

A Bayesian Approach to Unsupervised One-Shot Learning of Object Categories

Li Fei-Fei Rob Fergus Pietro Perona

Presenter:

Tuyen Huynh – Duy Vu

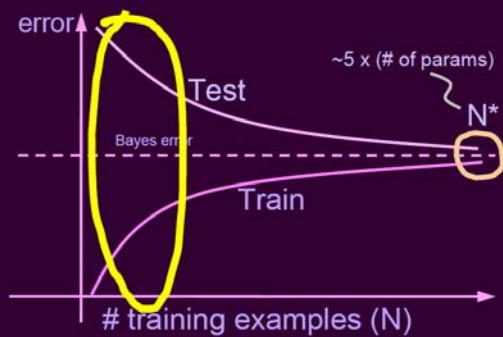
Slides from Li Fei-Fei http://www.vision.caltech.edu/feifeili/Fei-Fei_KUL2003.pdf

Goal

One-Shot learning of object categories.

Algorithm	Training Examples	Categories
Rowley et al.	~500	Faces
Schneiderman, et al.	~2,000	Faces, Cars
Viola et al.	~10,000	Faces
Burl, et al. Weber, et al. Fergus, et al.	200 ~ 400	Faces, Motorbikes, Spotted cats, Airplanes, Cars

Words of wisdom from statisticians




Summary of main points...


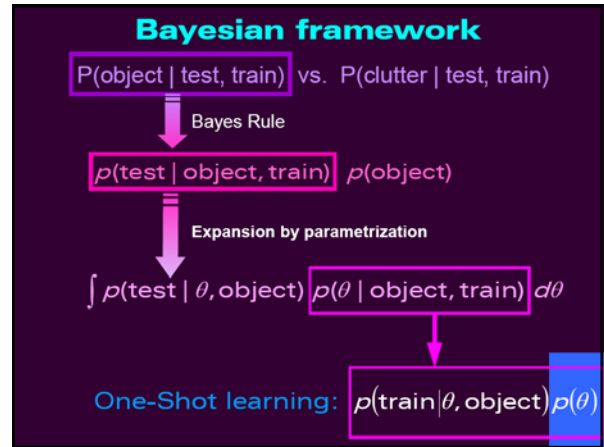
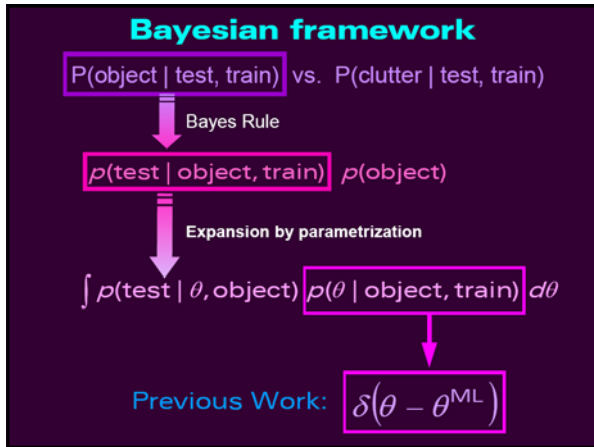
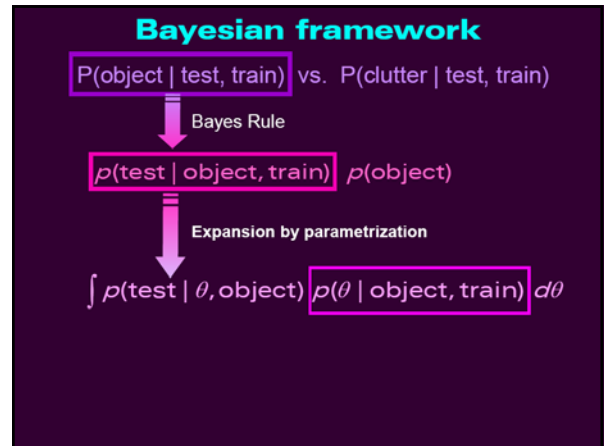
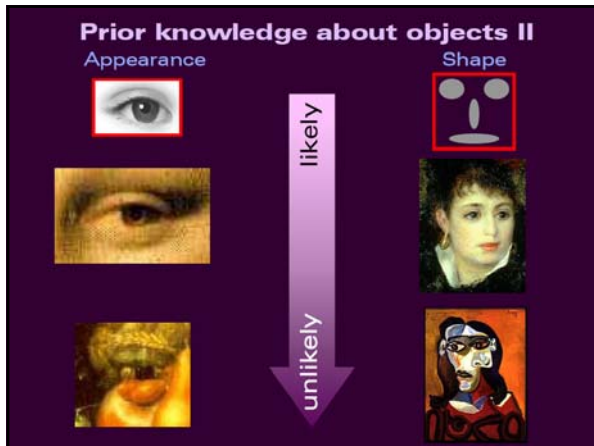
- Goal: one-shot learning
- Intuition: use general prior knowledge
- Approach: full Bayesian
 - Marginalize over all θ
- Implementation:
 - Probability models
 - Variational EM

