Boundaries as Contours of Optimal Appearance and Area of Support

Christina Pavlopoulou and Stella X. Yu

Computer Science Department
Boston College

Acknowledgments:
NSF CAREER IIS-0644204
Clare Luce Boothe Professorship

Contour-specific Observations

Bayesian contour formulation:

$$P(C|O) = \frac{P(O|C) \cdot P(C)}{P(O)}$$

1. When can it be omitted?
2. Is it useful to include it?

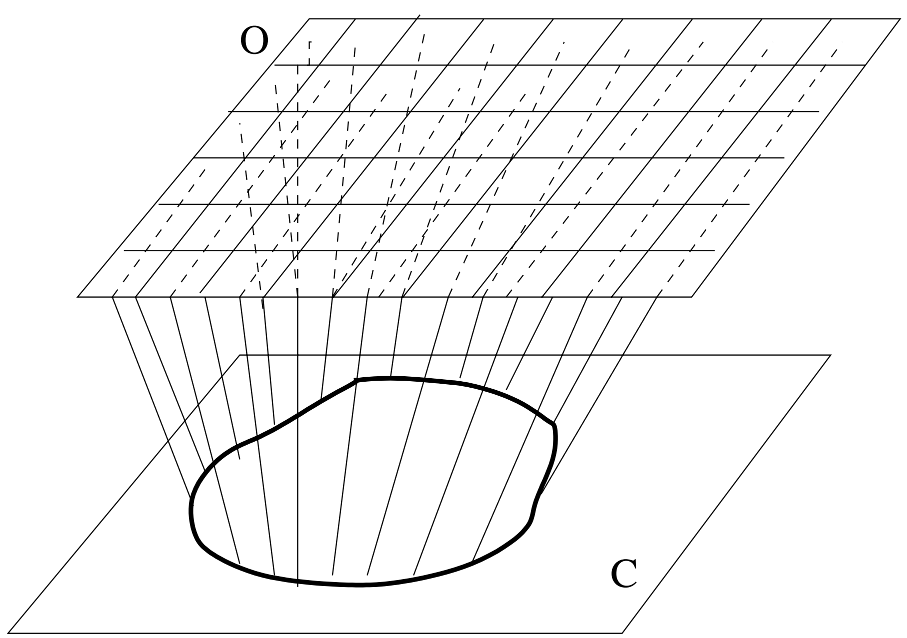

$O =$ entire image
$P(O)$ is irrelevant

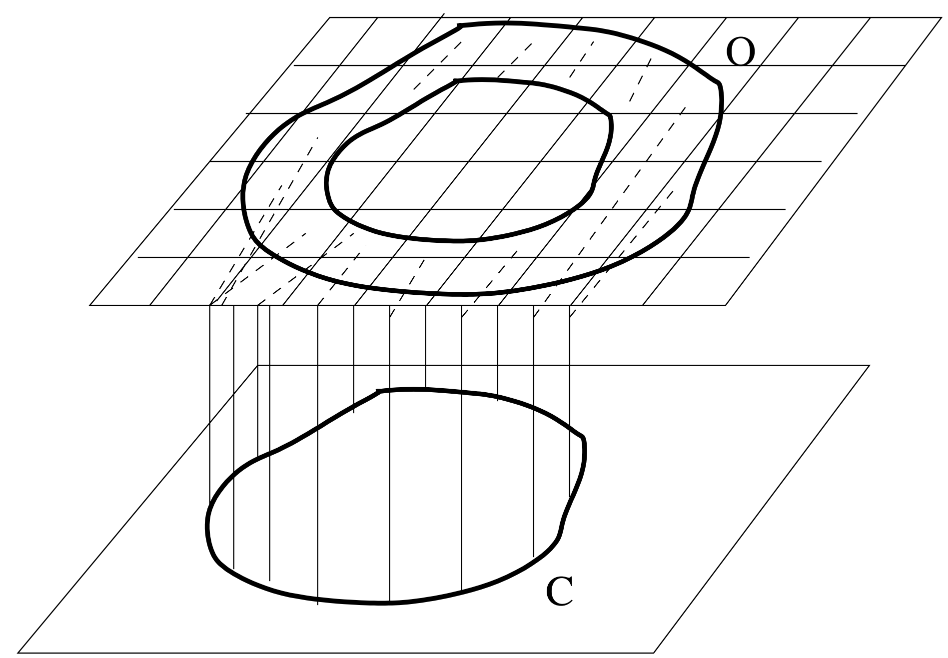

$O =$ depends on contour *C*
$P(O)$ can **not** be omitted

$P(O)$ as a Confidence Measure

$P(O)$ is computed as a marginal:

