MIT CSAIL

6.869 Advances in Computer Vision

Fall 2018

Problem Set 6: Belief Propagation

Posted: Thursday, Oct 18, 2018 **Due:** Thursday, Nov 1, 2018

Please submit **two files:** 1) a **PDF** report named {your_kerberos}.pdf, including your answers to all required questions with images and/or plots showing your results, and 2) a file named {your_kerberos}.zip, containing relevant source code.

Late Submission Policy: We do not accept late submissions. The submission deadline has a 50-minute soft cut-off; after after midnight Thursday, submissions are penalized 2% per minute late.

Problem 1 *Markov Network*

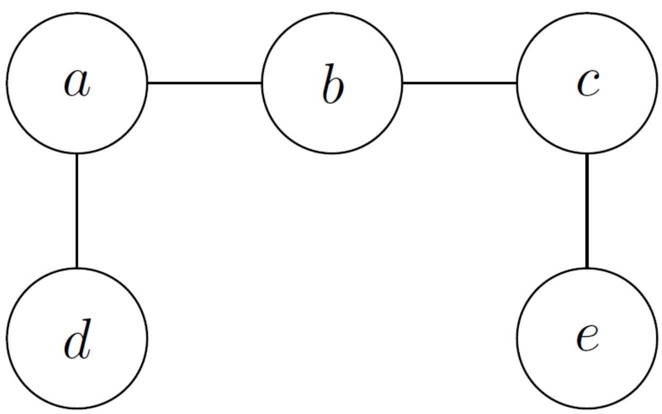

Figure 1: A Markov network

Consider the Markov network in Figure 1. Each variable is binary and can be in state 0 or state 1. d is observed to be in state 1 and e is observed to be in state 0. Additionally, the compatibility matrices Φ and Ψ are given by

$$\Phi(a, d) = \begin{pmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{pmatrix} = \Phi(c, e)$$

$$\Psi(a, b) = \begin{pmatrix} \alpha & 1-\alpha \\ 1-\alpha & \alpha \end{pmatrix} = \Psi(b, c)$$

(a) For $\alpha = 0.99$ find P(a), the marginal probabilities of variable a being in either of its two possible states, 0 and 1.

(b) Do the same for $\alpha = 0.6$. Discuss why the result is different in these two cases.

Problem 2 *Belief Propagation* Many vision problems consist of measuring local evidence, then propagating it across space. Belief propagation is often useful for such tasks. This homework problem was presented as a belief propagation example by Yair Weiss in a NIPS paper [1].

A task of early vision is to make a figure/ground assignment: which side of a contour is the foreground object, and which side is the background? A good cue for that assessment is convexity. Contours typically encircle the object, rather than form holes within it, so the foreground side is often on the inside of a contour's curve.

Locally, a complex contour may bend both ways and only a global assessment of convexity can tell us the right answer. We define a Markov chain of points along a contour in an image (we assume this contour has already been detected). The hidden states are the side of the foreground assignment for the contour (+1 means to the right as you traverse the contour, incrementing the node index; -1 is to the left). The local evidence at each node is based on the local curvature, defined by the angle θ_j based on the local three adjacent points on the curve (nodes $j-1$, j, and $j+1$). Let $\theta_j = 0$ correspond to a straight line, and $\theta_j = \frac{\pi}{2}$ correspond to a 90° right bend, and $\theta_j = -\frac{\pi}{2}$ correspond to a 90° left bend.

Let the local evidence for a positive or negative curvature curve be:

$$\phi(x_j, y_j) = \begin{pmatrix} \frac{\pi-\theta_j}{2\pi} \\ \frac{\pi+\theta_j}{2\pi} \end{pmatrix}$$

This favors figure/ground evidence in proportion to the acuteness of the local angle of bending.

The hidden state compatibility requires that hidden states have the same value as that of the neighboring node: $\psi\left(x_j, x_{j+1}\right) = I_2$ (i.e. the identity matrix).

The joint probability of figure/ground estimates conditioned on the observed curve is given by the product of the local evidence and the node capabilities:

$$P\left(\vec{x}|\vec{y}\right) = \prod_j \phi\left(x_j, y_j\right) \psi\left(x_j, x_{j+1}\right)$$

In the resources folder, you will find three images and their curves, which are mat files containing x and y coordinates of points along the boundary. In general we can use computer vision algorithms (some of which you have already encountered) for extracting such contours from images, however here we produced them manually using the supplied `DrawCurve` function.

(a) Load and plot the supplied curves in the 2D plane. Note that the coordinate system has its origin as the top left corner of the image.

(b) Indicate the local direction of figure (draw a small arrow to the foreground side at every node) based on local evidence alone, before running belief propagation.

(c) Show the final estimated direction of figure after running belief propagation.

References

[1] Yair Weiss. "Interpreting images by propagating Bayesian beliefs". In: Advances in Neural Information Processing Systems 9. `<http://www.cs.huji.ac.il/~yweiss/nips96.pdf>`. 1996, pp. 908-915.