A Continuous Shape Prior for MRF-based Segmentation

Dmitrij Schlesinger

Dresden University of Technology

Continuous

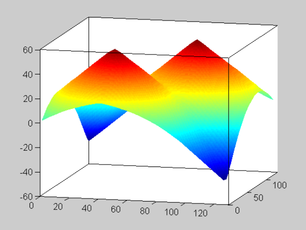

+ Low-level properties – length, curvature etc.
+ Shape modeling
– Statistical interpretation, learning

Discrete

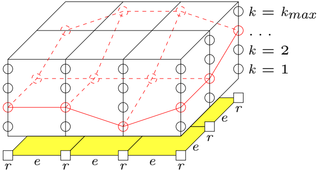

– Metrication artifacts
+ Spatial relations
+ Multi-label problems
– Locality
+ Sound statistical interpretation

Try to combine ...

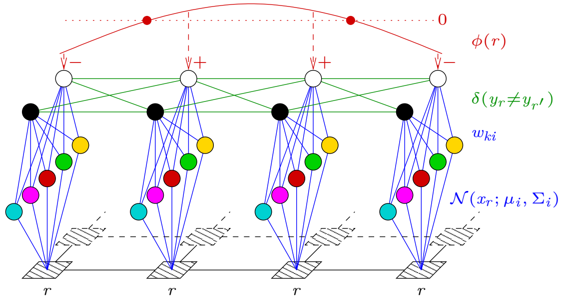

x – observation, y – labeling, ϕ – shape function

$$p(x, y; \phi) \propto \exp\left[\alpha \cdot \sum_{rr'} \delta(y_r \neq y_{r'}) + \sum_r y_r \cdot \phi(r)\right] \times \prod_r p(x_r|y_r)$$

The prior consists of two parts – Potts terms and shape terms.
Conditionally independent observation model (Gaussian mixtures).

TECHNISCHE UNIVERSITÄT DRESDEN

Segmentation:

Let ϕ be known. Bayesian Decision Task with Hamming Loss
$\rightarrow$ Maximum Marginal Decision:

$$y_r^* = \arg\max_k p(y_r{=}k|x;\phi)$$

Gibbs Sampling for approximate estimation

Appearances:

Standard EM-scheme:

$$\sum_r p(y_r{=}k|x;\phi) \cdot \ln \sum_i w_{ik} \mathcal{N}(x_r;\mu_k,\Sigma_k) \rightarrow \max_{w_{ik},\mu_k,\Sigma_k}$$

Posterior marginal label probabilities are again necessary

Inference or learning ?

Regularized Likelihood:

$$L(\phi) + \mathcal{R}(\phi) =$$

$$= \ln \sum_{y} p(x, y; \phi) - \int_{\Omega} \left(\lambda_1 \cdot \|\nabla\phi\|^2 + \lambda_2 \cdot \triangle\phi^2\right) d\omega \to \max_{\phi}$$

The gradient with respect to ϕ:

$$\frac{\partial}{\partial\phi(r)} = p(y_r{=}1|x;\phi) - p(y_r{=}1;\phi) + \frac{\partial\mathcal{R}(\phi)}{\partial\phi(r)}$$

- Posterior marginal label probabilities (Gibbs Sampling)
- Prior marginal label probabilities (an approximation)
- Gâteaux derivative (convolution)

TECHNISCHE UNIVERSITÄT DRESDEN

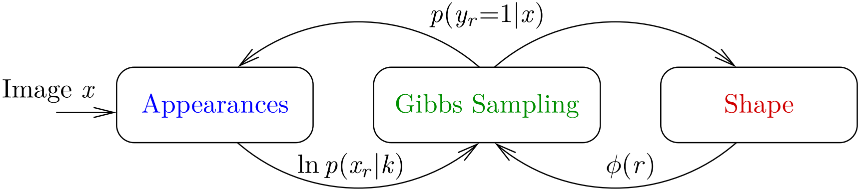

Three blocks working in parallel:

- Gibbs Sampling: computes marginal posterior label probabilities
- Appearances: performs the learning of the appearances via EM
- Shape: estimates the shape function via gradient ascent

Experiments

Weizmann Horse Database, 328 images

Good (< 1%) ← Satisfactory (8 – 9%) → Bad (>25%)

Original images

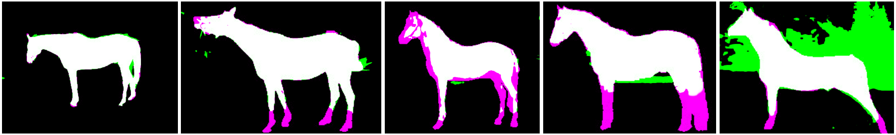

Segmentations (green – false positive, magenta – false negative)

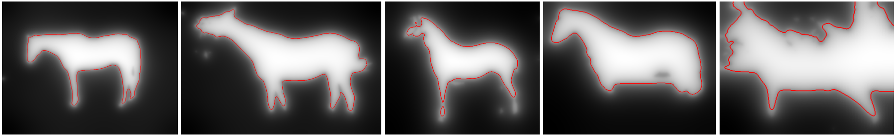

Estimated shape functions (zero-level in red)

Approach	Accuracy (%)	
ObjCut	96.0	Complex methods
Levin	95.5	
Borenstein	93.6	
Zhang	95.4	
Bertelli	94.3	
Domke	92.0	
Our	92.2	
GrabCut	85.5	Simple methods
Co-segmentation	80.1	
MNcut	51.0	

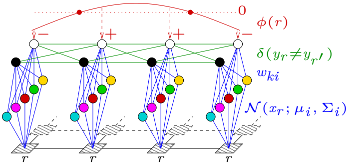

- $+$ Low-level properties – length, curvature etc.
- $+$ Shape modeling (possible)
- $+$ Multi-label problems (one shape function per label)
- $+$ Spatial relations (possible, easy)
- $+$ Statistic interpretation (for the discrete part)

- – Quite slow (actually, under circumstances)
- – Non-convex optimization, initialization is necessary
- – Approximate inference

Outlook

Compare

$$\text{“usual” cont. optimization:} \quad \int_\Omega \big[\mathrm{Data}(x,\phi) + \mathrm{Regularizer}(\phi)\big]\, d\omega$$

$$\text{ours:} \quad L(x;\phi) + \int_\Omega \mathrm{Regularizer}(\phi)\, d\omega$$

In both cases the gradient looks like

$$\frac{\partial E(\phi)}{\partial \phi(r)} = \mathrm{Something}(\phi(r)) + \frac{\partial\ \mathrm{Regularizer}(\phi)}{\partial \phi(r)}$$

$\Rightarrow$ other regulizers can be easily adopted
$\Rightarrow$ advanced methods from continuous optimization can be used

---

Multi-label problems $\rightarrow$ part-based shape segmentation