НИУ Высшая школа экономики – Нижний Новгород

Распознавание образов и приближенные множества

По материалам конференций:

1. Iberian Conference on Pattern Recognition and Image Analysis (IbPRIA 2017), г. Фару (Португалия), 20.06 - 23.06,
   Web: http://www.ibpria.org/2017/

2. International Joint Conference on Rough Sets (IJCRS 2017), Ольштын (Польша),
   03.07 - 07.07,
   Web: http://ijcrs2017.uwm.edu.pl

1. IbPRIA 2017
2. IJCRS 2017


Faro

Weakly-supervised learning from images and video

Ivan Laptev
WILLOW, INRIA/ENS/CNRS, Paris

Задача weakly-supervised object detection

Training input

\+

image-level labels:

✓ Person	✓ Reading
✓ Chair	× Riding bike
× Airplane	× Running
…	…

Test output

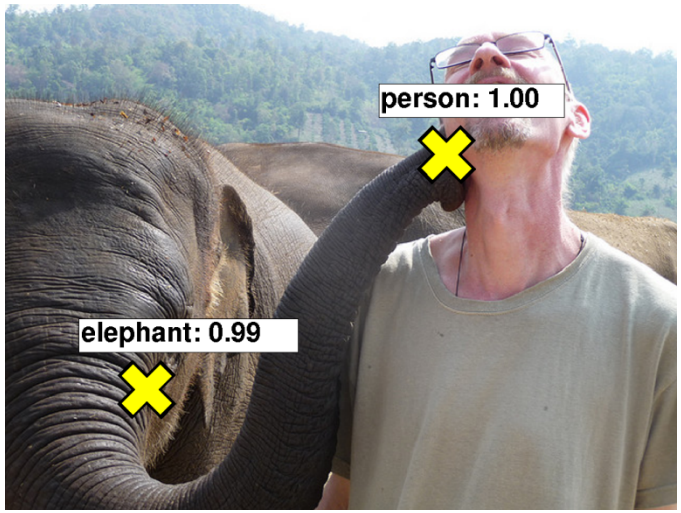

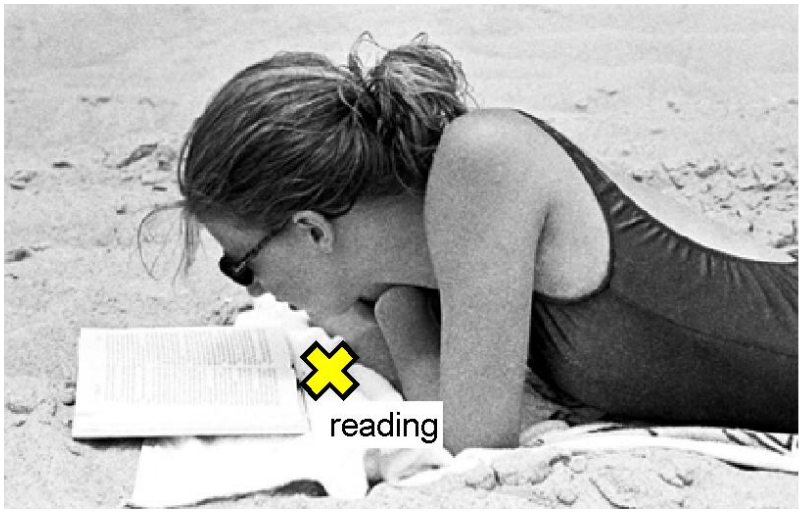

More details in http://www.di.ens.fr/willow/research/weakcnn/

Approach: search over object's location at the training time

Oquab, Bottou, Laptev and Sivic CVPR 2015

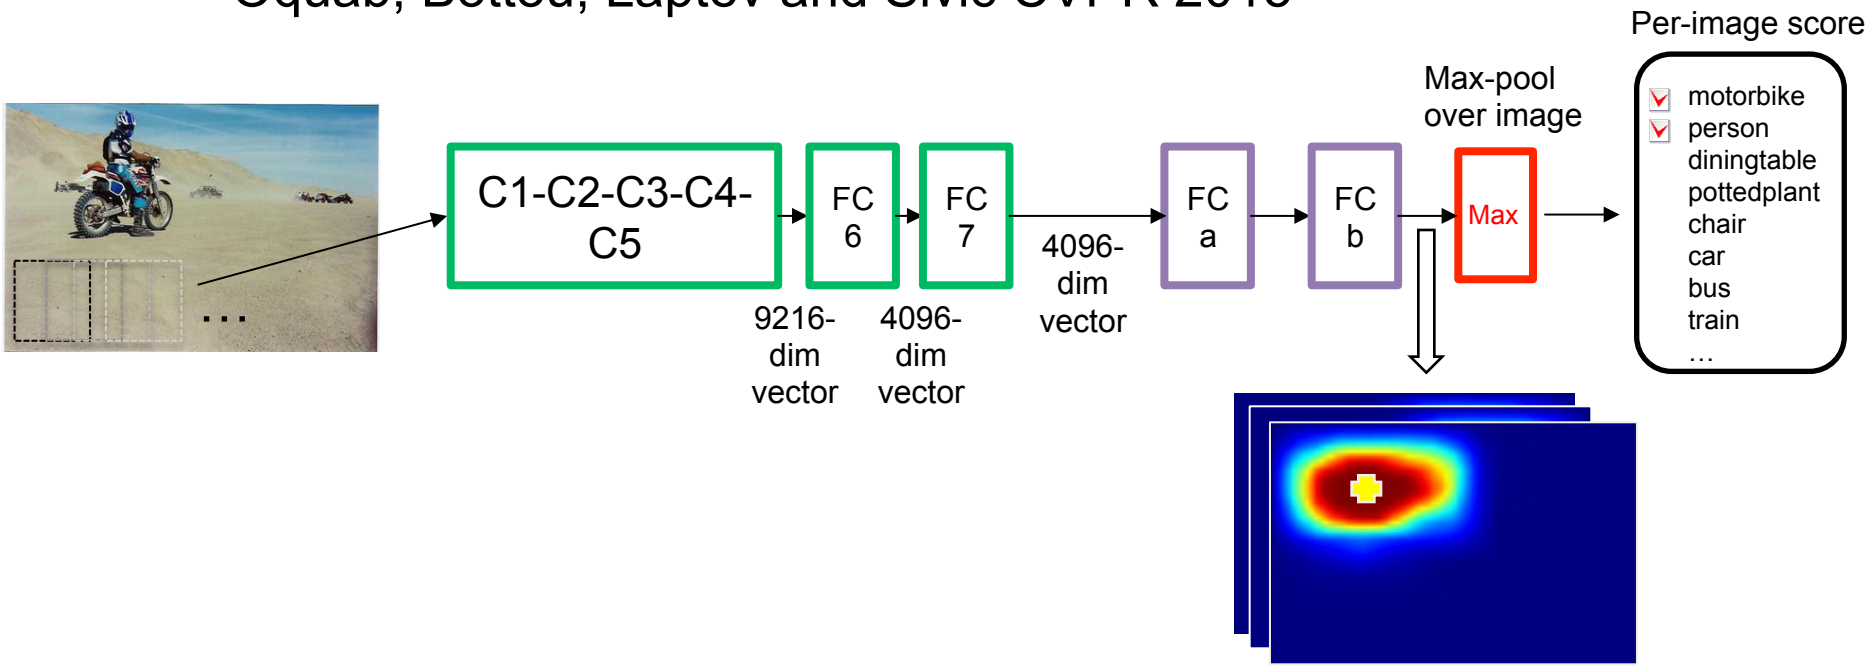

1. Fully convolutional network
2. Image-level aggregation (max-pool)
3. Multi-label loss function (allow multiple objects in image)

See also [Papandreou et al. '15, Sermanet et al. '14, Chaftield et al.'14]

