text). (b) with $n = 6$.

1. use ground truth on human segmented images to establish prior and measurement models for boundary contours in natural scenes;
2. propose a multi-scale probabilistic model for contour completion;
3. evaluate the power of our contour completion algorithm on a wide range of images.

In this paper we made somewhat *ad hoc* choices of parametric forms for the terms in the local model of image measurements. These can be replaced by suitable non-parametric estimates.

Acknowledgements. We thank Yair Weiss for discussions on the probabilistic formulation of contour completion. The local models for contour measurements are related to joint work with Charless Fowlkes and David Martin. This research was supported by NSF through a Digital Library Grant IRI-9411334.

References

1. S. Belongie and J. Malik. Finding boundaries in natural images: A new method using point descriptors and area completion. In *Proc. 5th Euro. Conf. Computer Vision*, Freiburg, Germany, June 1998.
2. J.H. Elder and S.W. Zucker. Computing contour closures. In *Proc. Euro. Conf. Computer Vision*, volume I, pages 399–412, Cambridge, England, Apr 1996.
3. W. S. Geisler, J. S. Perry, B. J. Super, and D. P. Gallogly. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41:711–724, 2001.
4. S. Geman and D. Geman. Stochastic relaxation, gibbs distribution, and the bayesian retoration of images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6:721–41, Nov. 1984.
5. G. Guy and G. Medioni. Inferring global perceptual contours from local features. *Int'l. Journal of Computer Vision*, 20(1-2):113–33, Oct. 1996.
6. F. Heitger and R. von der Heydt. A computational model of neural contour processing. In *Proc. Int. Conf. Computer Vision*, pages 32–40, Berlin, Germany, May 1993.
7. J. Huang and D. Mumford. Statistics of natural images and models. In *Proc. IEEE Conf. Comput. Vision and Pattern Recognition*, pages I:541–547, 1999.
8. S. Konishi, A.L. Yuille, J.M. Coughlan, and S.C. Zhu. Fundamental bounds on edge detection: An information theoretic evaluation of different edge cues. pages I:573–579, 1999.

