Personal Sensing

Tarek Abdelzaher
Dept. of Computer Science
University of Illinois at Urbana Champaign

Review: Localization with a Single LED

- Can you simultaneously localize a large number of optical receivers using a single "smart" LED?

Song Liu and Tian He, "SmartLight: Light-weight 3D Indoor Localization Using a Single LED Lamp," ACM Sensys, Delft, The Netherlands, November 2017

Idea #1

- Your location determines what you see:

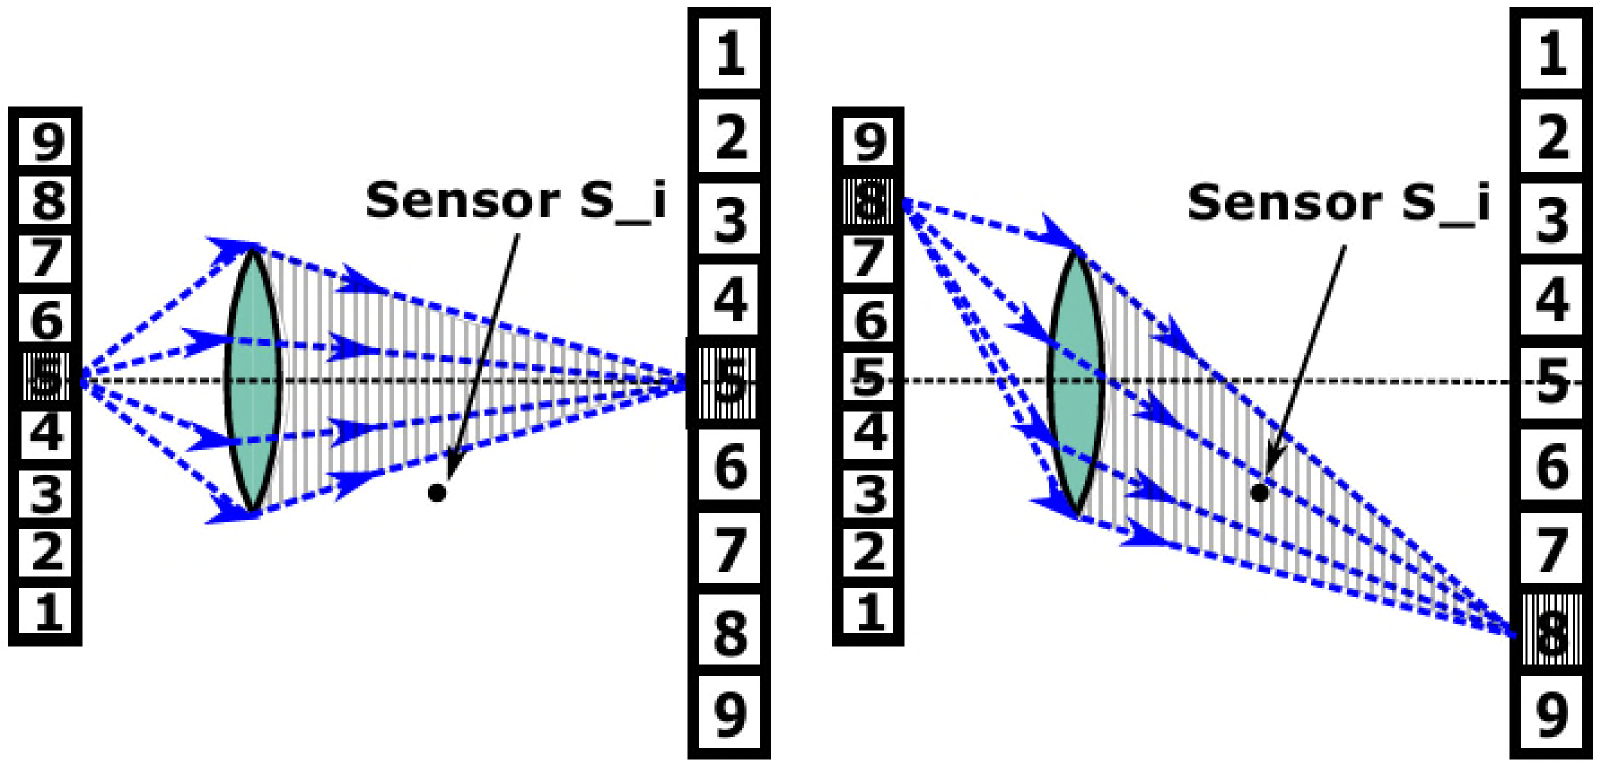

Idea #2

- Distance determines size of visible area

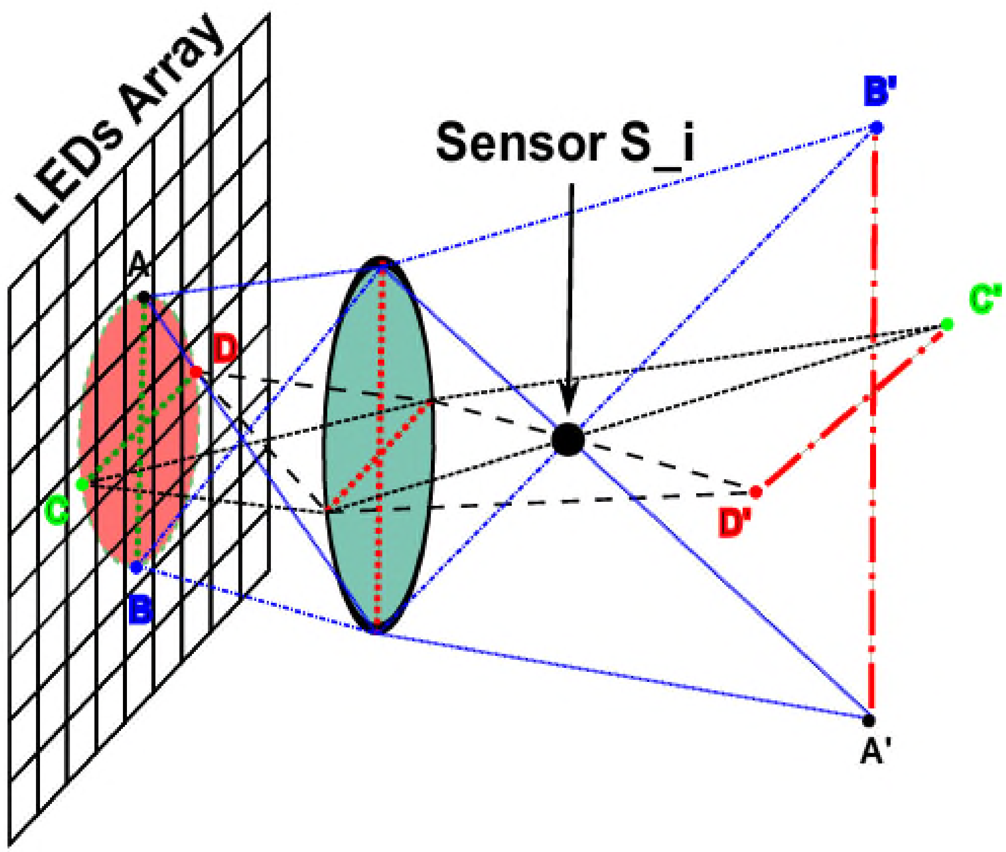

Testbed

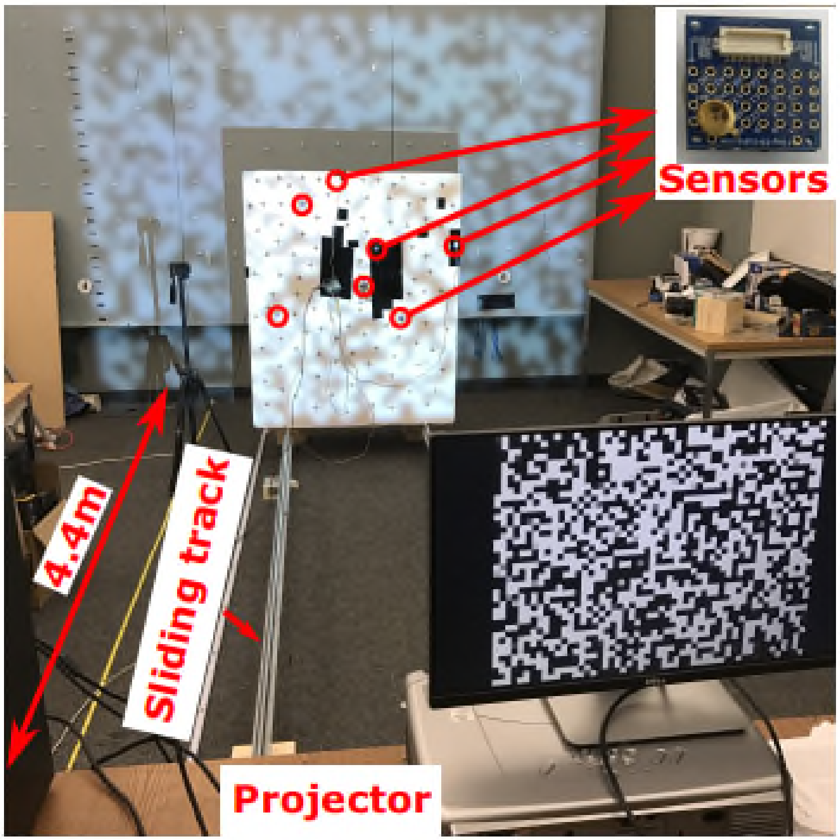

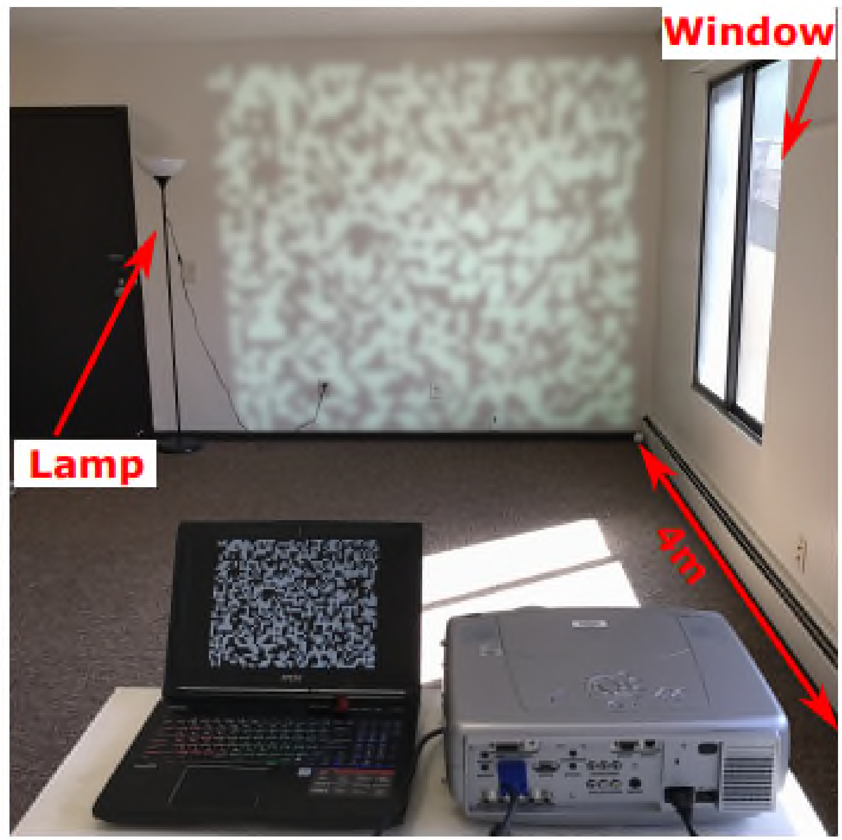

Evaluation

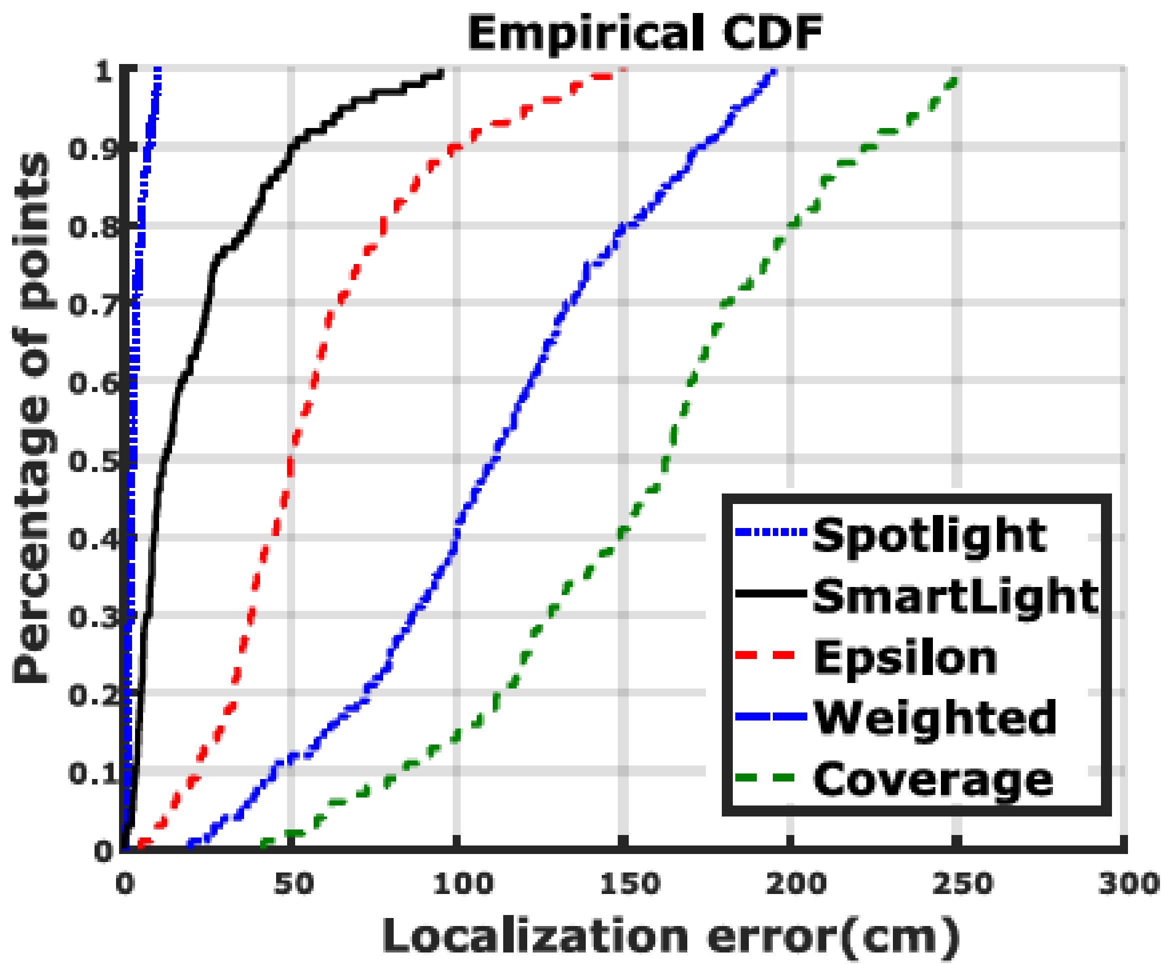

Empirical CDF
Percentage of points
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
50
100
150
200
250
300
Localization error(cm)
Spotlight
SmartLight
Epsilon
Weighted
Coverage

Localization with Aircraft Signals

- Can we use aircraft signals to localize mobile devices even in-doors (where there is no GPS)?

Manuel Eichelberger, Kevin Luchsinger, Simon Tanner, and Roger Wattenhofer, “Indoor Localization with Aircraft Signals,” ACM Sensys, Delft, The Netherlands, November 2017

