CVPR 2010 Tutorial
Semi-Supervised Learning in Vision

A, Saffari, Ch. Leistner, H. Bischof

Inst. for Computer Graphics and Vision
Graz University of Technology

http://www.icg.tugraz.at/Members/Saffari/ssl-cvpr2010

Typical Vision Tasks/Trends

Internet: Image Search/Classification Categorization

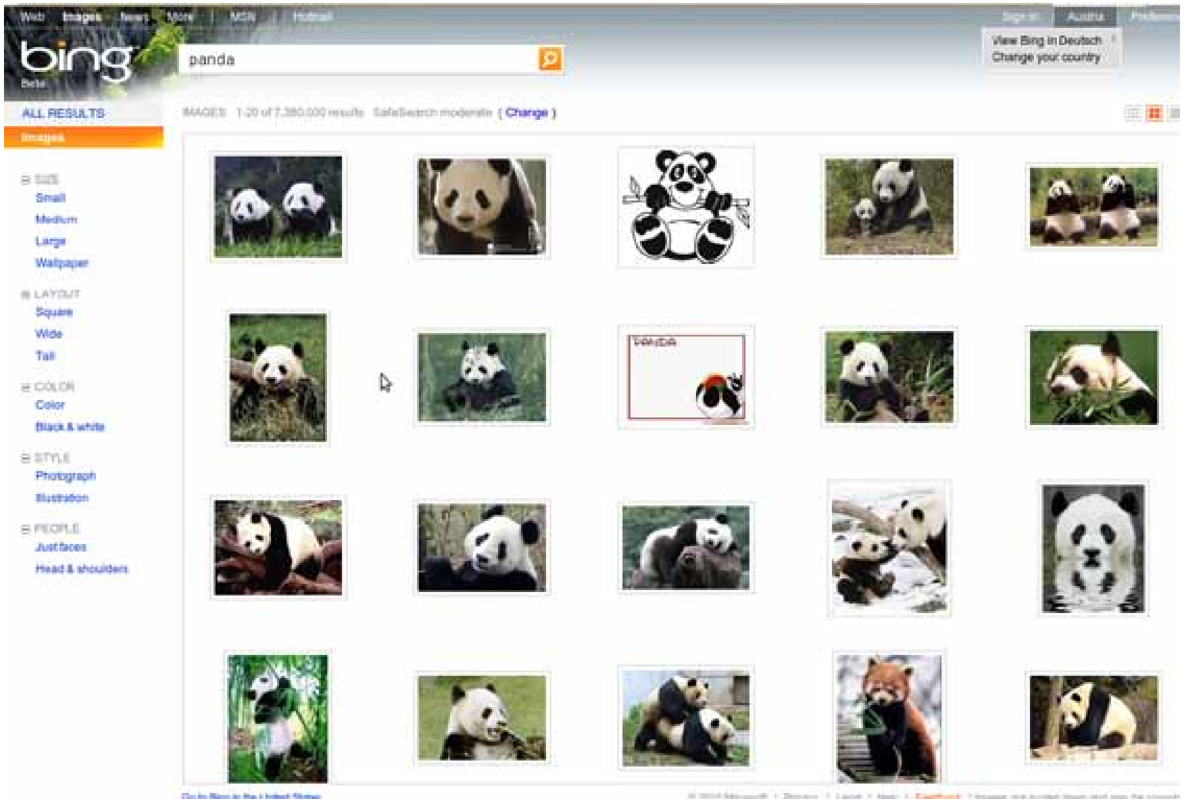

Huge Amounts of data, Huge labeling effort

Typical Vision Tasks/Trends

Internet: Various Video and Image Databases

Picasa

flickr

ImageShack

YouTube
Broadcast Yourself

facebook

myspace
a place for friends

