Part Discovery from Partial Correspondence

Subhransu Maji Gregory Shakhnarovich
Toyota Technological Institute at Chicago

Project goals

Semantic part-based models of categories

We propose a novel *correspondence* driven *annotation* and *learning* framework for part discovery that overcomes some of the drawbacks of existing techniques

Current techniques for annotation

Obtaining annotations can be hard

anatomical landmarks

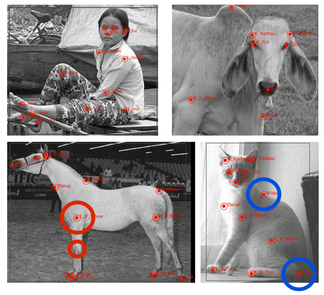

Names can be *misleading*
Where is the *elbow* of a horse?
Hard to localize

diverse categories

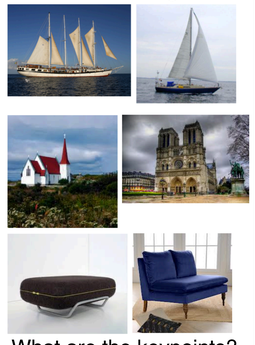

What are the keypoints?
Can you name them?

Proposed annotation framework

Mark semantic correspondence

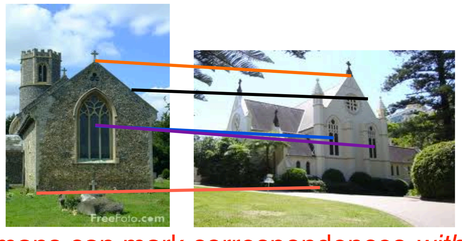

Humans can mark correspondences *without* knowing the names of the parts [HCOMP 12]

Annotators are shown pairs of images, examples of landmarks and GUI instructions

Example annotations collected

Annotations induce a semantic graph

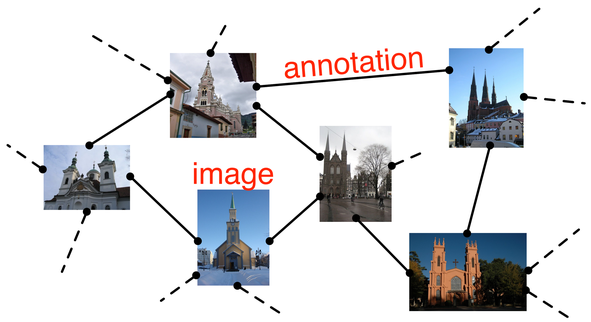

Exploring the semantic graph

Partial correspondence between a pair

Obtain patch correspondences

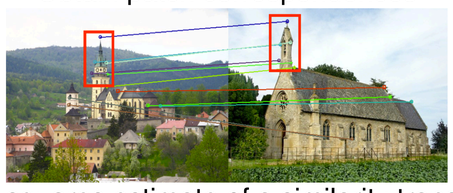

least squares estimate of a similarity transform

Depth-first exploration of the graph

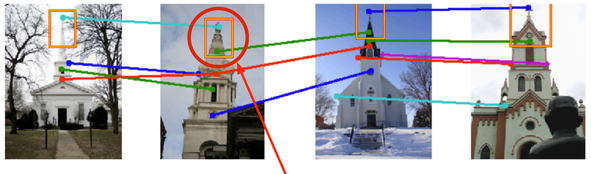

Automatically corrects for annotation bias

Breadth-first exploration of the graph

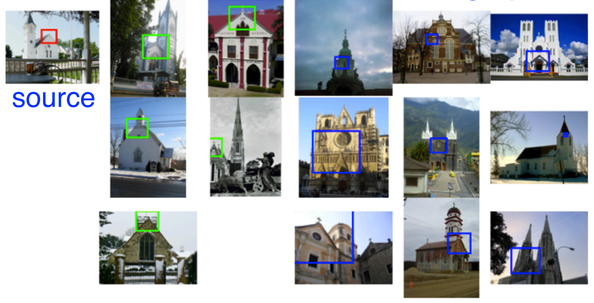

Can find a match as long as there is a path from the source in the semantic graph

Learning part detectors

Sample seed windows
Find similar patches using the semantic graph
Learn a robust appearance model

Where to sample seed windows?