$$P(O) = \sum_{C} P(O|C) \cdot P(C)$$

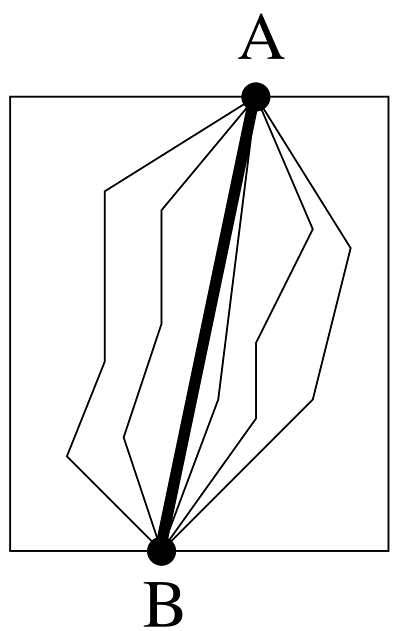

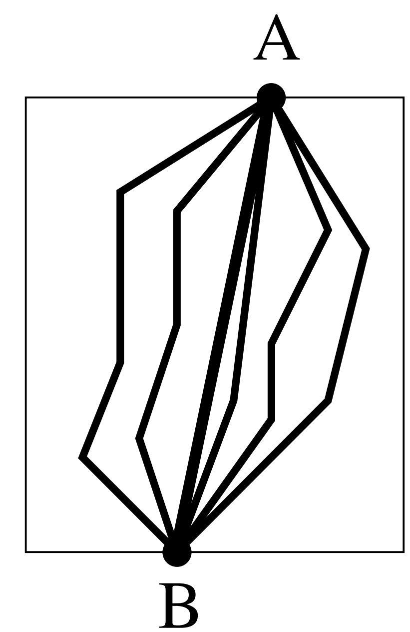

small $P(O) \rightarrow$ large $P(C|O)$

favor singularly best contour

large $P(O) \rightarrow$ small $P(C|O)$

discount numerous good contours

Entropy Interpretation

Let $P(O|C)P(C) = \beta_j$. Then:

$$\log P(C|O) = E(\beta_j) = \log \beta_j - \log \sum_i \beta_j$$

- maximum at 0, when there is a single good candidate
- minimum at $-\infty$, when there are many good candidates
- inverse entropy behavior

$$E(\beta_j) \geq \underbrace{\log \beta_j}_{\text{data term}} - \sum_j \underbrace{\beta_j \log \beta_j}_{\text{entropy of contour distribution}}$$

- naturally combines uncertainty measure with contour quality
- the area of support of the optimal contour has low uncertainty

Approaches to the Length Bias Problem

In traditional models, longer contours have higher energy. This creates a *strong bias towards short boundary segments*.

- Good initialization, user interaction (Terzopoulos87, Geiger95, Mortensen98, Udupa98)
- Heuristics during optimization (Fua97, Mortensen01, Mortensen04)
- Image features (Cohen91, Paragios02, Cohen02)
- Stronger contour priors (Blake01, Ziou07, Srivastava09)
- Mean ratio weight cycles (Ishikawa01, Cremers07) (Address directly the formulation of the objective criterion. Are applicable to closed contours and do not admit interaction.)
- *Our normalization term alleviates the bias problem.*