Huge Amounts of data, Partially/weakly labeled data

Typical Vision Tasks/Trends

Object Categorization

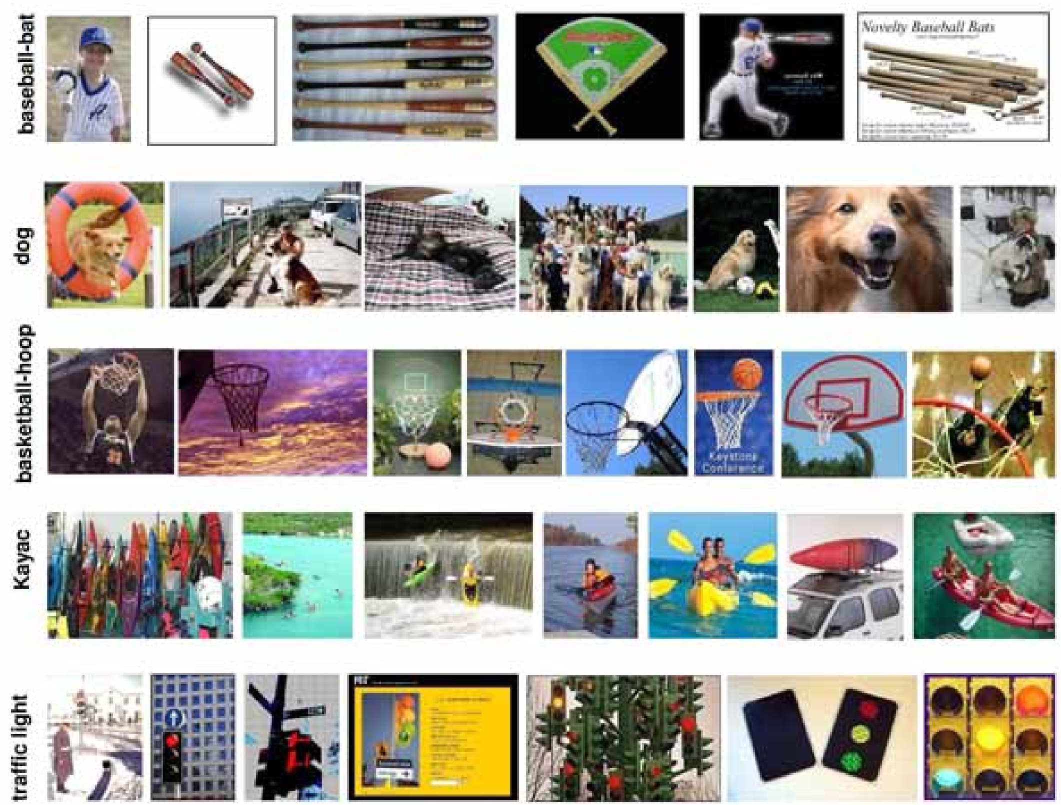

Huge Amounts of data, Huge labeling effort

Perona 2009

Visipedia

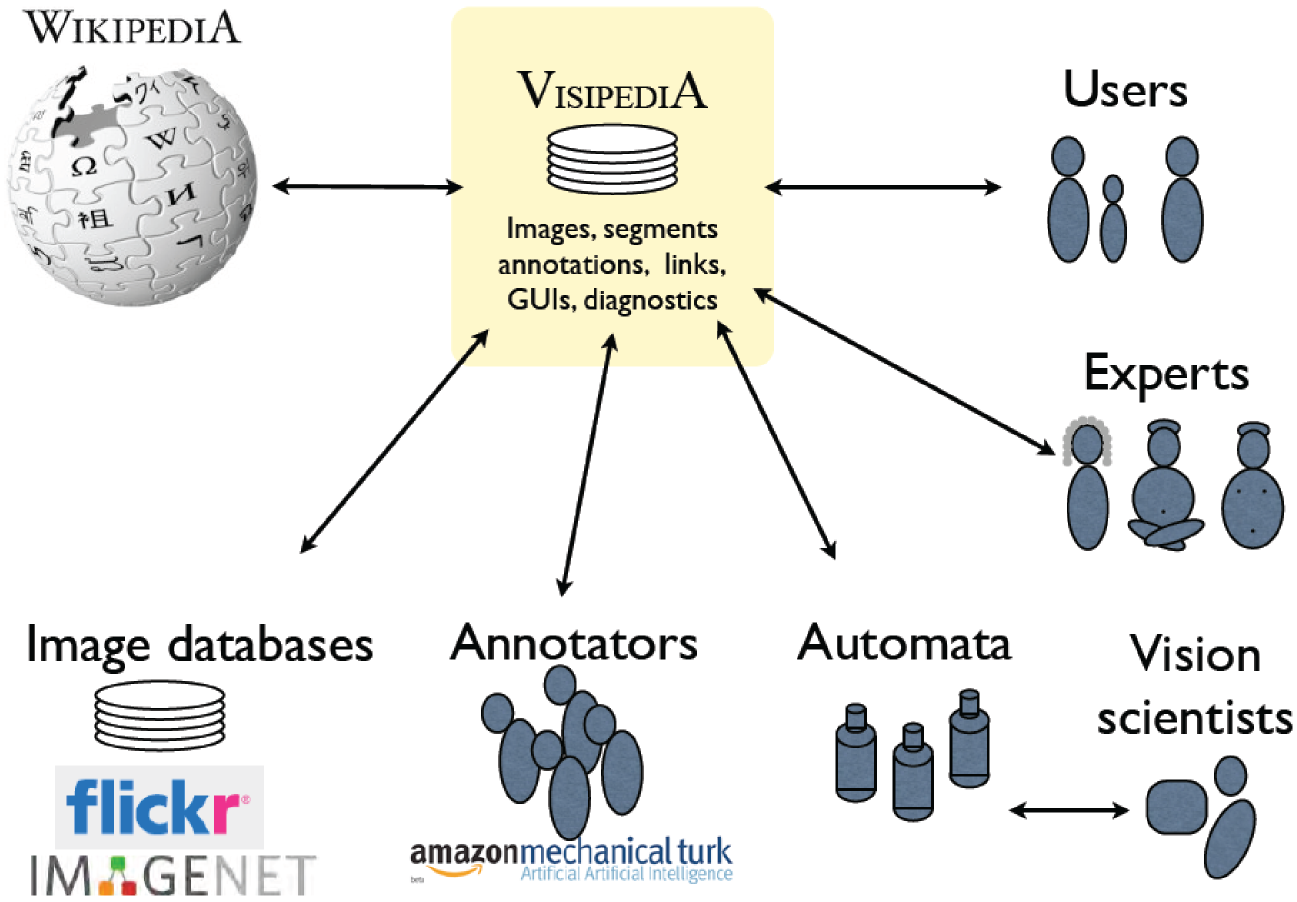

Typical Vision Tasks/Trends

Surveillance: On-line data/Detection-Tracking

Huge Amounts of data, On-line processing, Scene adaptation

Typical Vision Tasks/Trends

Many other tasks:

- Tracking: On-line adaptation to object
- Interactive segmentation: Changing model on the fly
- Interactive labeling: Suggestions as you label
- …..