Action recognition

Weak supervision from scripts: Joint Learning of Actors and Actions

[Bojanowski et al. ICCV 2013],[Miech, Alayrac, Bojanowski, Laptev, Sivic, 2017]

$$\min_{Z,T,w,b,v,c} \frac{1}{N}\|Z - \phi(X)w - b\|_F^2 + \lambda_1\ Tr(w^T\ w) + \frac{1}{N}\|T - \psi(X)v - c\|_F^2 + \lambda_2\ Tr(v^T\ v)$$

$$\begin{bmatrix} z_{11} & \dots & z_{1p} & \dots & z_{1P} \\ \vdots & & \vdots & & \vdots \\ z_{n_1 1} & \dots & z_{n_1 p} & \dots & z_{n_1 P} \\ z_{n_2 1} & \dots & z_{n_2 p} & \dots & z_{n_2 P} \\ z_{n_3 1} & \dots & z_{n_3 p} & \dots & z_{n_3 P} \\ \vdots & & \vdots & & \vdots \\ z_{N1} & \dots & z_{Np} & \dots & z_{NP} \end{bmatrix} \quad \begin{bmatrix} t_{11} & \dots & t_{1a} & \dots & t_{1A} \\ \vdots & & \vdots & & \vdots \\ t_{n_1 1} & \dots & t_{n_1 a} & \dots & t_{n_1 A} \\ t_{n_2 1} & \dots & t_{n_2 a} & \dots & t_{n_2 A} \\ t_{n_3 1} & \dots & t_{n_3 a} & \dots & t_{n_3 A} \\ \vdots & & \vdots & & \vdots \\ t_{N1} & \dots & t_{Na} & \dots & t_{NA} \end{bmatrix}$$

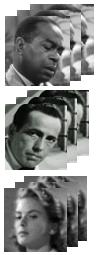

$p = Rick$

$a = Walk$

Person p appears in clip N :

$$\sum_{n \in \mathcal{N}_i} z_{np} \geq 1$$

Action a appears in clip N :

$$\sum_{n \in \mathcal{N}_i} t_{na} \geq 1$$

Person p **and** Action a appear in clip N :

$$\sum_{n \in \mathcal{N}_i} z_{np}\ t_{na} \geq 1$$

Learning from narrated instruction videos

J.-B. Alayrac, P. Bojanowski, N. Agrawal, J. Sivic, I. Laptev and S. Lacoste-Julien

CVPR 2016

Given a set of narrated instruction videos of a task

- Discover main steps
- Learn their visual and linguistic representation
- Temporally localize each step in input videos

"How to" instruction videos: changing tire

Approach: two linked clustering problems

1. Text clustering into a sequence of common steps

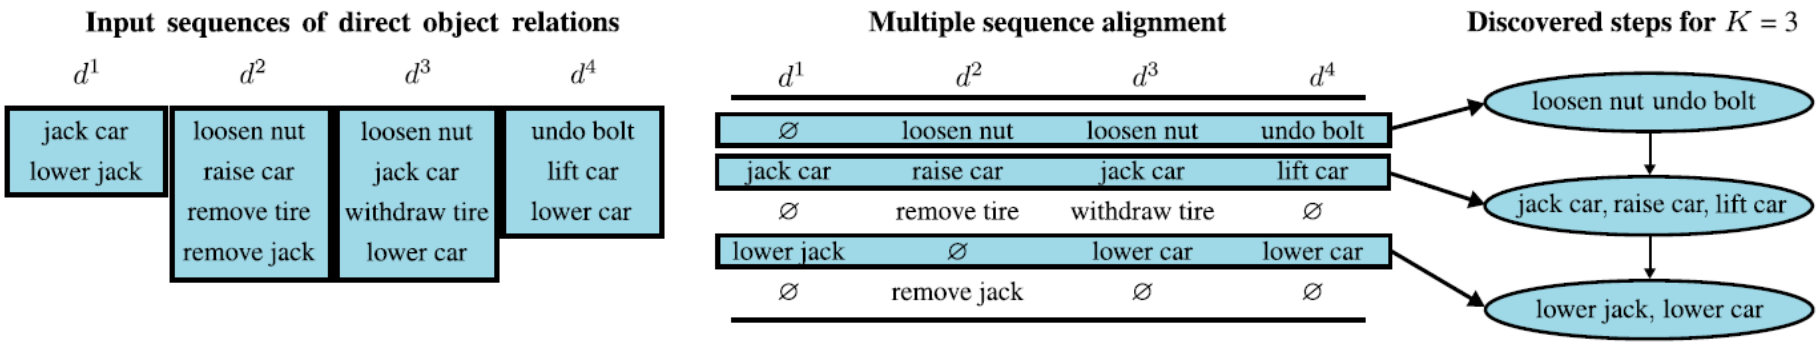

2. Video clustering to localize the actions with text constraints

$$h(Z) = \min_{W \in \mathbb{R}^{K \times d}} \underbrace{\frac{1}{2T}\|Z - XW\|_F^2}_{\text{Discriminative loss on data}} + \underbrace{\frac{\lambda}{2}\|W\|_F^2}_{\text{Regularizer}} \quad \text{s.t.} \quad \underbrace{Z \in \mathcal{Z}}_{\text{ordered script}} , \quad \underbrace{AZ \geq R}_{\text{weak textual constraints}} .$$

Discovered temporal localization [TxK matrix]

Representation of video chunks (IDTF,CNN) [Txd] matrix

Linear action classifier [dxK] matrix

Temporal constraints from text

[Bach and Harchaoui'08, Xu et al.'04, Bojanowski et al.'13,'14,'15]