Graph-based Optimization

Calculation of $P(O)$

for $d_j, d_k \in D = \{0, \ldots, 7\}$

calculation of normalization is performed simultaneously with optimization

$$\alpha_i(d_k) = P(O_{c_1}, \ldots, O_{c_i}, q_i = d_k)$$

$$\alpha_1 = P(O_{c_1} | c_1 = d_k)$$

$$\alpha_{i+1}(d_k) = \left(\sum_{d_j} \alpha_i(d_j) P(c_{i+1} = d_j | c_i = d_k) \right) P(O_{c_{i+1}} | c_{i+1} = d_k)$$

$$P(O) = \sum_{d_j} \alpha_n(d_j)\,,\ n \text{ number of contour points}$$

Graph Structure

for connected neighboring pixels *u, v*

$$w(u,v) = \log P(O_u, O_v | u, v) + \log P(v|u) - \log \sum_{d_j} \alpha_v(d_j)$$

for points a priori known on the object boundary, Dijkstra's algorithm finds the global optimum very efficiently

Feature Calculations

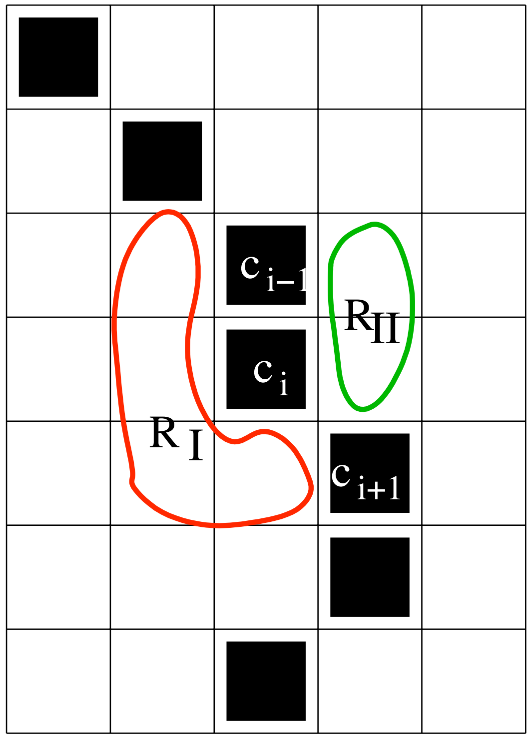

$$P(O_{\{c_i\}}|c_{\{i\}}) = \sum_{p \in R_I} M_I(p) + \sum_{p \in R_{II}} M_{II}(p)$$

$M_I(p), M_{II}(p)$ confidence of classifying p in I or II

$P(c_{i+1}|c_i, c_{i-1})$ penalizes angular configurations

Requires knowledge of interior and exterior of object.
Provided by the user in small parts of the image.
Non-parametric modeling is used for interior and exterior.