Prior knowledge about objects I

Appearance

Shape



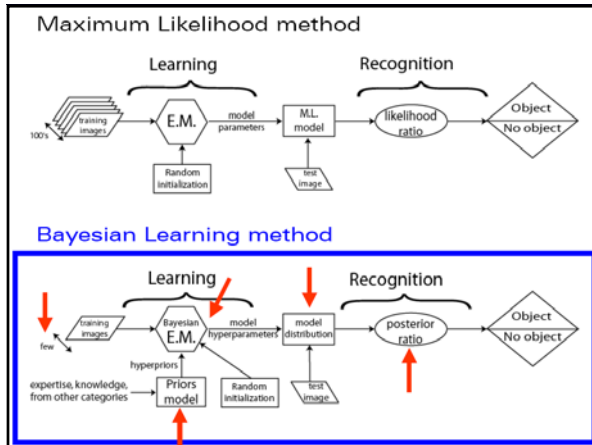
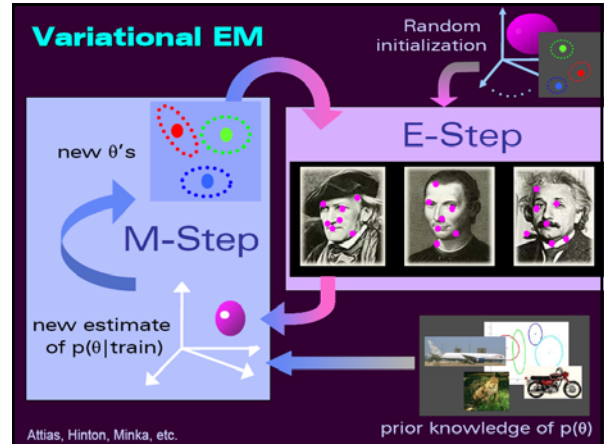
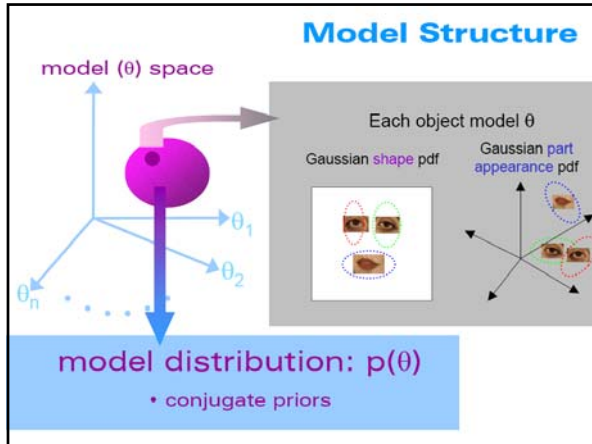


Representing θ

Main issues:

- measuring the similarity of parts
- representing the configuration of parts

- † Fischler & Elschlager 1973
- † Yuille '91
- † Brunelli & Poggio '93
- † Lades, v.d. Malsburg et al. '93
- † Cootes, Lanitis, Taylor et al. '95
- † Amit & Geman '95, '99
- † Perona et al. '95, '98, '98, '00, '03
- † Agarwal & Roth '02



Experiments

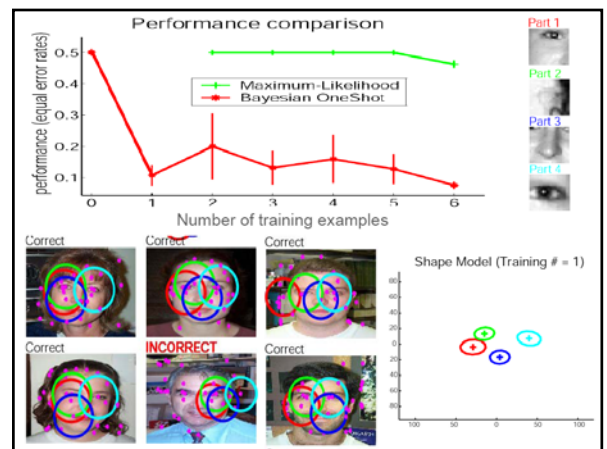
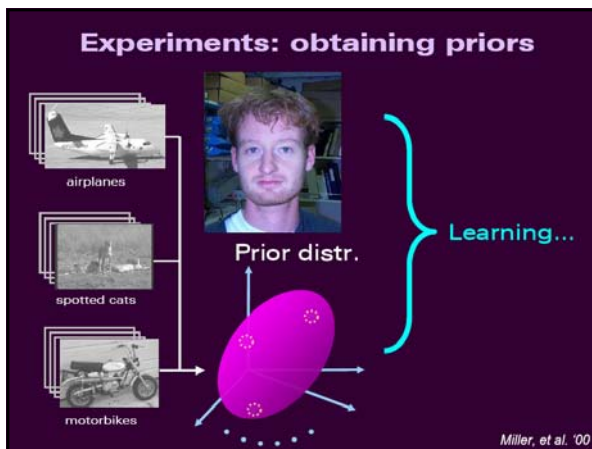
Training: 1- 6 images (randomly drawn)