Large-scale video tagging

A. Miech, J. Sivic, I. laptev, 2017

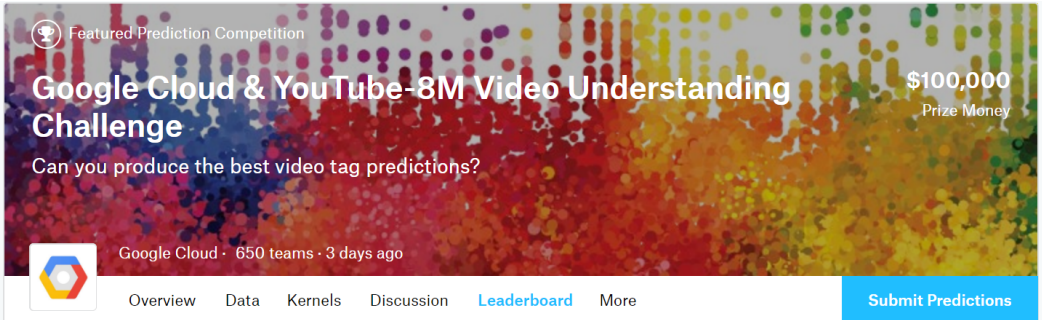

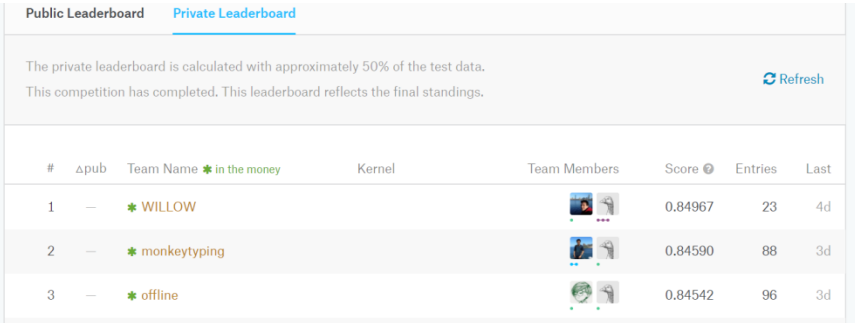

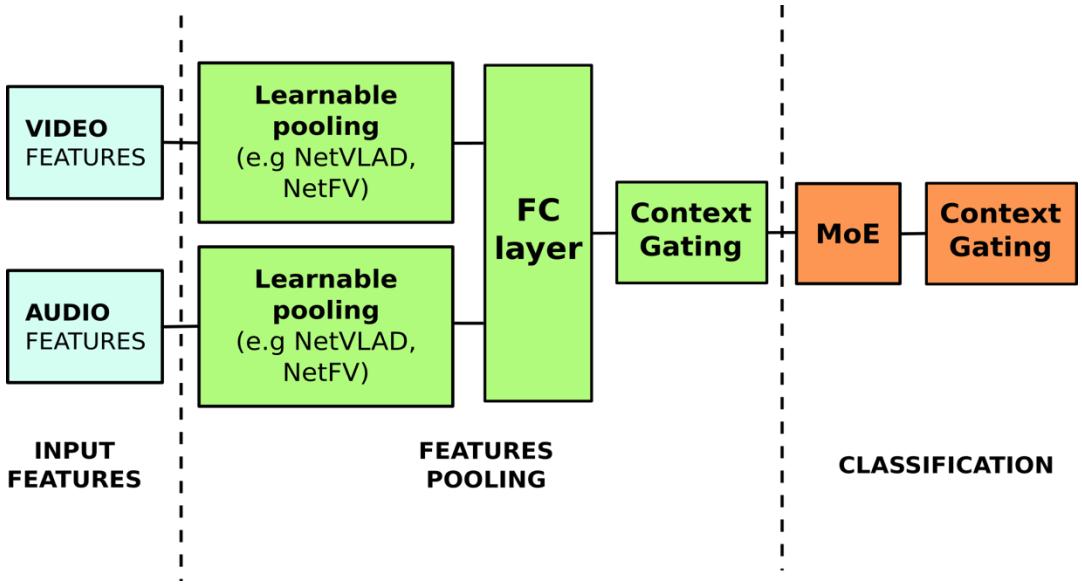

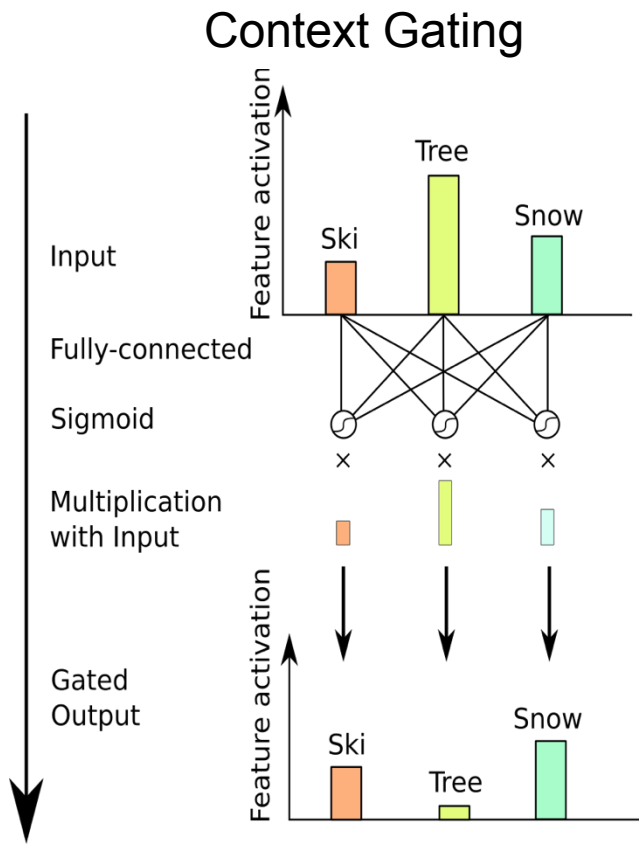

Details in http://arxiv.org/abs/1706.06905

Cost Sensitive Boosting Algorithms: Do we really need them?

Gavin Brown, University of Manchester

Cost sensitive problems – differing cost for a False Positive / False Negative

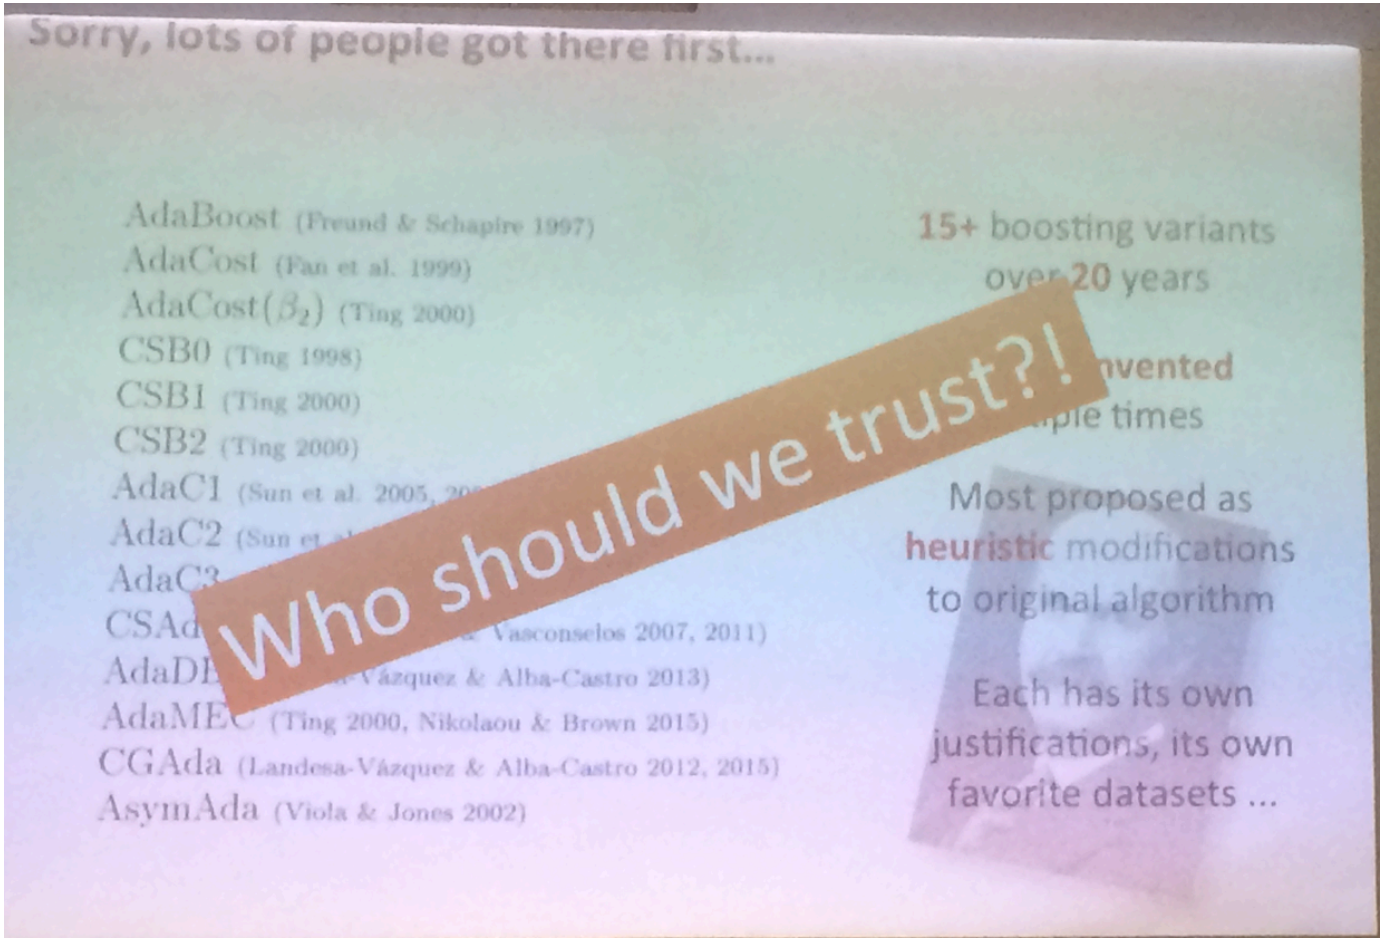

http://www.di.uoa.gr/sites/default/files/BoostingTalkAthens.pdf

https://link.springer.com/article/10.1007/s10994-016-5572-x

Почему AdaBoost интересен?

- Functional Gradient Descent (Mason et al., 2000)
- Decision Theory (Freund & Schapire, 1997)
- Margin Theory (Schapire et al., 1998)
- Probabilistic Modeling (Lebanon & Lafferty 2001; Edakunni et al 2011)

Gödel Prize 2003

	My new algorithm
Functional Gradient Descent	✓
Decision Theory	✓
Margin Theory	✓
Probabilistic Modelling	✓

"Does my new algorithm still follow from each?"