Results: Reduction in User Input

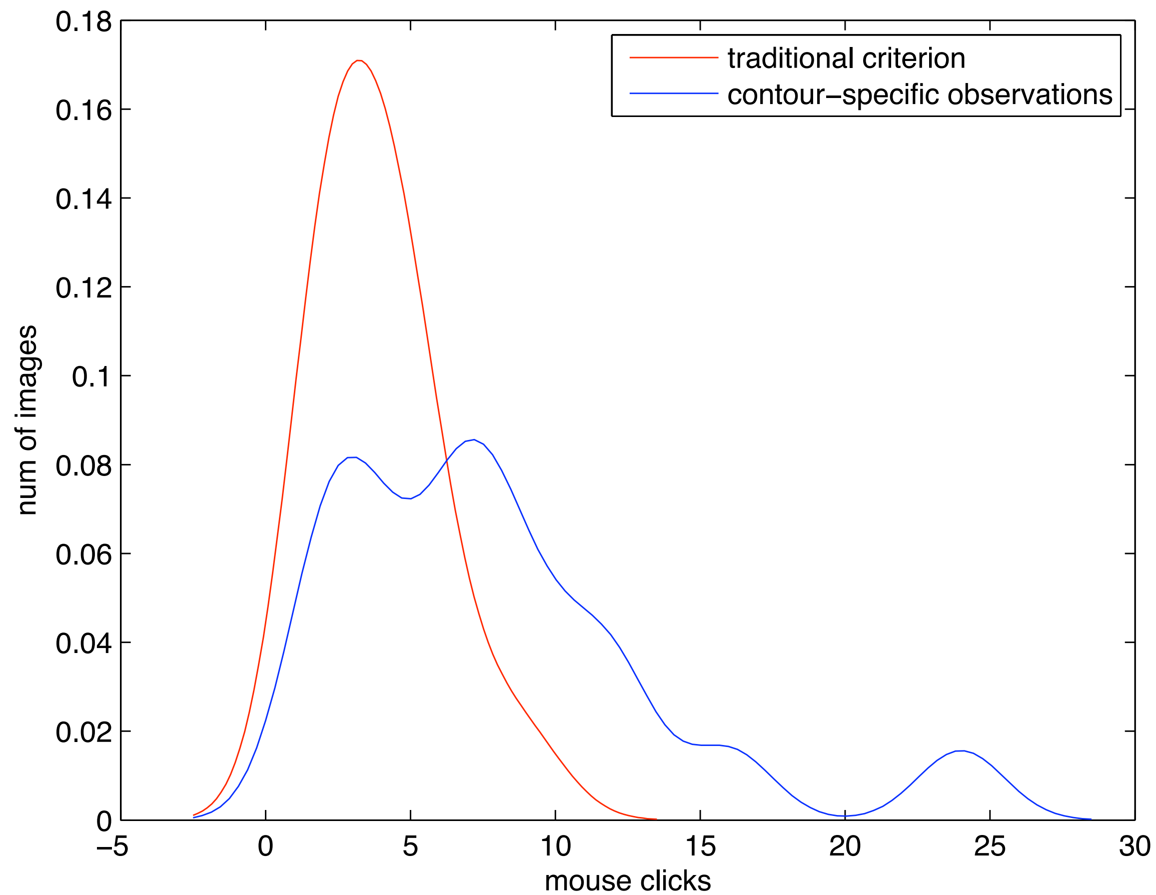

- Counted the amount of mouse clicks required to delineate a boundary in 17 images.
- Mouse clicks were inserted sequentially.
- Same training data for interior and exterior were used for both methods.
- Traditional criterion (without normalization) requires more user input than the proposed one because of the boundary length bias.

Geometrically Complex Objects

Because of the bias towards short segments, the traditional criterion favors straight lines and does not follow complex geometry even when the object is very distinct from its background.