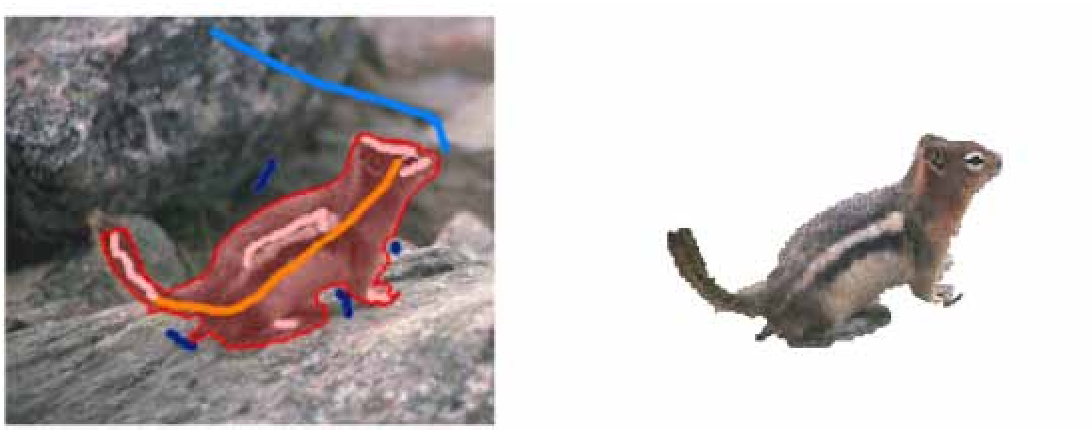

Requirements

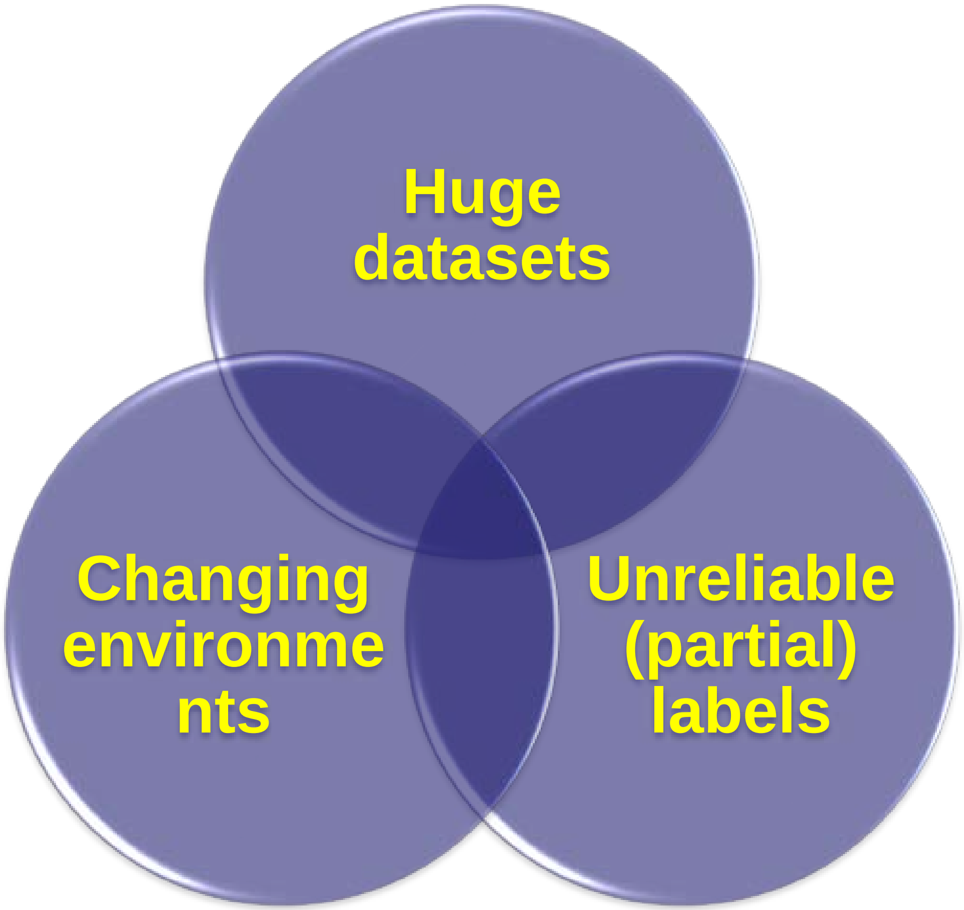

Requirements on Learner

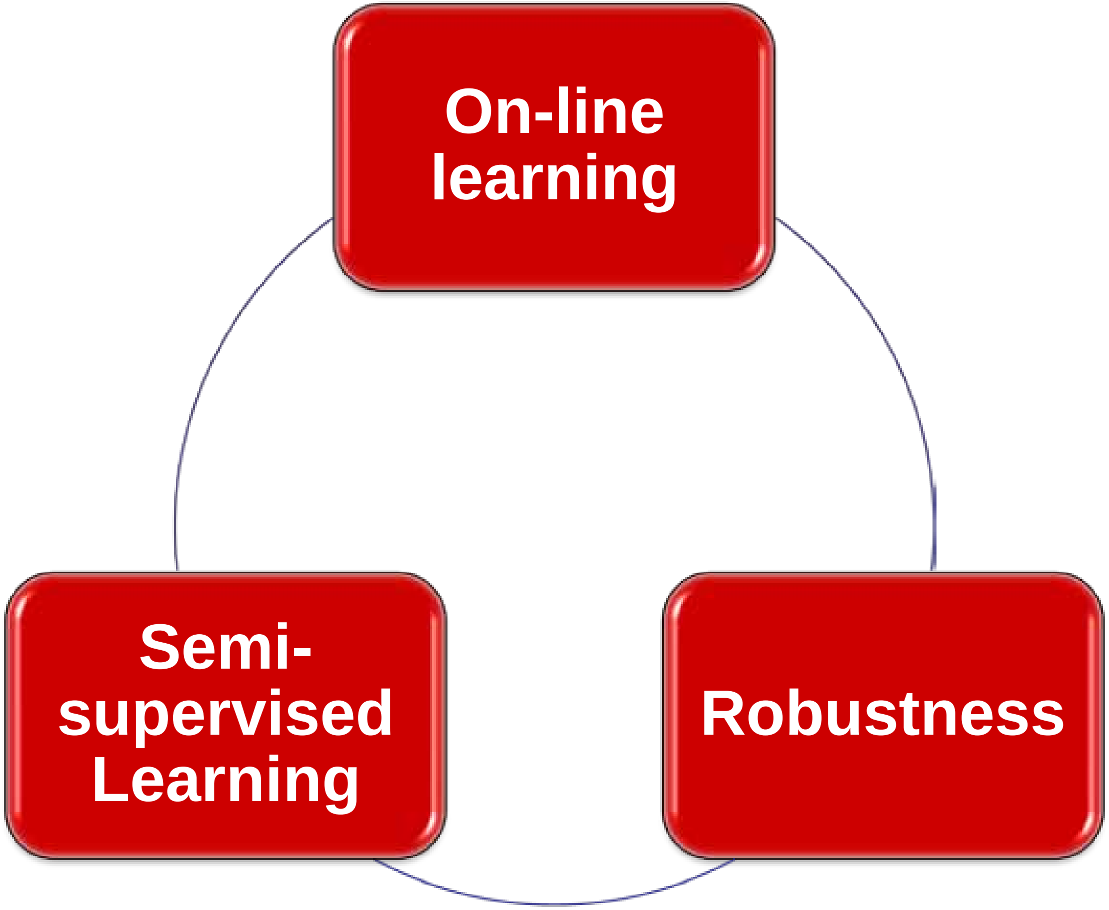

Theory/Methods Covered

1. Semi-Supervised Learning
2. Self-Training
3. Generative Models
4. Margin Assumption
5. Cluster and Manifold Assumption
6. Multi-View Learning
7. Large-Scale, Multi-Class SSL and Online Learning
8. Transfer Learning, Domain Adaptation, and Weakly Related Data
9. Multiple-Instance Learning

Applications Covered

1. Object Detection
2. Categorization
3. Tracking
4. Activity Recognition
5. Segmentation

Using:

Boosting

Random Forests

Learning Tasks

- Unsupervised learning
  - Density estimation, Clustering, Dimensionality reduction.
- Semi-Supervised Clustering
  - Clustering with pair-wise constraints: must-link, cannot-link.
- Semi-Supervised Classification and Regression
- Supervised Learning
  - Classification, Regression.