Калибровка оценок апостериорных вероятностей (1)

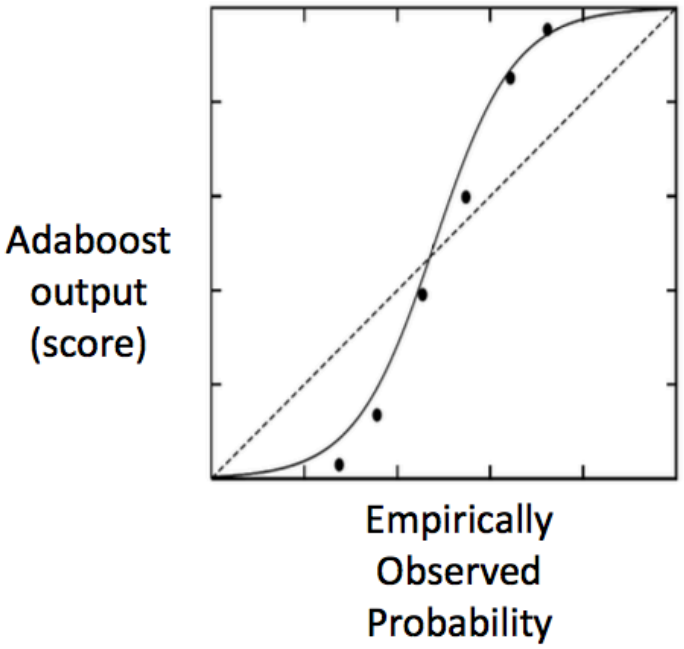

Adaboost tends to produce probability estimates close to 0 or 1.

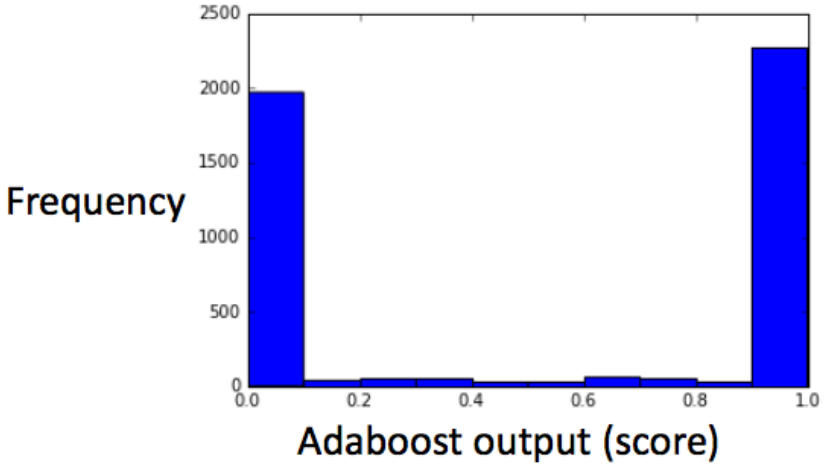

Property: Calibrated estimates

Does the algorithm generate "calibrated" probability estimates?

All algorithms produce uncalibrated probability estimates!

Калибровка оценок апостериорных вероятностей (2)

Platt scaling (logistic calibration)

(Platt, 1999)

Training: Reserve part of training data (here 50% -more on this later) to fit a sigmoid to correct the distortion:

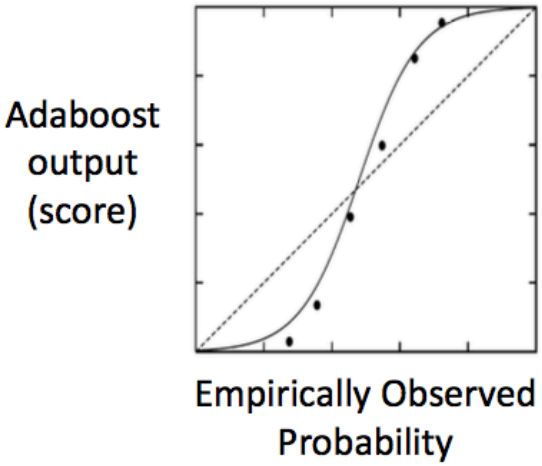

$$\hat{p}(y=1|\mathbf{x}) = \frac{1}{1+e^{As(\mathbf{x})+B}}$$

Prediction: Apply sigmoid transformation to score (output of ensemble) to get probability estimate

Параметры A,B определяются с помощью метода максимума правдоподобия

Recognizing Activities of Daily Living from Egocentric Images

Alejandro Cartas, Juan Marın, Petia Radeva, Mariella Dimiccoli
University of Barcelona

Egocentric (first-person) wearable cameras

Drinking/eating alone

AlexNet Top 5

#	Activity	Score
1	Public Transport	0.1864
2	Cooking	0.1464
3	Eating together	0.1382
4	**Drinking/eating alone**	**0.1223**
5	Cleaning and chores	0.1067

GoogLeNet Top 5

#	Activity	Score
1	Plane	0.2004
2	Public Transport	0.1943
3	Cleaning and chores	0.1450
4	Cooking	0.0925
5	**Drinking/eating alone**	**0.0748**

Meeting

AlexNet Top 5

#	Activity	Score
1	Talking	0.2034
2	**Meeting**	**0.1701**
3	Cooking	0.1090
4	Cleaning and chores	0.0903
5	Shopping	0.0878

GoogLeNet Top 5

#	Activity	Score
1	**Eating together**	**0.3065**
2	Talking	0.1215
3	Socializing	0.1062
4	Cleaning and chores	0.08126
5	Meeting	0.0664

Reading

AlexNet Top 5

#	Activity	Score
1	**Drinking/eating alone**	**0.5002**
2	Cleaning and chores	0.1511748880
3	Eating together	0.1263086796
4	Shopping	0.0589886233
5	Drinking together	0.0251834411

GoogLeNet Top 5

#	Activity	Score
1	**Cleaning and chores**	**0.4259**
2	Eating together	0.1145323068
3	Drinking/eating alone	0.1137253270
4	Drinking together	0.0841688886
5	Reading	0.06203

The result of the ensembles AlexNet+RF on FC6 and GoogLeNet+RF on Pool5/7x7+prob is highlighted on color in its corresponding table.
The green and red colors means true positive and false positive classification

1. IbPRIA 2017
2. IJCRS 2017


2


20.VII.1944 - ZAMACH - ATTENTAT - ASSASINATION ATTEMPT
Płk. Claus von Stauffenberg
8:00 Szef sztabu armii zapasowej płk.Claus von Stauffenberg wraz z adiutantem por. W. Heaftenem z lotniska Rangsdorf pod Berlinem samolotem udaje się na naradę do Wilczego Szańca. Ma tam przedstawić Hitlerowi koncepcje użycia armii rezerwowej na froncie wschodnim.
10:15 Samolot ląduje na lotnisku pod Kętrzynem, skąd samochodem obaj przyjeżdżają do kwatery.Stauffenberg ma w teczce dwie bomby z zapalnikami chemicznymi.
11:30 Narada wstępna u feldmarszałka W.Keitela.
12:20 Stauffenberg przy pomocy adiutanta uzbraja jedną bombę (drugiej nie zdążył) i udaje się z nią do baraku narad. Tam teczkę z bombą
Barak narad po wybuchu
Blick in den Lagebsprechungsraum in der Lagebaracke nach dem Attentat am 20. Juli 1944
Kartentisch
Radio-truhe
Przybliżone miejsca, zajmowane przez uczestników narady 20 lipca 1944 r.
1. Adolf Hitler
2. Gen. Adolf Heusinger, szef oddziału operacyjnego wojsk lądowych
† 3. Gen. Gunther Korten, szef sztabu generalnego wojsk lądowych
† 4. Płk Heinz Brandt z-ca gen. Adolfa Heusingera
5. Gen. Karl Bodenschatz, przedstawiciel naczelnego dowództwa wojsk lądowych
6. Płk Heinz Waizenegger, adiutant feldmarszałka Keitla
Situationsskizze der Lagebesprechung am 20. Juli 1944
1. Adolf Hitler
2. Generalleutnant A. Heusinger
3 General der Flieger G. Körten
4 Oberst i.G. H. Brandt
5 General der Flieger K. H. Bodenschatz
6 Oberstleutnant G. H. Waizengger
7 Generalleutnant R. Schmundt
Zdjęcie drewnianego braraku obudowanego cegłą i betonem
H. Göring im Gespräch mit A. Hitlers persönlichen Adjutanten SS-Führer Julius Schaub; vorn: General der Flieger Bruno Loerzer (links) und Reichsleiter M.Bormann (rechts), am 20. Juli 1944

Toward data-driven education (1)

Rakesh Agrawal, Microsoft

Задачи:

1. A study plan is the choice of concepts and the organization and sequencing of the concepts in an educational course.
Task: Identify which concepts should be studied together and how students should move from one group of concepts to another.
A data-driven method, which given a list of concepts can automatically propose candidate plans to cover all the concepts.
[Journal of Educational Data Mining, 2016]
2. How to enhance the quality of the electronic textbooks? [ACM DEV 2010]
3. Grouping Students for Maximizing Learning from Peers [Educational Data Mining 2017]