clicked landmarks saliency map

Sample uniformly on the seeds
Reflects the underlying frequency of parts

Learning an appearance model

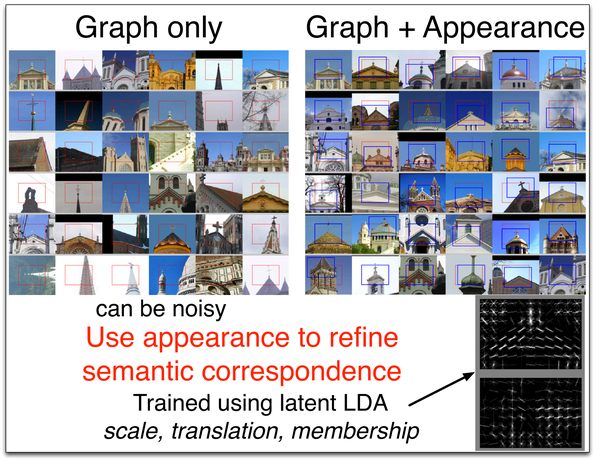

can be noisy

Use appearance to refine semantic correspondence

Trained using latent LDA
scale, translation, membership

Examples of discovered parts

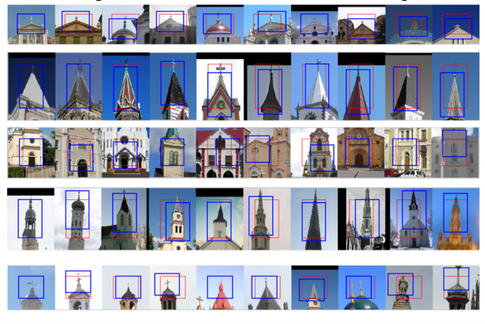

Evaluating parts

Dataset: 288 images of churches collected from **Flickr**, 1000 pairs of correspondence

Object localization

Hough voting based detector

Evaluating individual parts

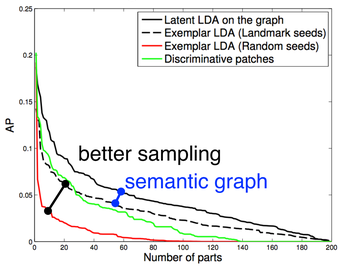

Example detections

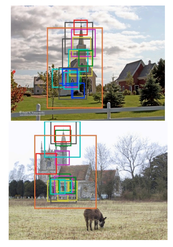

Each part is shown with a different color

Combining parts AP=**39.90%**; DPM AP=**34.75%**

Predicting landmark saliency

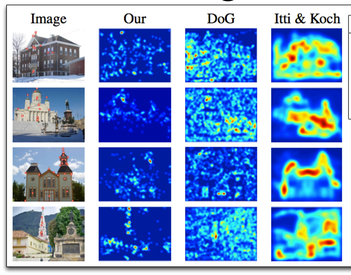

Method	MAL
Difference of Gaussian	1.23
Itti and Koch	1.86
Discriminative patches	6.14
Exemplar LDA (Landmark seeds)	5.79
Latent LDA on the graph	**7.84**

Task: Predict the locations of the clicked landmarks

Evaluated as the *Mean Average Likelihood* (MAL) of the clicked landmarks

Uniform saliency has MAL=1

Visualizing parts on images

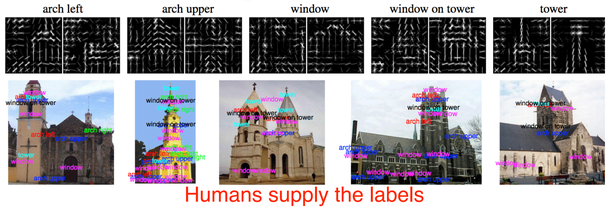

Humans supply the labels

References

[HCOMP 12] *Part annotation via pairwise correspondence*, Subhransu Maji and Gregory Shakhnarovich, AAAI Human Computation Workshop, 2012
[Discriminative Patches] *Unsupervised discovery of mid-level discriminative patches*, S.Singh, A. Gupta, and A. Efros, ECCV 2012
[DPM] *Object detection with discriminatively trained part-based models*, P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan, IEEE TPAMI, 2010
[Itti & Koch] *Computational modeling of visual attention*, L. Itti, and C. Koch, Nature reviews neuroscience, 2(3), 2011