Localization with Aircraft Signals

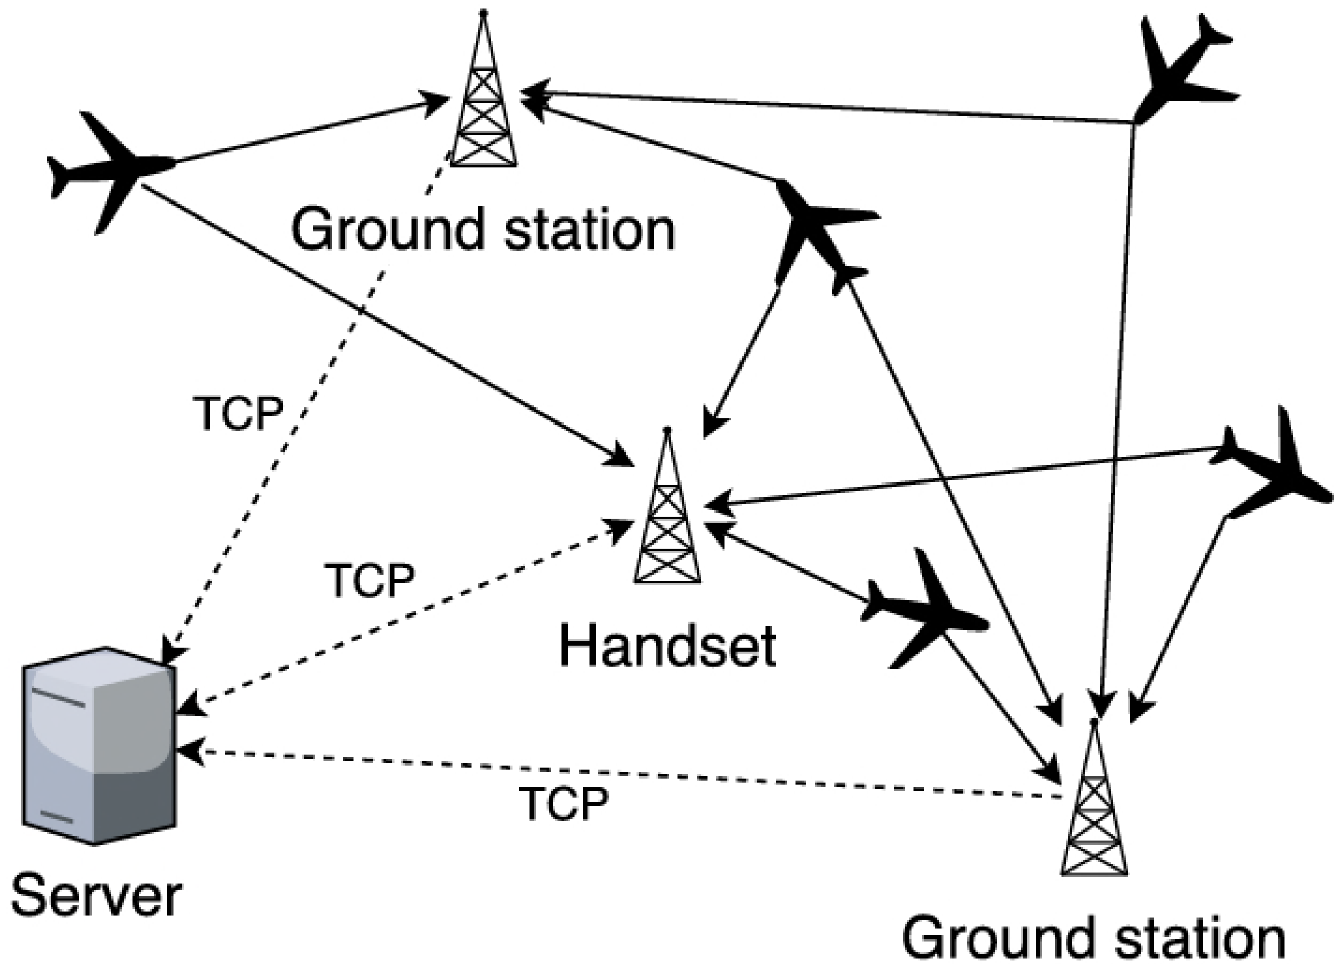

Evaluation

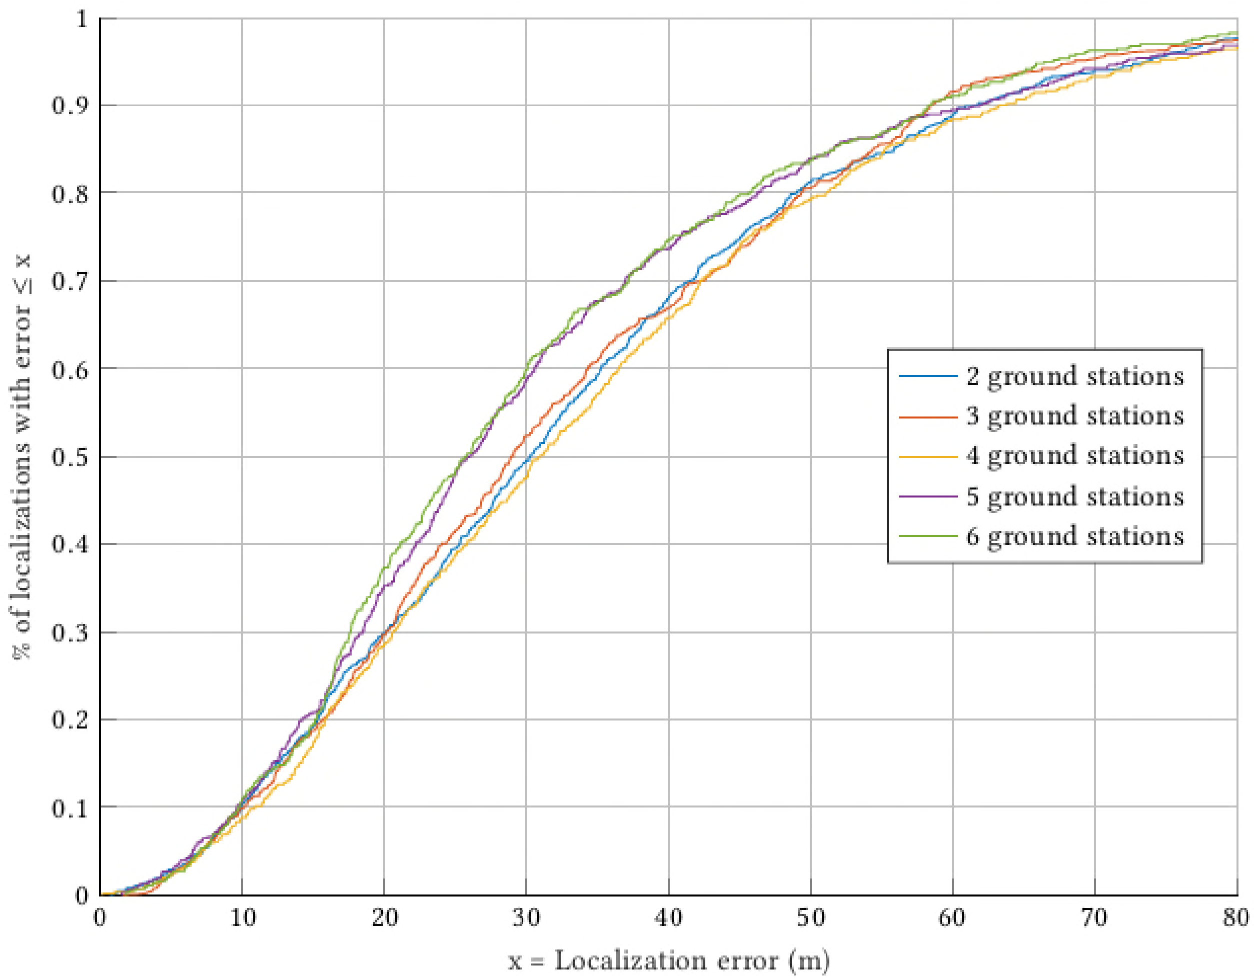

% of localizations with error ≤ x
x = Localization error (m)
2 ground stations
3 ground stations
4 ground stations
5 ground stations
6 ground stations
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
10
20
30
40
50
60
70
80

What Can Smart Phones Do (In Personal Sensing)?