Labels

How do we get labels?

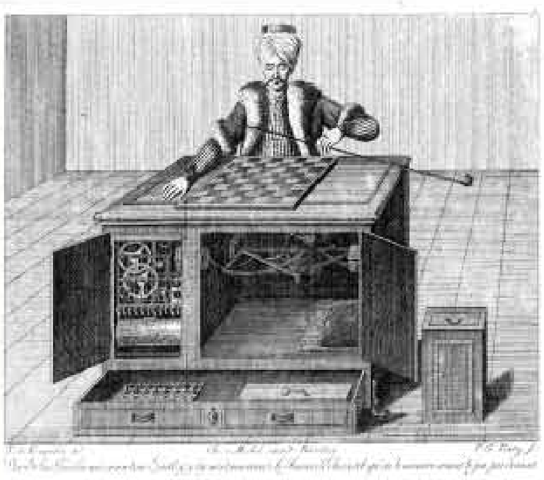

Semi-Supervised Learning

SSL is Supervised Learning...

Goal: Estimate P(y|**x**) from Labeled Data
$D_l=\{ (\mathbf{x}_i,y_i) \}$

$$p(y \mid x) = \frac{p(x \mid y)\, p(y)}{P(x)}$$

But: Additional Source tells about P(**x**)
(e.g., Unlabeled Data $D_u=\{\mathbf{x}_j\}$)

The Interesting Case:

$$|D_l| = n \ll m = |D_u|$$

Supervised learning

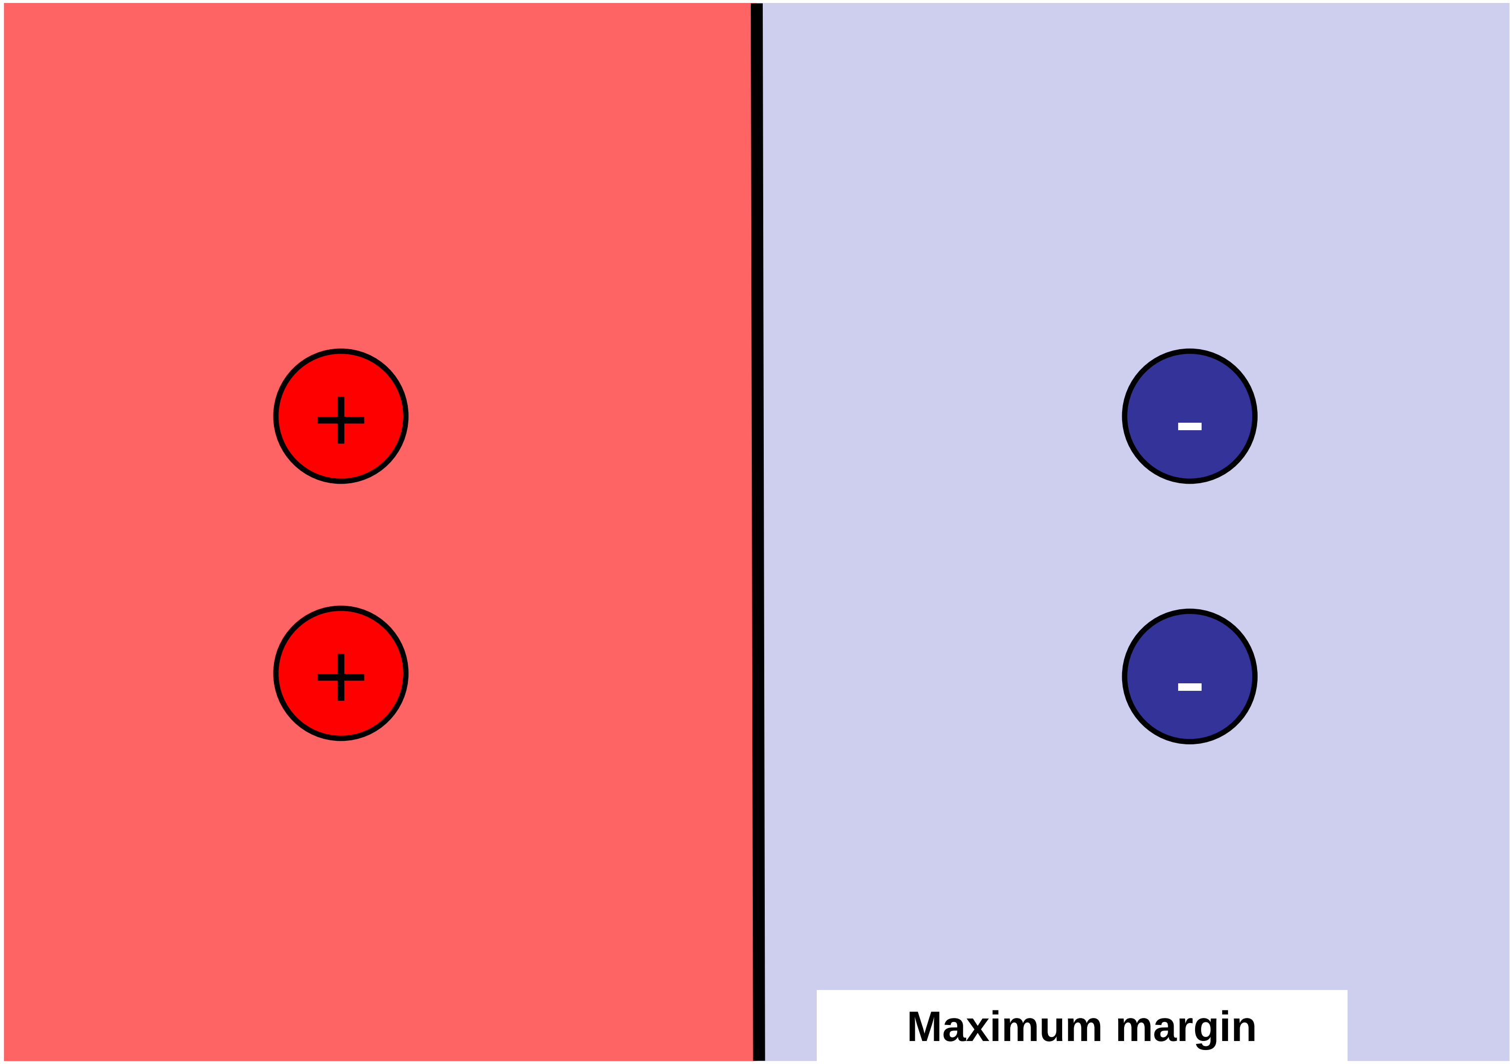

Can Unlabeled Data Help?