Non-uniform Interior or Exterior

training data

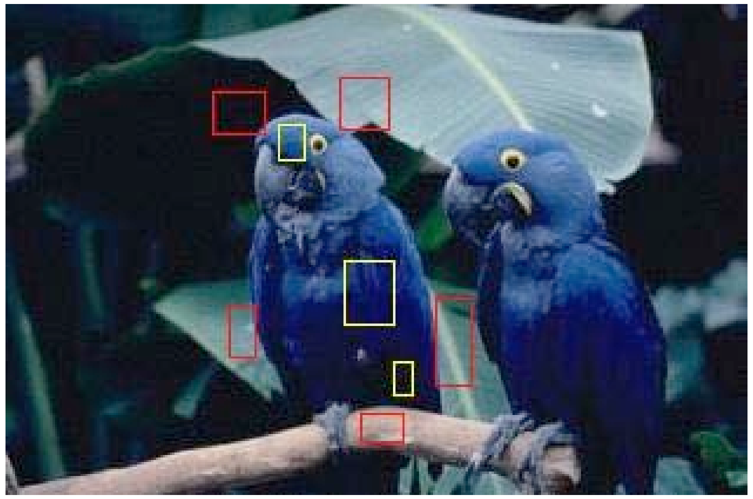

old method

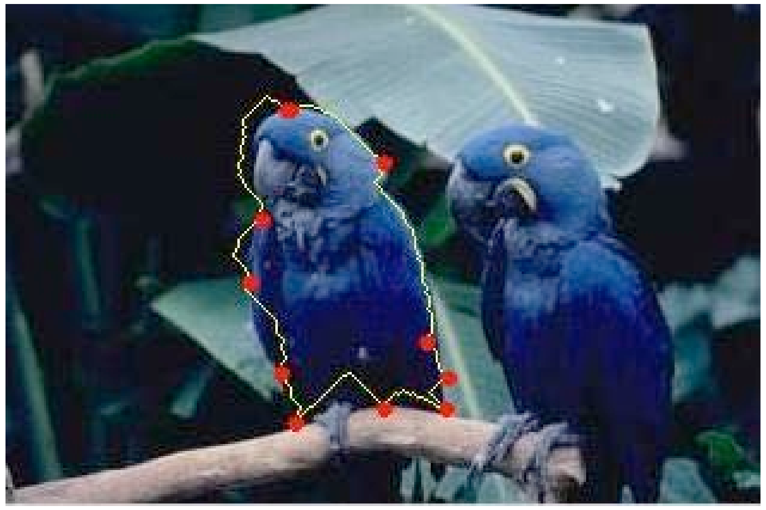

new method

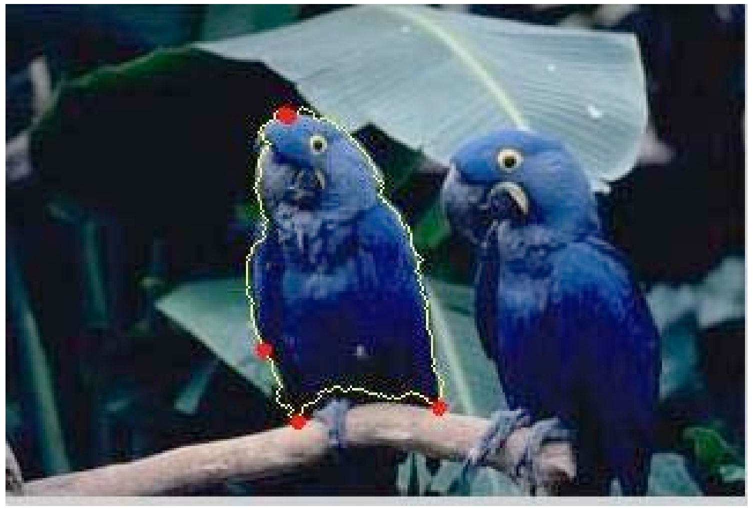

Proposed criterion snaps more accurately to the boundary than traditional one for the same training data.