What's next in personal/context sensing?

- Deep learning?
- Augmented reality?
- Virtual Reality?

Processors in the Smartphone Market

Platform and Datasets

- Snapdragon 800 SoC

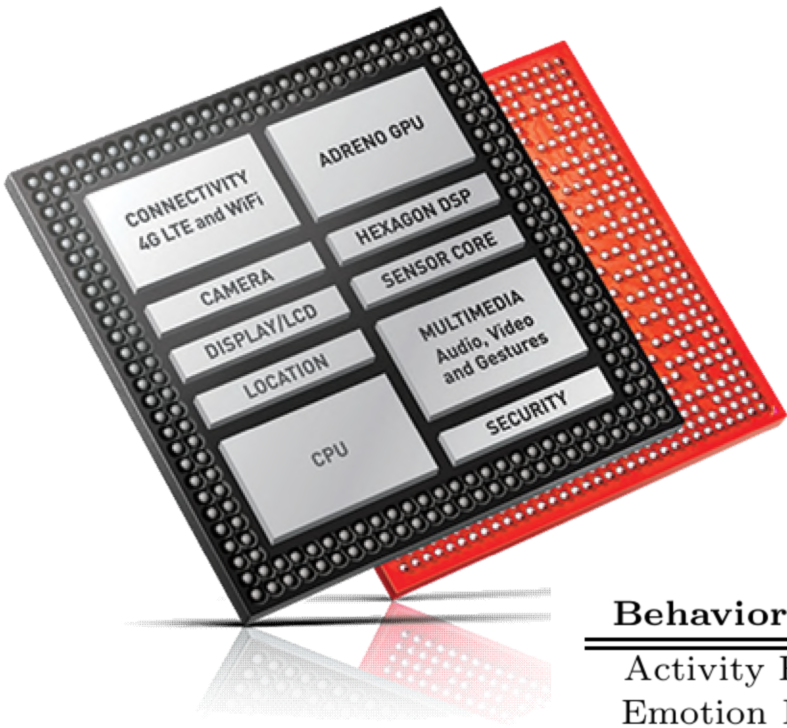

Behavioral Context	Dataset Description
Activity Recognition	wrist-worn accelerometer activities [10]
Emotion Recognition	emotional prosody speech [19]
Speaker Identification	10-minute speech from 23 speakers each

Platform and Datasets

- Snapdragon 800 SoC

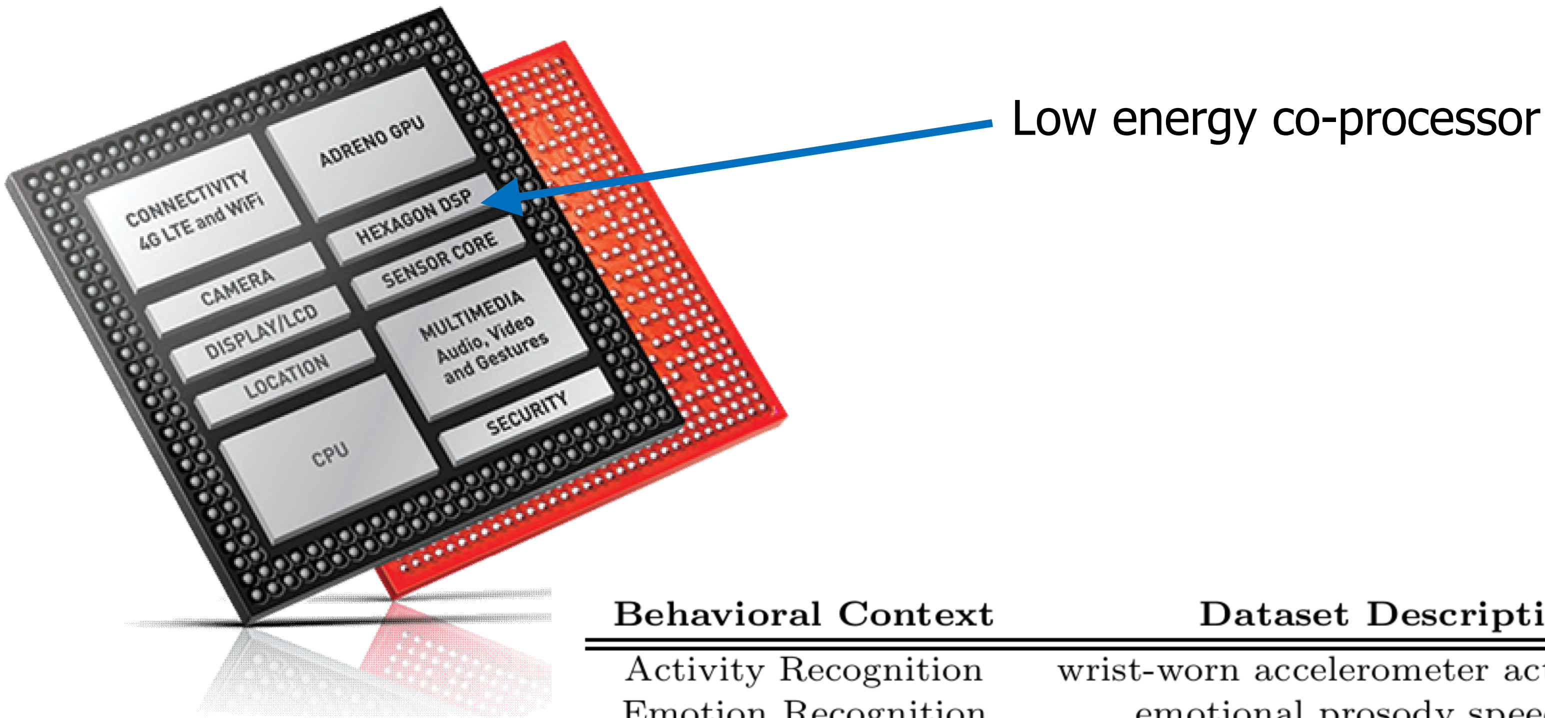

Behavioral Context	Dataset Description
Activity Recognition	wrist-worn accelerometer activities [10]
Emotion Recognition	emotional prosody speech [19]
Speaker Identification	10-minute speech from 23 speakers each

Question

- Can deep-learning-based classification be run on cell phones?