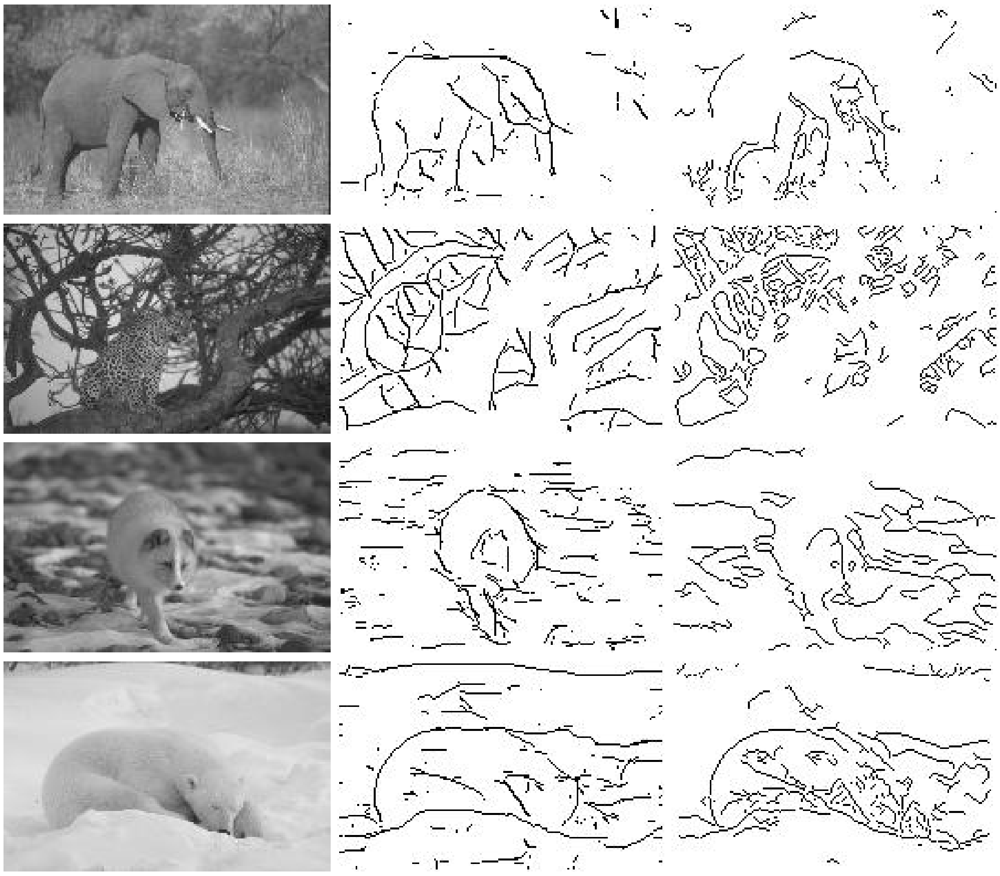

(a) (b) (c)

Fig. 13. More results on real images with low contrast contours and stochastic textures. (a) the original images, at the resolution of 179×115. (b) the thresholded results of our algorithm. (c) the Canny edges with approximately the same number of pixels. Notice how the algorithm recovers low-contrast contours and suppresses noise, for example, in the first row the back and the left leg of the elephant; in the second row the branches on the right and the cheetah in the middle; and in the third row the outline of the wolf.

9. T. Leung and J. Malik. Contour continuity in region-based image segmentation. In H. Burkhardt and B. Neumann, editors, *Proc. Euro. Conf. Computer Vision*, volume 1, pages 544–59, Freiburg, Germany, June 1998. Springer-Verlag.
10. J. Malik, S. Belongie, T. Leung, and J. Shi. Contour and texture analysis for image segmentation. *Int'l. Journal of Computer Vision*, 43(1):7–27, June 2001.
11. D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l. Conf. Computer Vision*, volume 2, pages 416–423, July 2001.
12. M.C. Morrone and R.A. Owens. Feature detection from local energy. *Pattern Recognition Letters*, 6:303–13, 1987.
13. D. Mumford. Elastica and computer vision. In Chandrajit Bajaj, editor, *Algebraic Geometry and Its Applications*, pages 491–506. Springer Verlag, 1994.
14. P. Parent and S.W. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(8):823–39, Aug. 1989.
15. P. Perona and J. Malik. Boundary detection using quadratic filters: Performance criteria and experimental assessment. In *SPIE*, 1708.
16. P. Perona and J. Malik. Detecting and localizing edges composed of steps, peaks and roofs. In *Proc. Int. Conf. Computer Vision*, pages 52–7, Osaka, Japan, Dec 1990.
17. L. Rosenthaler, F. Heitger, O. Kubler, and R. von der Heydt. Detection of general edges and key-points. In *Proc. 2nd Europ. Conf. Comput. Vision, G. Sandini (Ed.), LNCS-Series Vol. 588, Springer-Verlag*, pages 78–86, 1992.
18. E. Sharon, A. Brandt, and R. Basri. Completion energies and scale. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(10):1117–1131, October 2000.
19. J. Shi and J. Malik. Normalized cuts and image segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 731–7, San Juan, Puerto Rico, June 1997.
20. L. Williams, J. Zweck, T. Wang, and K. Thornber. Computing stochastic completion fields in linear-time using a resolution pyramid. *CVIU*, 76(3):289–297, December 1999.
21. L.R. Williams and D.W. Jacobs. Stochastic completion fields: a neural model of illusory contour shape and salience. In *Proc. 5th Int. Conf. Computer Vision*, pages 408–15, Cambridge, MA, June 1995.
22. S.C. Zhu. Embedding gestalt laws in markov random fields. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 21(11):1170–1187, November 1999.