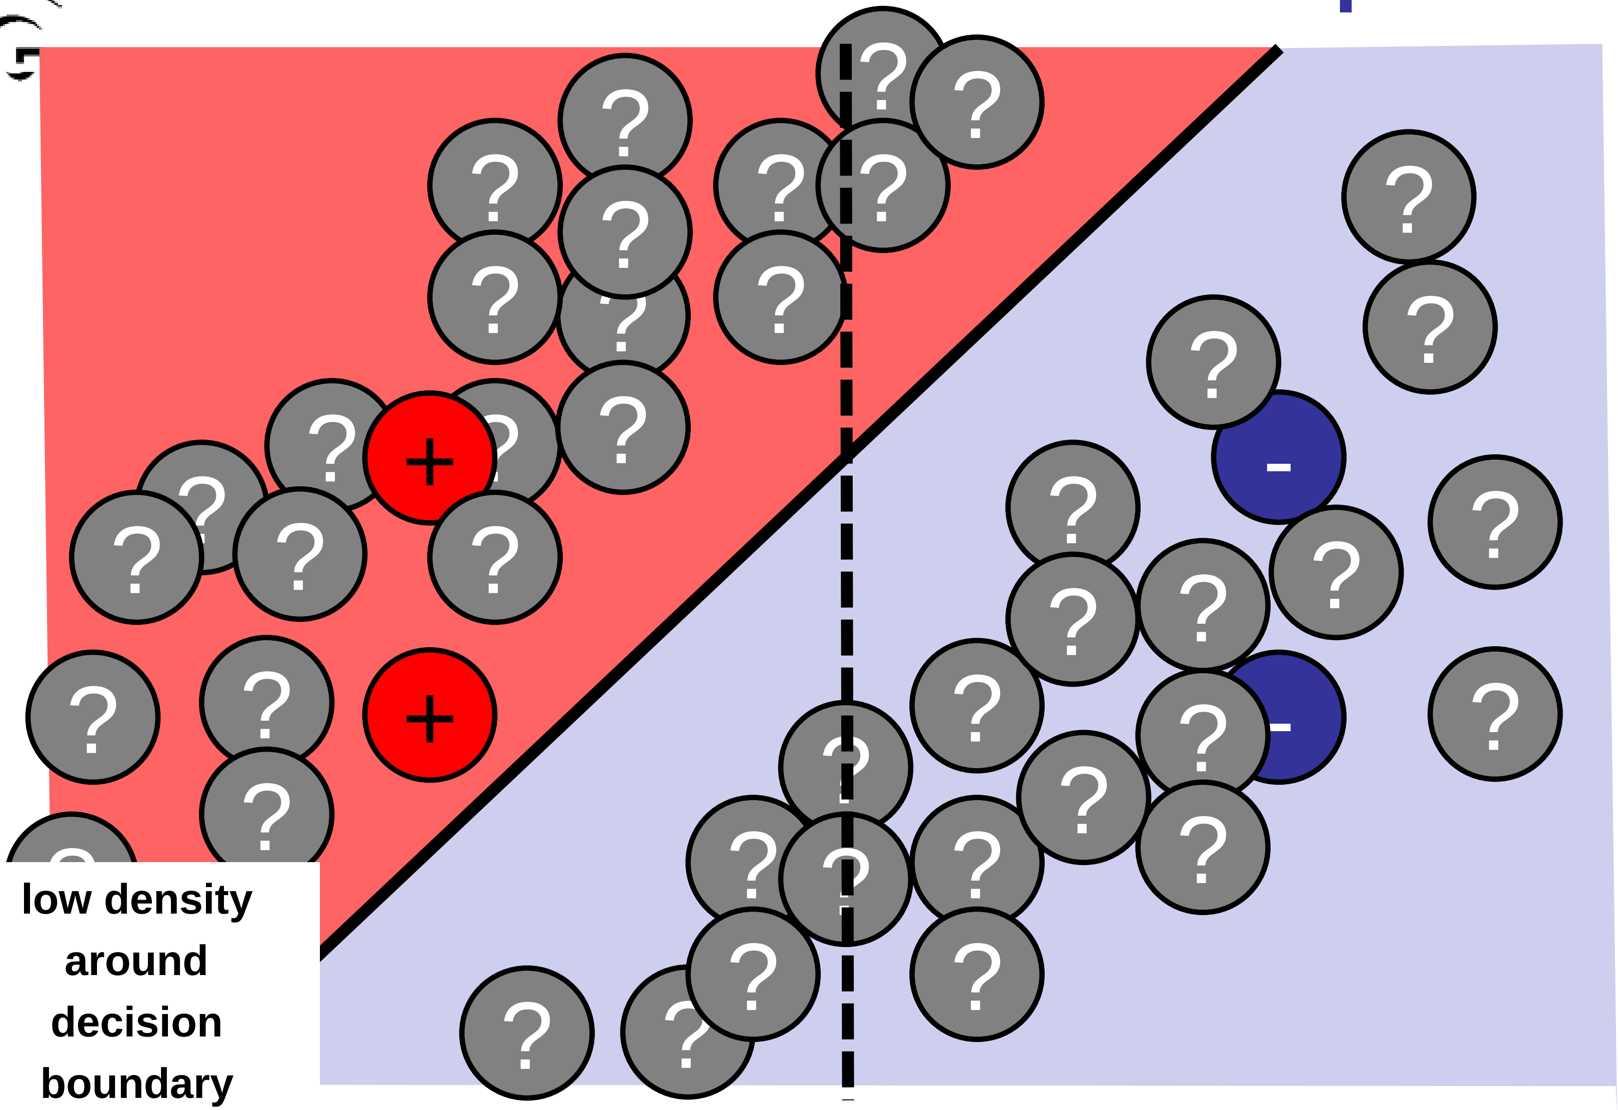

SSL is biologically plausible

- Co-training by infants [Bahrick et.al. 2002]
- Human change model once they see unlabeled data [Zahki et.al. 2007,Zhu et.al. 2007,Vandist et.al. 2007]
- Humans do On-line SSL [Zhu ICML 2010]

Why?

1. Unlabelled data is cheap/free

2. Labeled data is hard to get

   - human annotation is boring
   - labels may require experts
   - labels may require special devices
   - your graduate student is on vacation