Toward data-driven education (2). Структурирование учебника

Video Augmentation: Make inaccessible accessible

Table of contents for navigating the book *(automatically extracted)*

Re-rendered section: This section, about the laws of chemical combination, prescribes an activity for the chemistry lab, but the school might lack the lab to do the experiments

Augmentations panel: Video demonstrates the reaction for the second set of chemicals prescribed

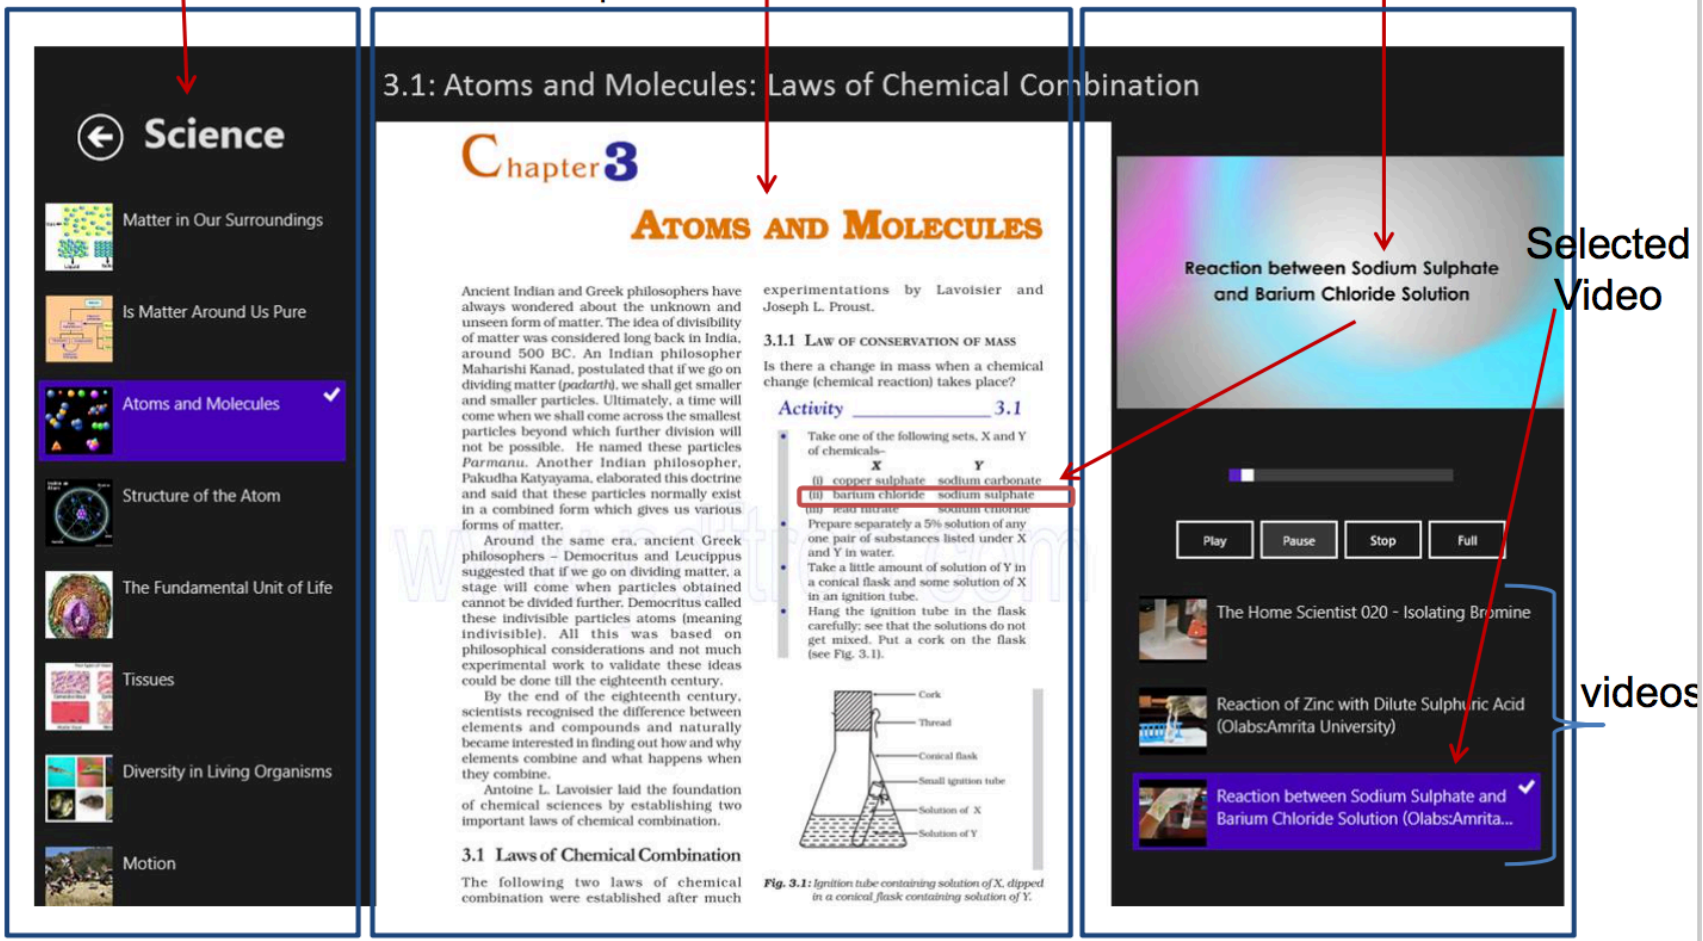

Toward data-driven education (3). Dispersion of Key Concepts in Section

Many unrelated concepts -> Hard to understand section

• V = set of key concepts discussed in section s

– Terminological noun phrases: Linguistic pattern A^*N^+ (A: adjective; N: noun)

– "concepti" Wikipedia titles

• Related(x,y) = Concept x is related to concept y – Co-occurrence

– true if Wikipedia article for x links to the article for y

• Dispersion(s):=Fraction of unrelated concept pairs

– (1 – Edge Density) of the concept graph

Toward data-driven education (4)

Avoid repetition across sections:

$$\max \sum_{i \in I} \sum_{j \in S} x_{ij} \cdot \lambda_{ij}$$

Relevance score of object *i* to section *j*

Total relevance score for the chapter: sum of relevance scores of objects assigned

s.t.