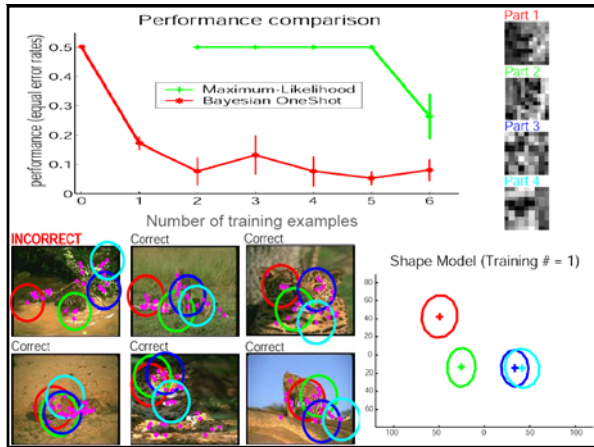
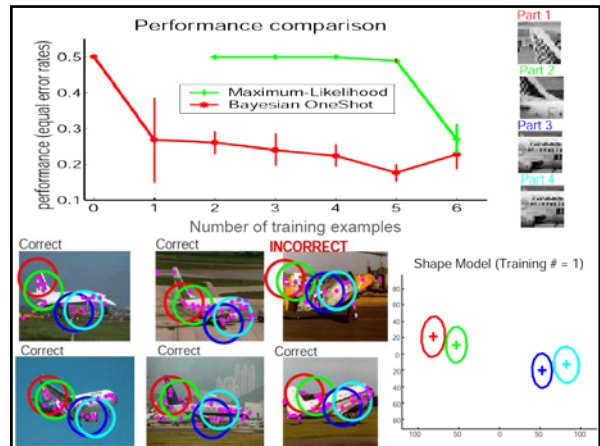
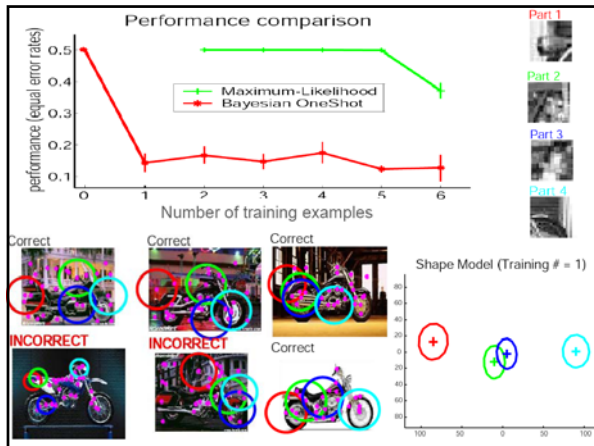
Testing: 50 fg/ 50 bg images (object present/absent)

Datasets

The datasets include: **faces**, **airplanes**, **spotted cats**, and **motorbikes**. Other images show various objects like a car, a building, and a person.

[www.vision.caltech.edu]





Algorithm	Training Examples	Categories	Results (error)
Burl, et al. Weber, et al. Fergus, et al.	200 ~ 400	Faces, Motorbikes, Spotted cats, Airplanes, Cars	5.6 - 10 %
Bayesian One-Shot	1 ~ 5	Faces, Motorbikes, Spotted cats, Airplanes	8 - 15 %

Summary

- Learning categories with one example is possible
- Can reduce # of training example from ~300 to 1~5 using prior from other categories

References

- L. Fei-Fei, R. Fergus, and P. Perona. A Bayesian approach to unsupervised One-Shot learning of Object categories. *IEEE Inter. Conf. Computer Vision*, 2003.
- L. Fei-Fei, R. Fergus and P. Perona. One-Shot learning of object categories. *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol28(4), 594 - 611, 2006
- Video of Fei-Fei's talk at ICCV 2003:
<http://www.robots.ox.ac.uk/~awf/iccvt03/videos/ICCV03-S11-P2-DIVX-MP3.avi>
- Variational Bayes: <http://www.variational-bayes.org/>