Deep Learning

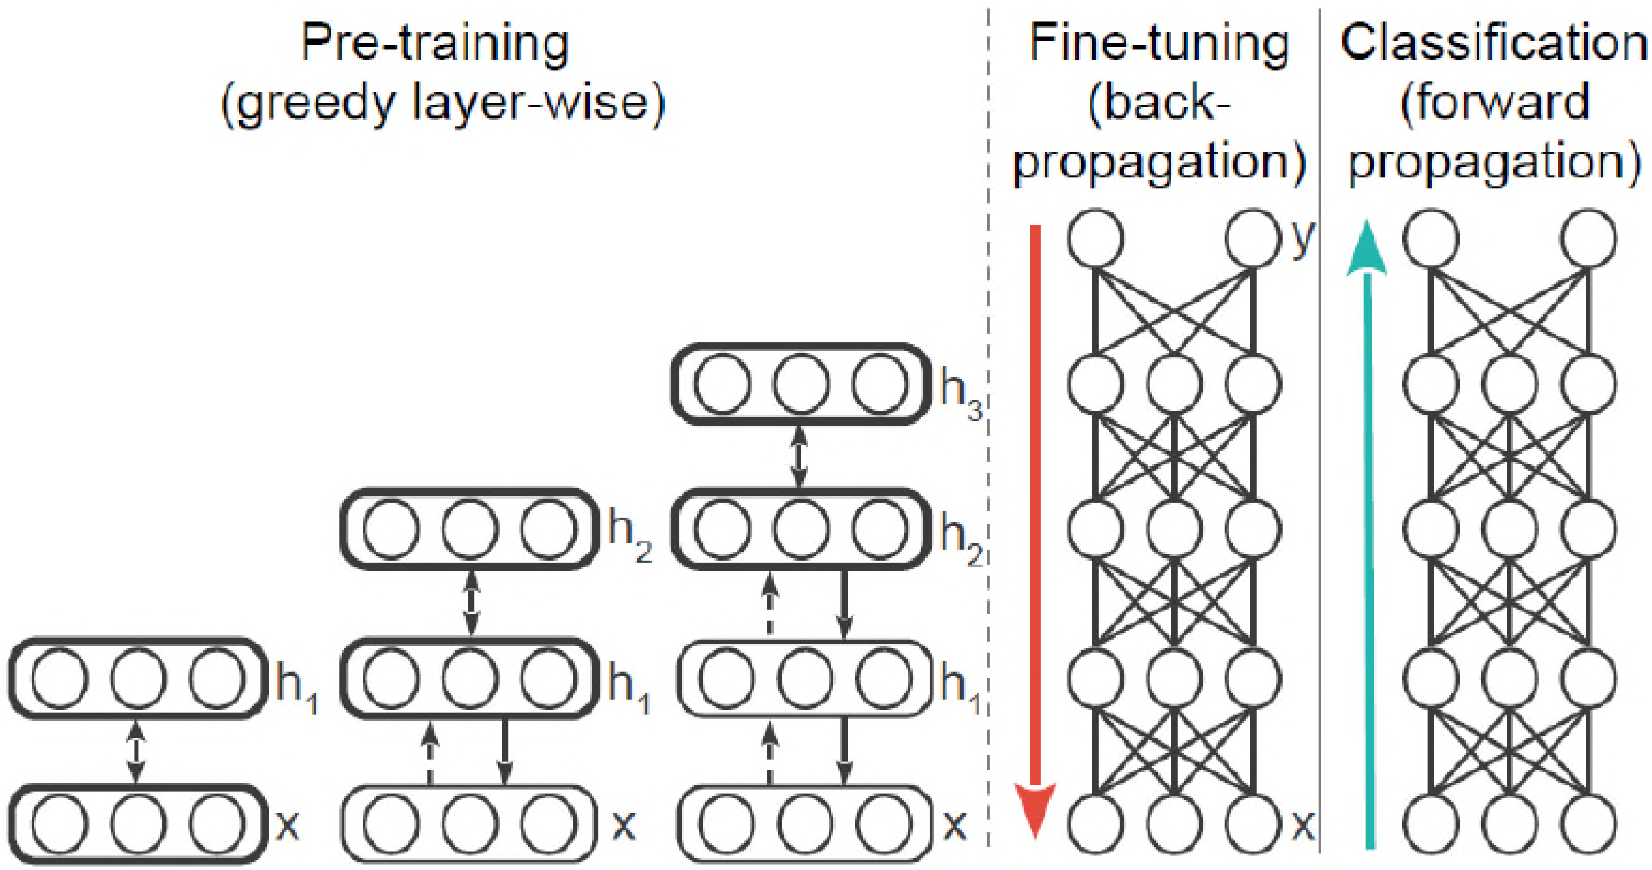

Nicholas D. Lane and Petko Georgiev, “Can Deep Learning Revolutionize Mobile Sensing?” ACM HotMobile, February 2015