Texture

training data

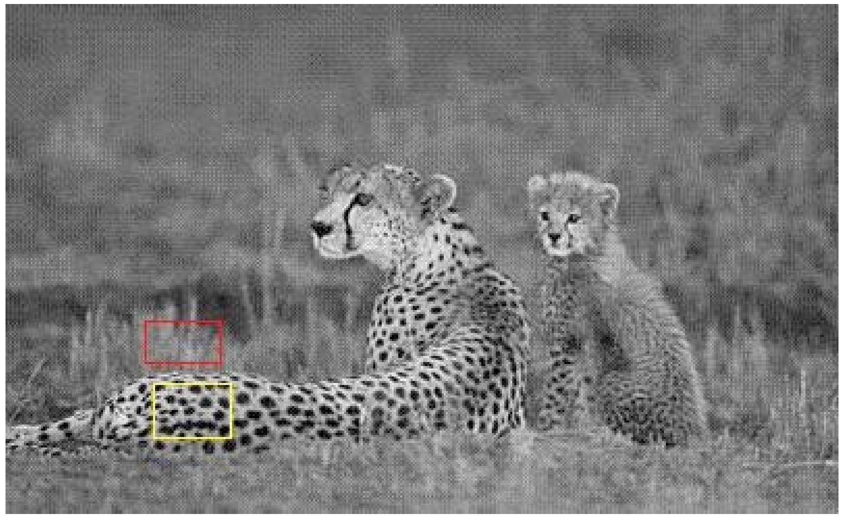

old method

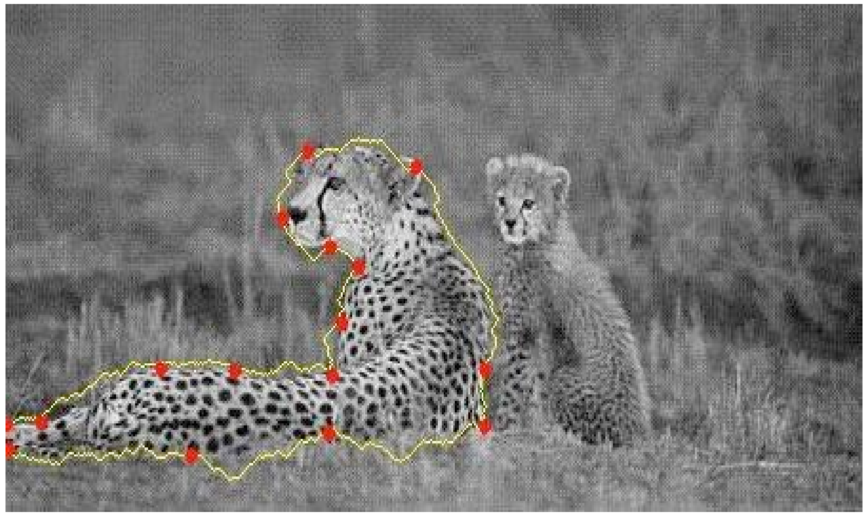

classification results

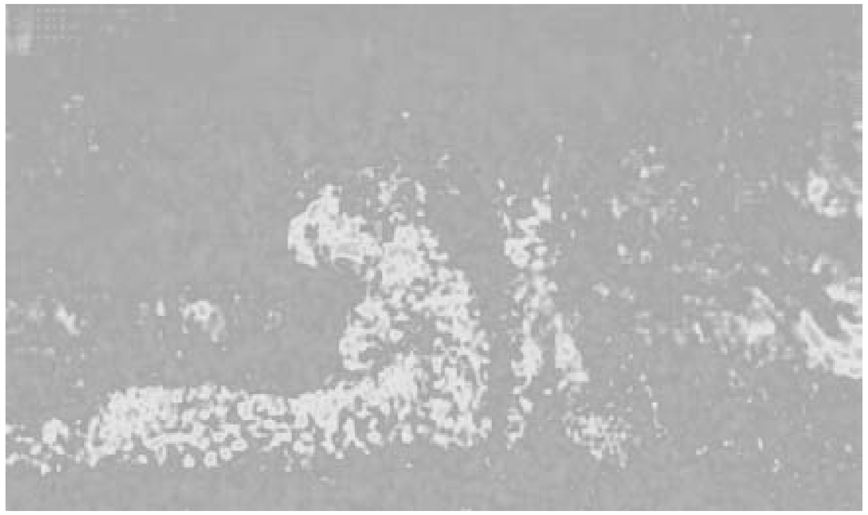

new method

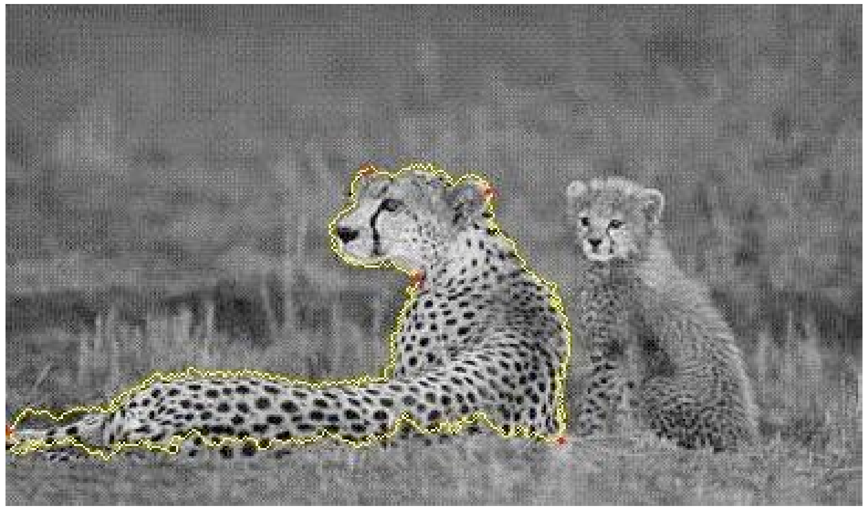

Proposed approach extracts the boundary of textured objects despite the lack of explicit texture modeling.

Conclusions

- Allowing for contour-specific observations requires inclusion of the normalization P(O) in the optimization.
- P(O) acts as a confidence measure and favors contours distinctly better than alternatives in their vicinity.
- The bias towards short boundary segments is alleviated.
- Global optimum can be found in low-order polynomial time.