Online Learning in Perspective

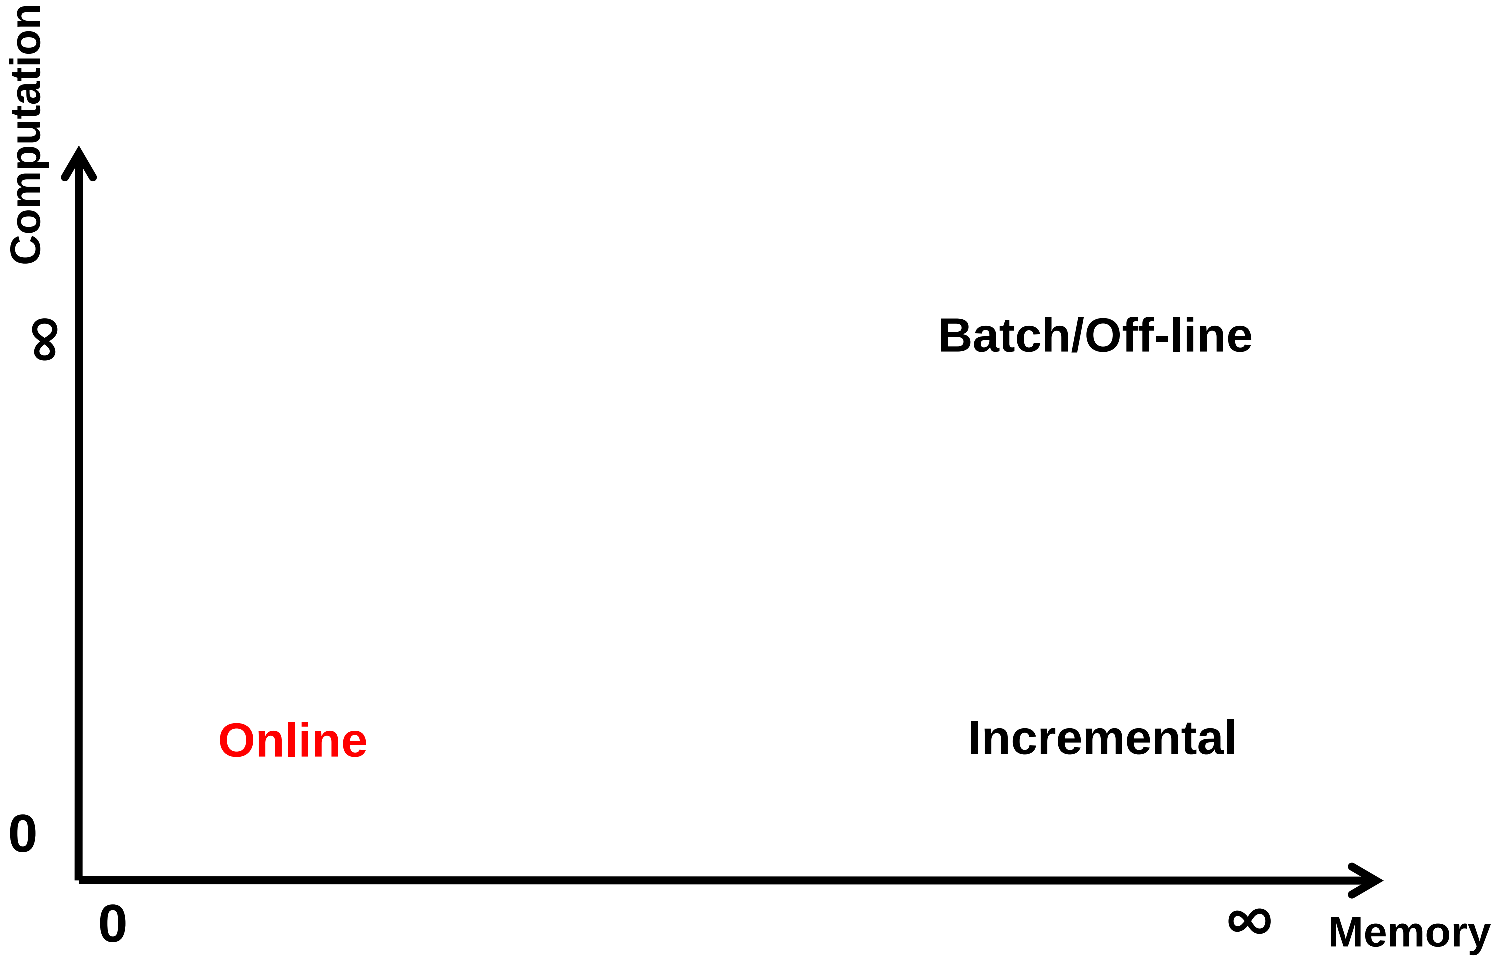

Why on-line learning?

Too much training data to fit in memory

– Internet!!!

Sample generation process

– Tracking, Co-Training

Changing processes

– Changing Environment

Why on-line learning?

Specializing

- Forget irrelevant information
- Specialize to current scene

Interactive Applications

- Data labeling
- Classifier Training
- Specializing (Human in the loop)
- Interactive Training (Segmentation)

Robustness in Learning

Noisy input data

Label noise

- Semi-supervised learning
- Weakly-labeled data
- Co-training
- Self-learning
- On-line learning

Amir Saffari
Theory

Christian Leistner
Algorithms/Applications