Approaches: Energy and Latency Comparison

- DNN: Deep learning
- GMM: Gaussian mixture models
- DT: Decision trees

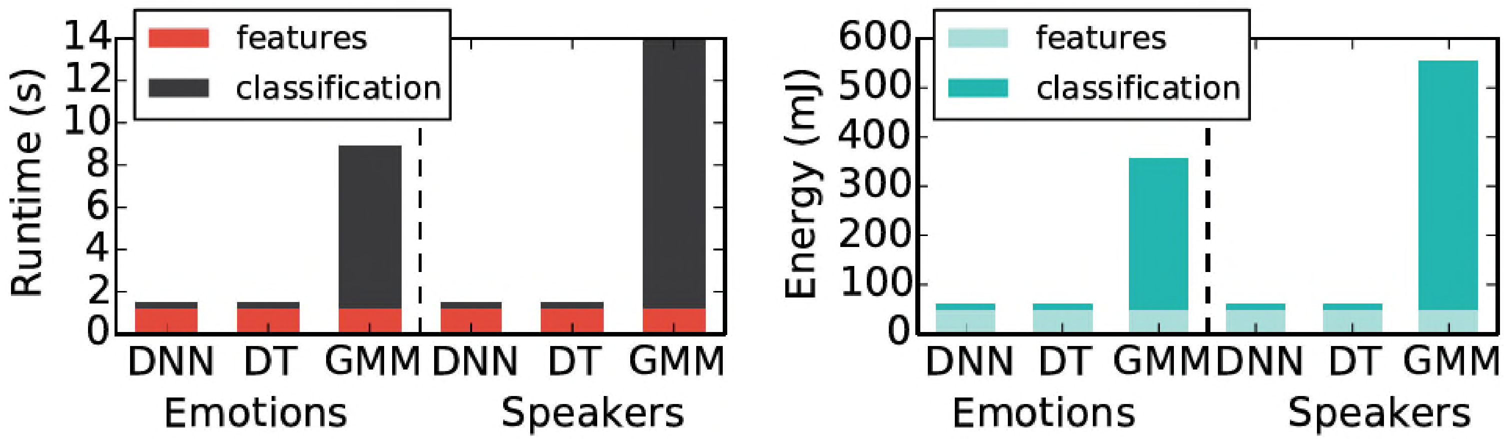

(a) Latency (b) Energy

Using the Microphone

Approaches: Energy and Latency Comparison

- DNN: Deep learning
- GMM: Gaussian mixture models
- DT: Decision trees

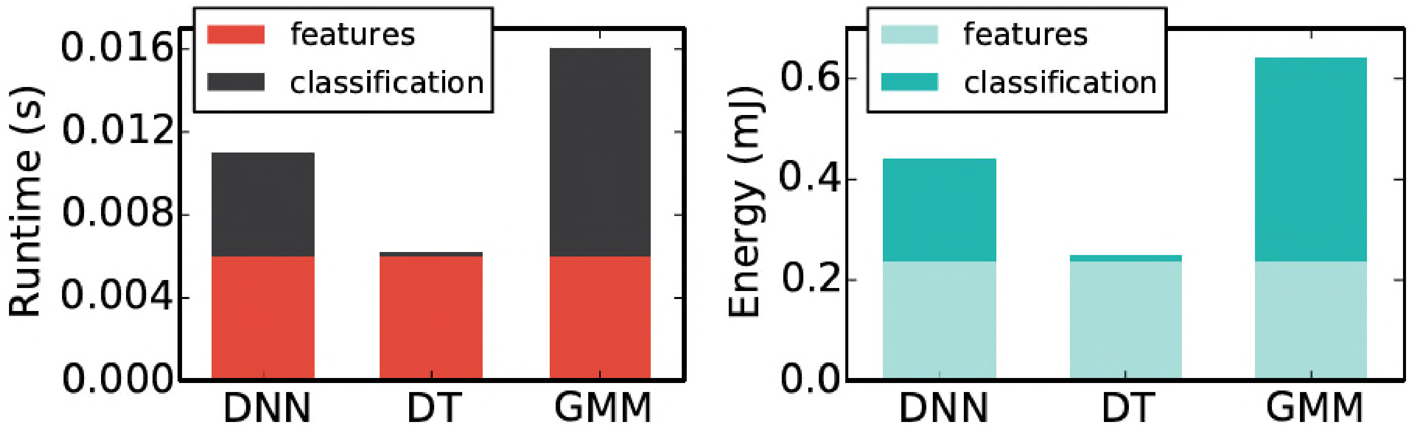

(a) Activity latency (b) Activity energy

Using the Accelerometer/Gyro

Approaches: Accuracy Comparison

- DNN: Deep learning
- GMM: Gaussian mixture models
- DT: Decision trees

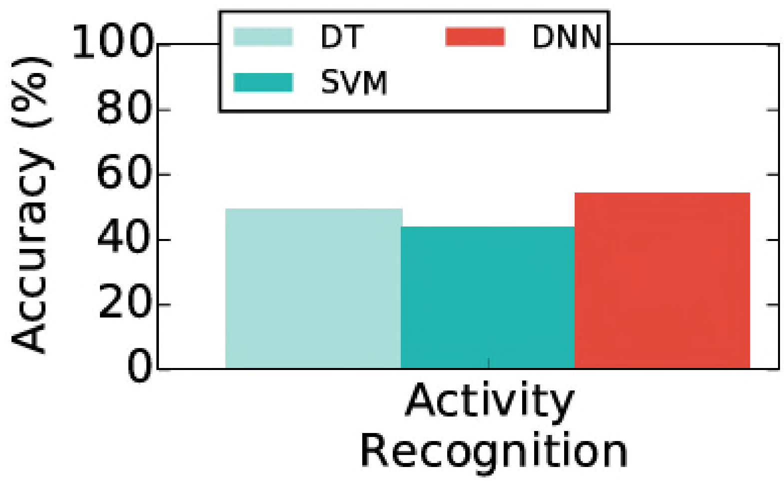

(a) Activity

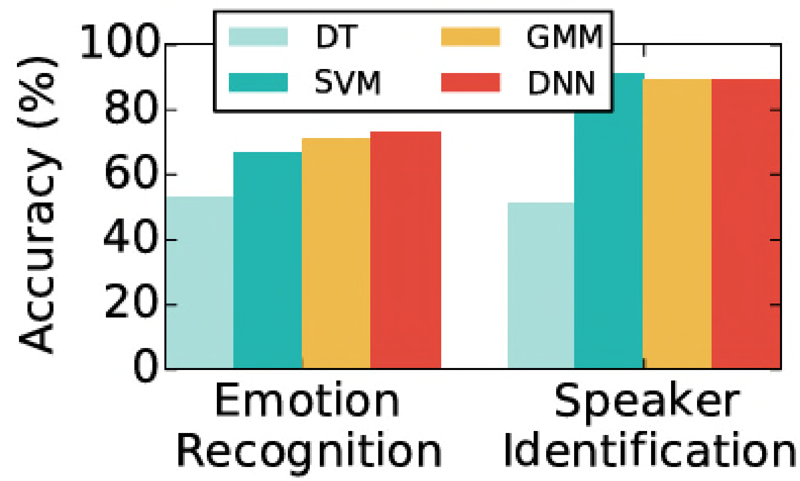

(b) Sound

Can Augmented Reality Applications Run on Cell Phones?