$$x_{ij} \in \{0, 1\} \ \forall i \in I \forall j \in S$$

=1 if image *i* is selected for section *j* else 0

$$\sum_{i \in I} x_{ij} \leq K_j \ \forall j \in S$$

Constraint: At most K_j images can be assigned to section *j*

$$\sum_{j \in S} x_{ij} \leq 1 \ \forall i \in I$$

Constraint: An image can belong to at most one section

Can be solved optimally in polynomial time

Toward data-driven education (5). Выделение подгрупп студентов

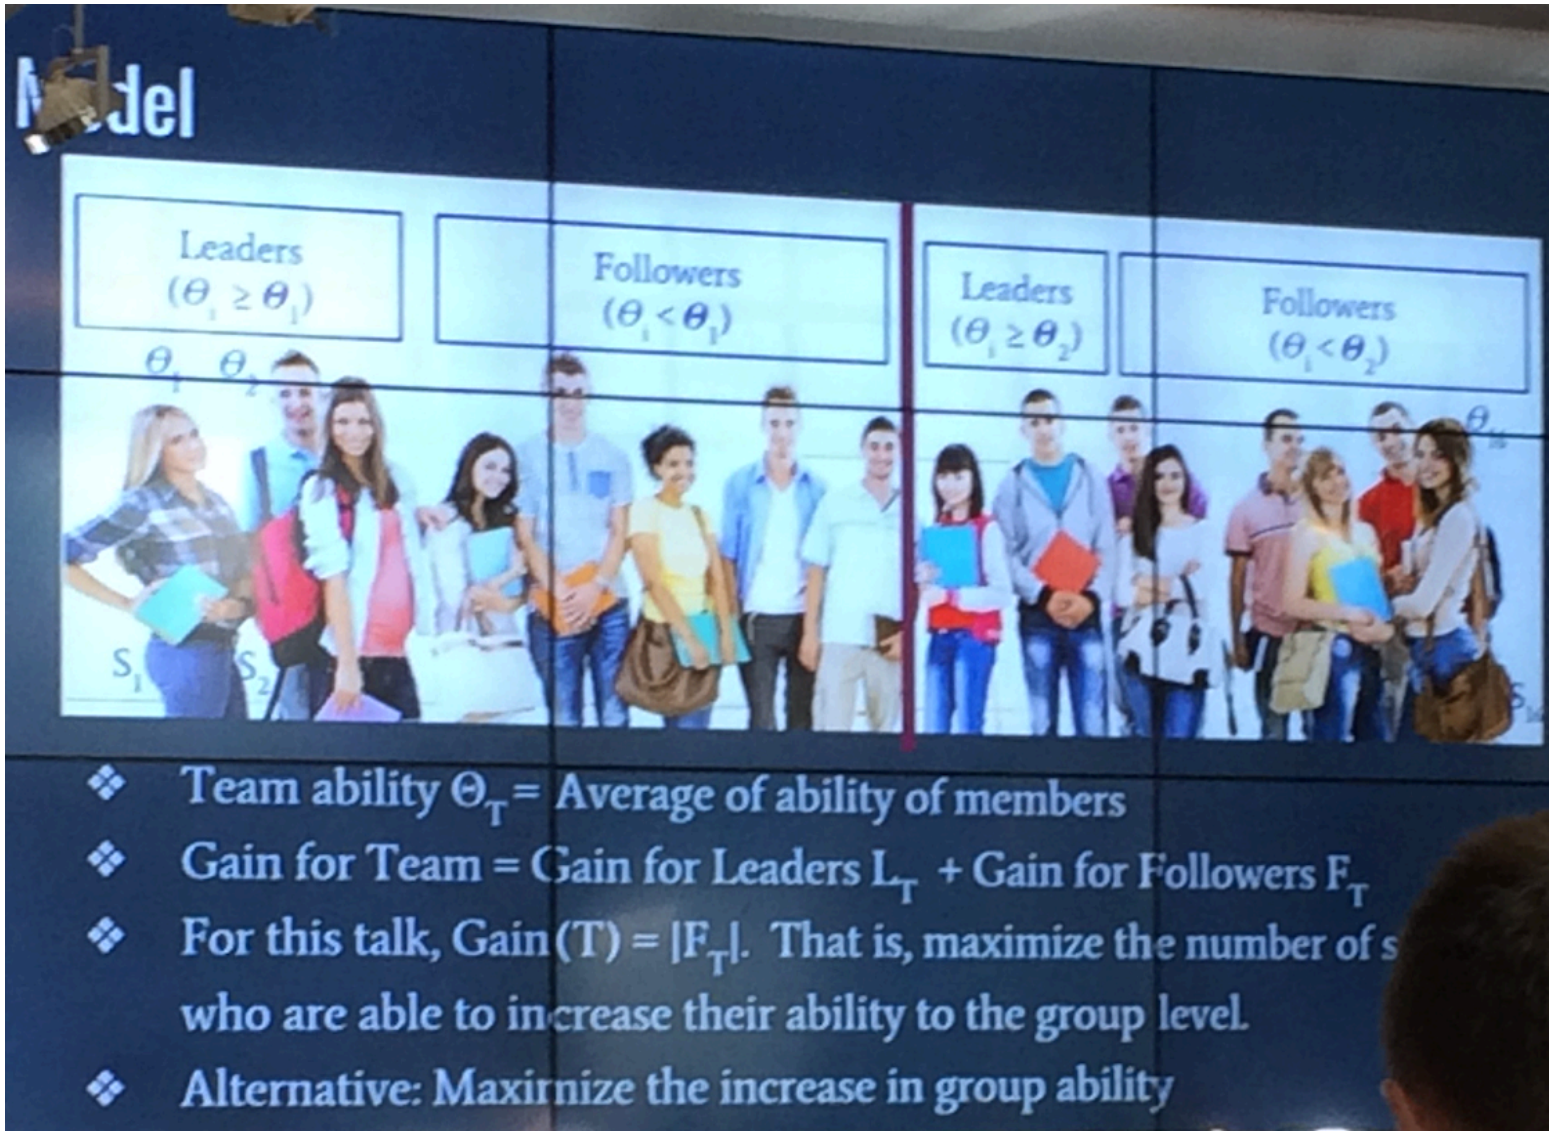

Toward data-driven education (4)

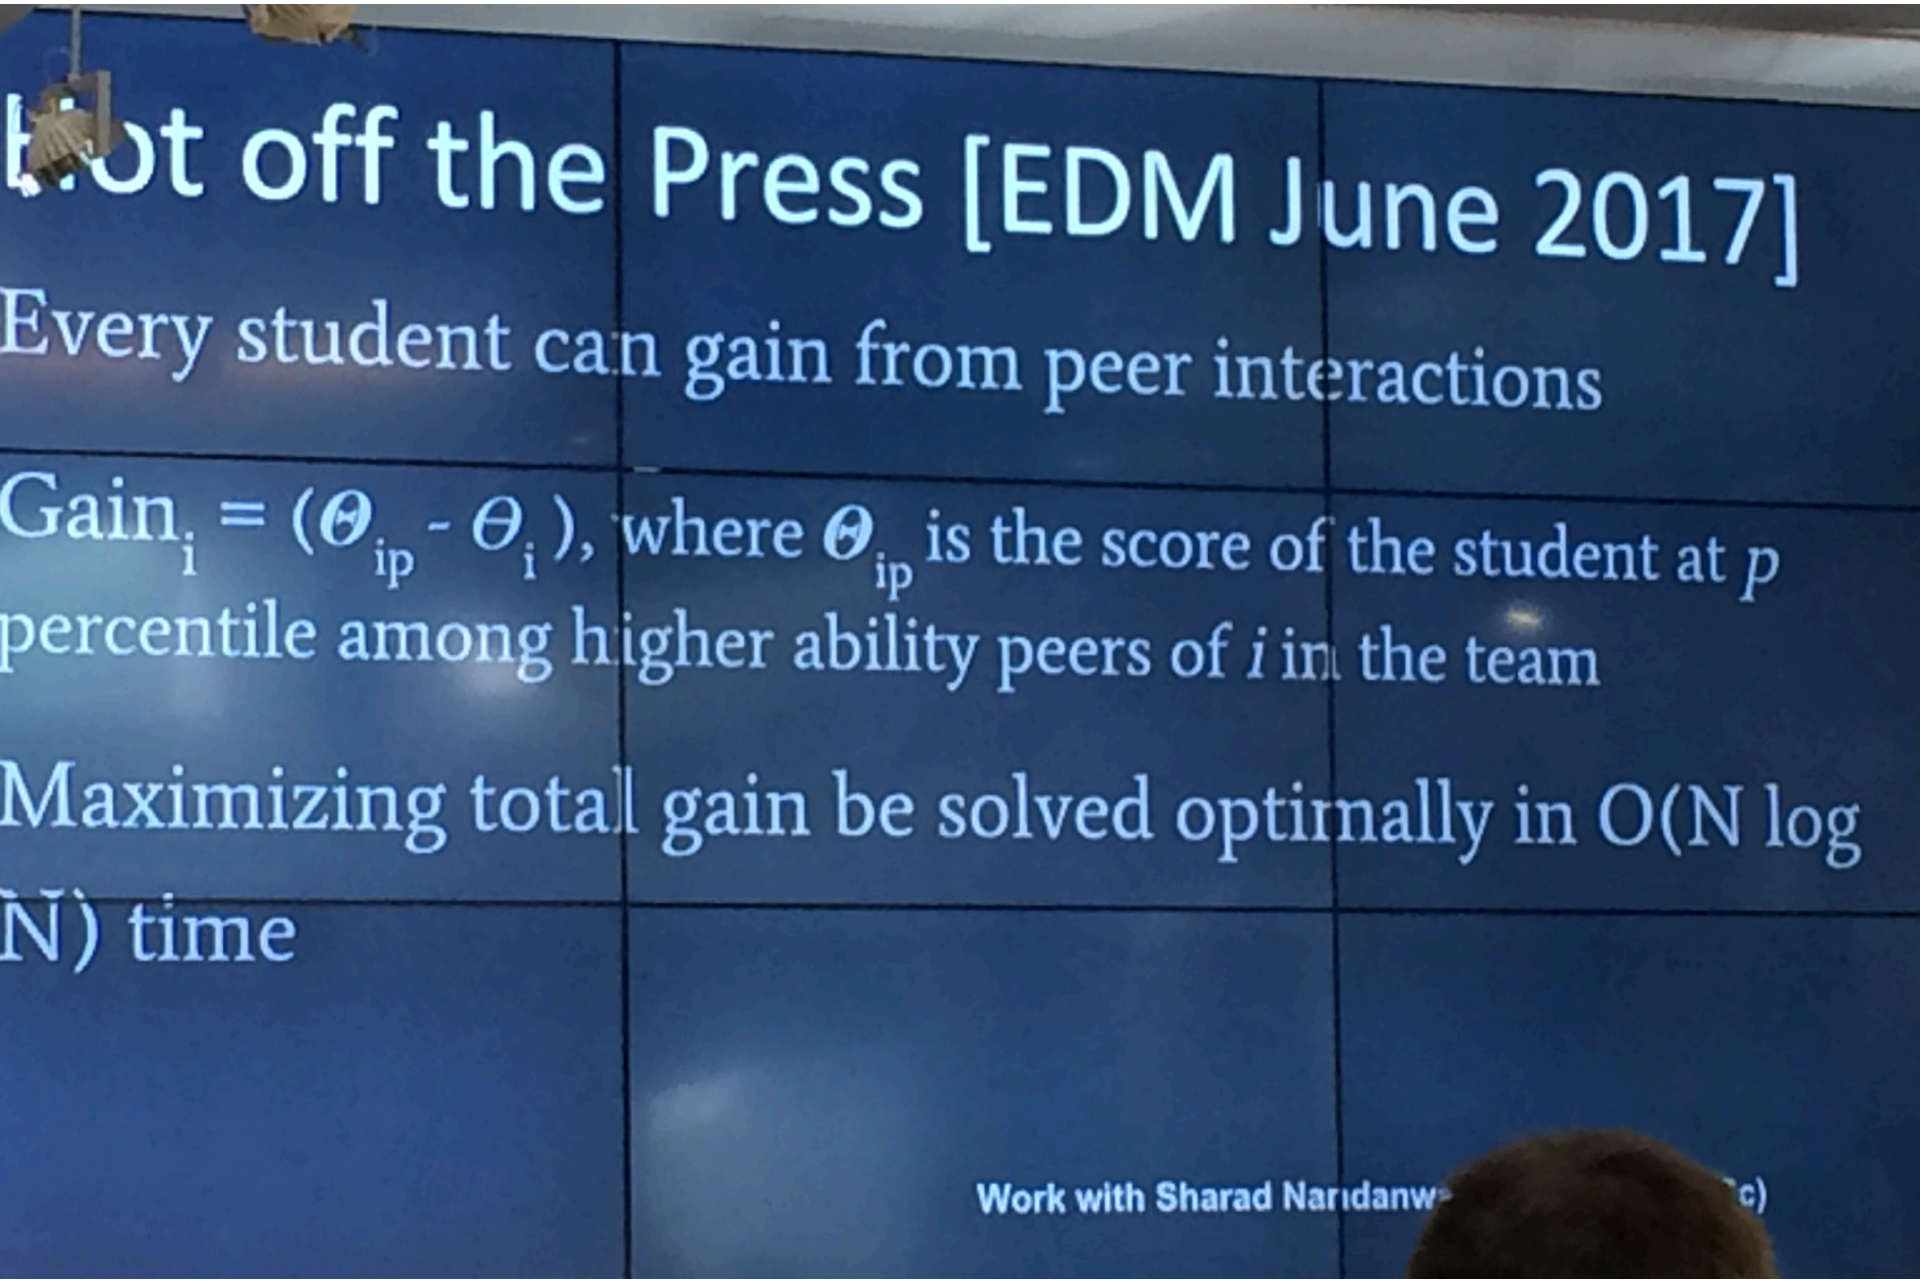

Advances in Rough Set Based Hybrid Approaches for Medical Image Analysis

Pradipta Maji, Indian Statistical Institute

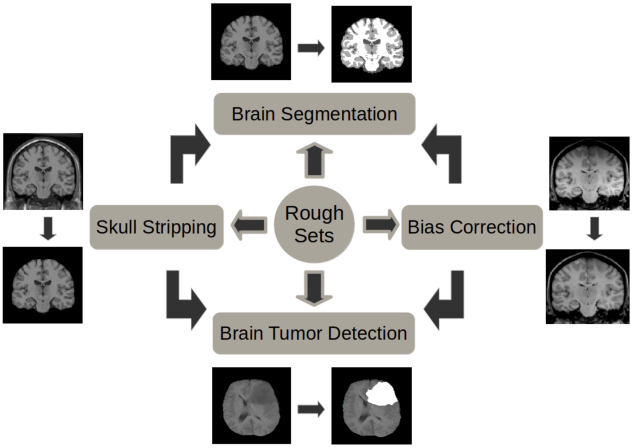

Stomped-t distribution: A. Banerjee, P. Maji/ Information Sciences 421 (2017) 104–125

$$f(y;\mu,\sigma^2,k)=\frac{1}{D}\frac{1}{\sigma}\phi(z),\ y\in\mathbb{R};\quad \text{where } z=\begin{cases}k, & \text{if } |\frac{y-\mu}{\sigma}|<k\\ \frac{y-\mu}{\sigma}, & \text{otherwise,}\end{cases}$$

$$D=2(1-\Phi(k)+k\phi(k))$$

φ() and Φ() are, respectively, the pdf and probability distribution functions of standard normal distribution

Rough-probabilistic clustering

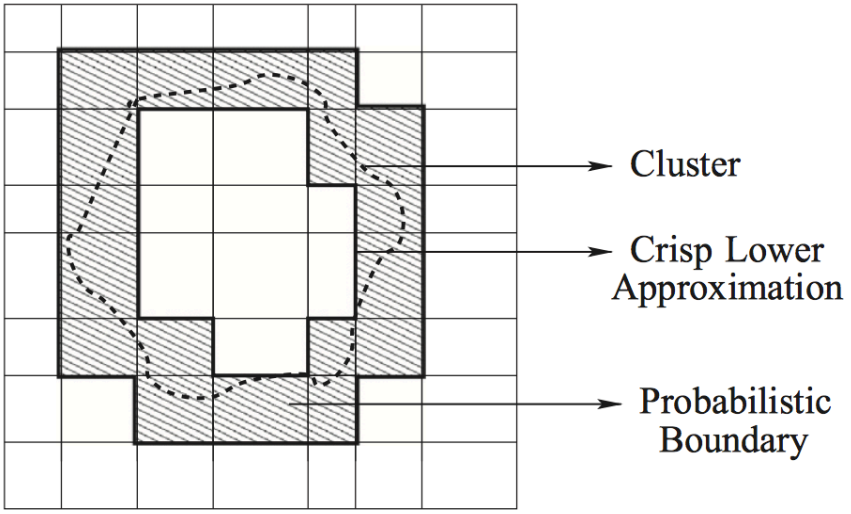

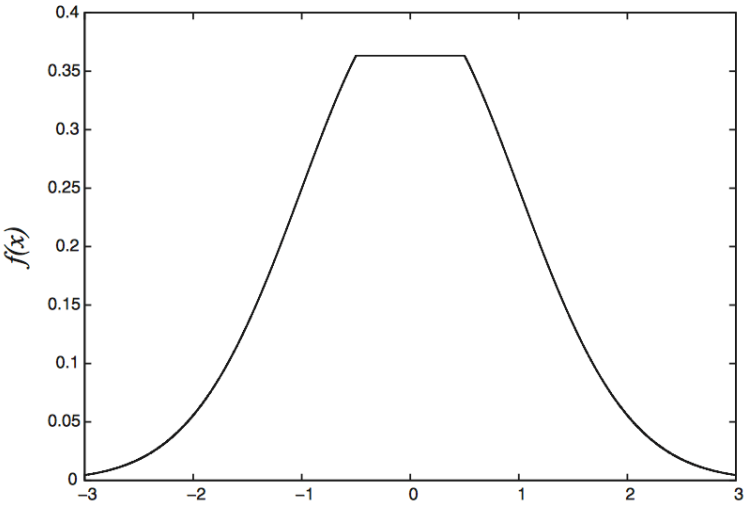

More with less: A new paradigm in modern Machine Learning

Nguyen Hung Son, University of Warsaw

Most of deep learning approaches:

- rely on the availability of huge amounts of data,
- often requiring millions of correctly labelled examples.

We will discuss the newest learning techniques for the case when we have

- a huge amount of data but
- very little amount of labelled data.

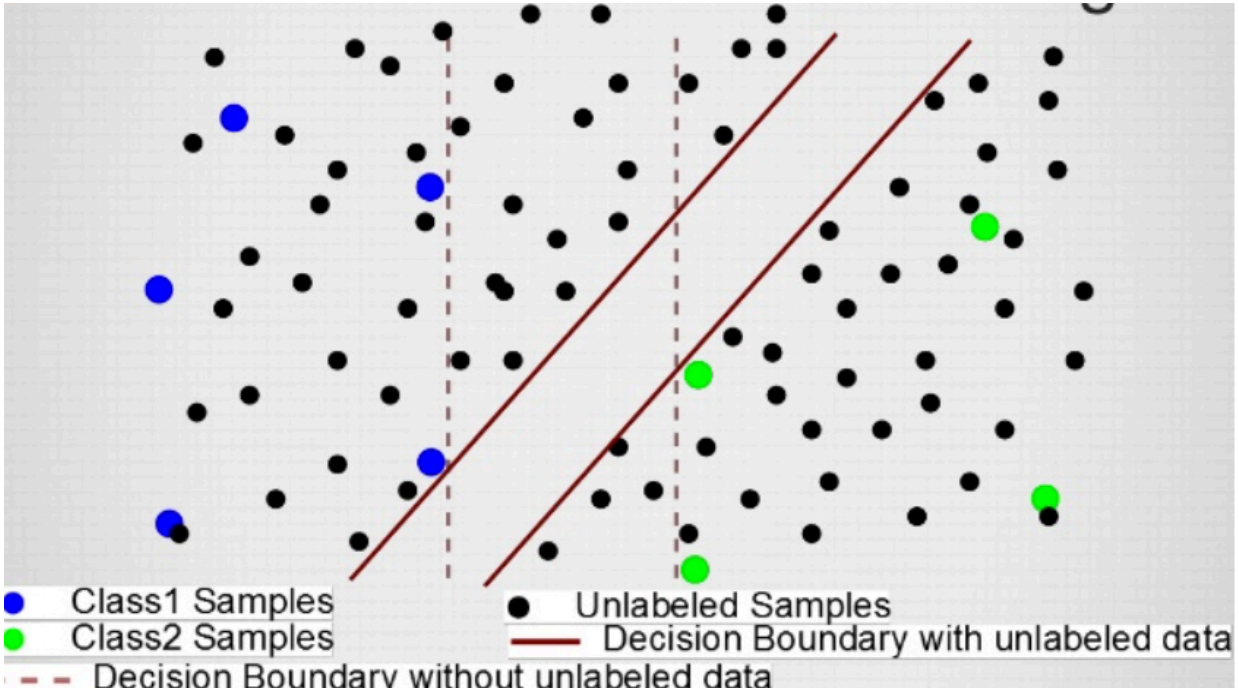

Big data analysis by rough sets and granular computing

Tianrui Li

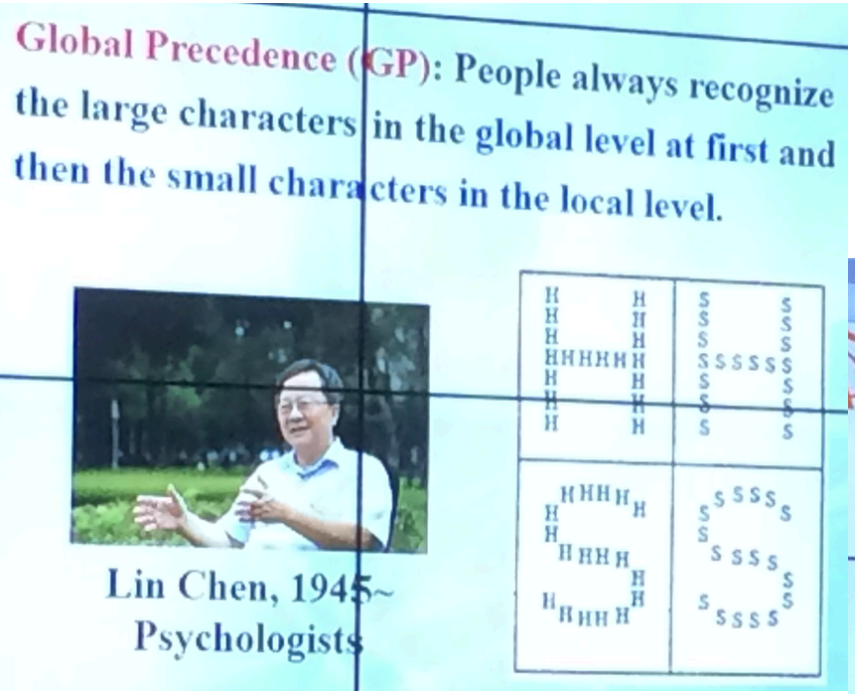

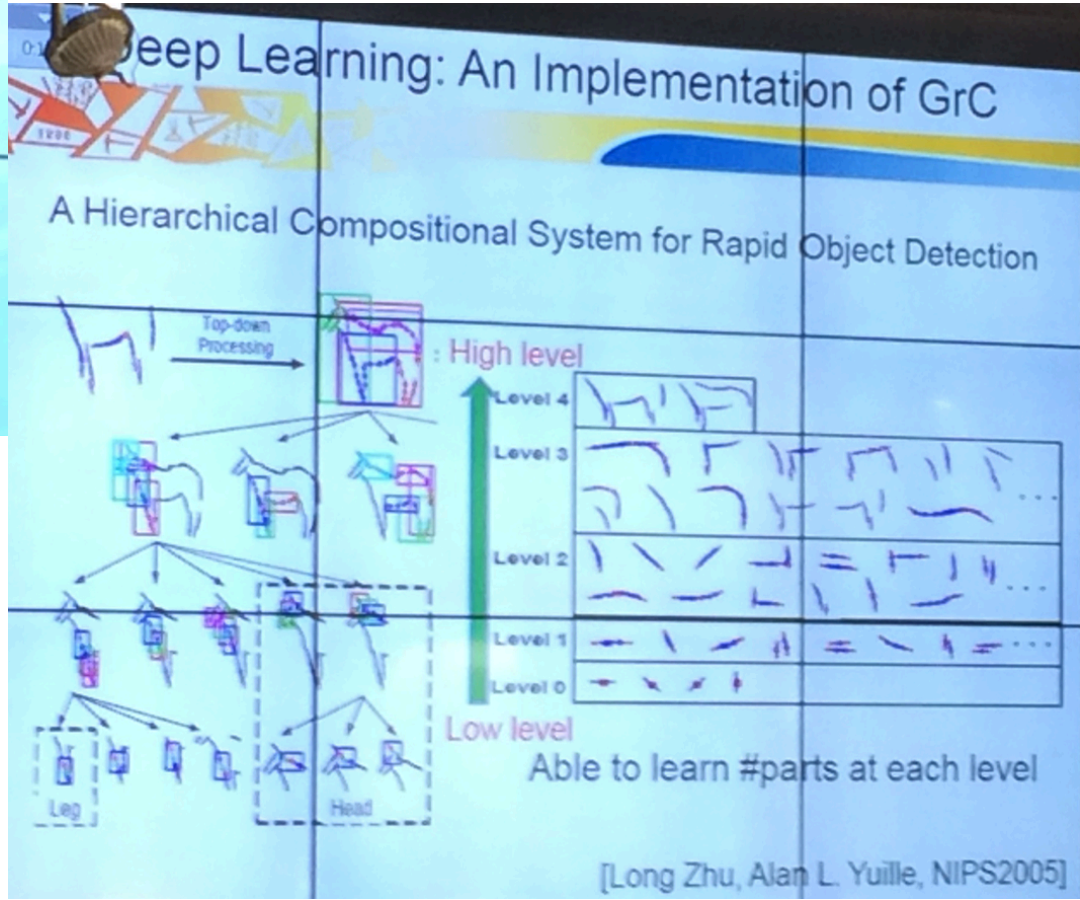

Three-way decisions (TWD)

Yiyu Yao, University of Regina (Canada)

Yiyu Yao, Three-way decisions with probabilistic rough sets, Information Sciences

"Rules constructed from the three regions are associated with different actions and decisions, which immediately leads to the notion of three-way decision rules. A **positive** rule makes a **decision of acceptance**, a **negative** rule makes a **decision of rejection**, and a **boundary** rule makes a **decision of abstaining**"