Can Augmented Reality Applications Run on Cell Phones?

- Prediction: Augmented reality market will grow from $500 million in 2016 to $5.7 billion in 2021
- Majority of revenue from subscription and license fee model
- Dominated by Tablets and Smart Glasses

http://www.forbes.com/sites/kevinanderton/2016/11/14/augmented-reality-the-future-and-pokemon-go-infographic/

Is the problem addressed in this paper already solved (e.g., see PokemonGo)?

- Why? Why not?
- Same? Different? How different?

Using the Camera: Augmented Reality

- Vision:
  - Annotate everyday objects using mobile phones.
  - Allow others to retrieve annotations previously associated with objects

Puneet Jain, Justin Manweiler, Romit Roy Choudhury, "OverLay: Practical Mobile Augmented Reality," ACM Mobisys, May 2015

Sensor-based: Localization

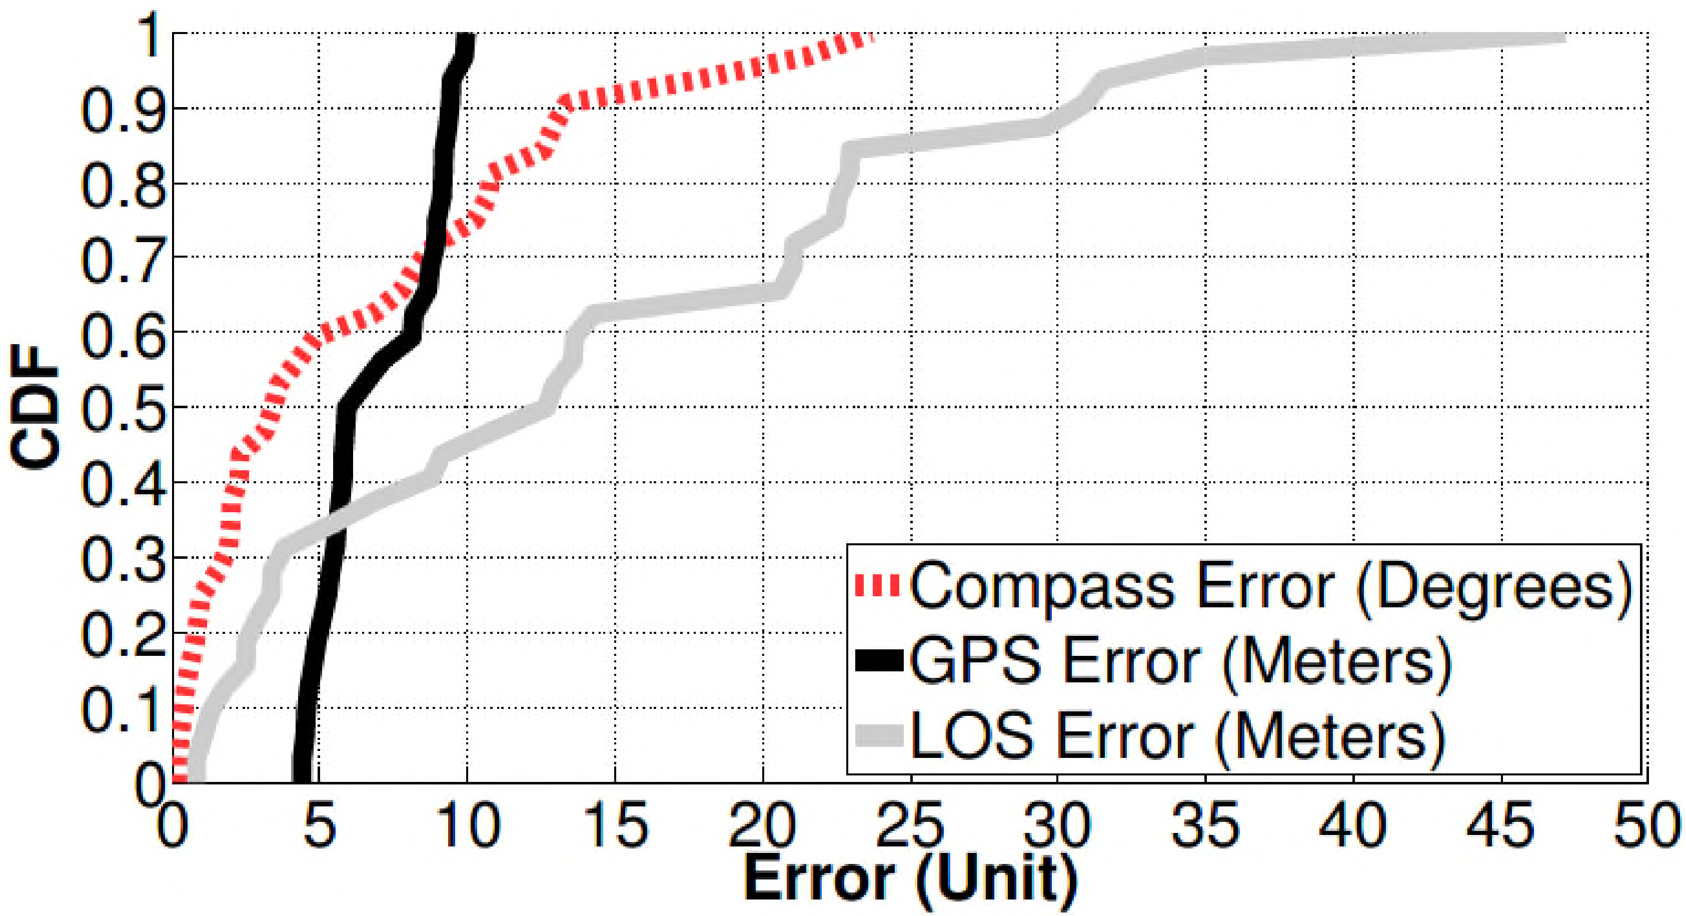

Vision-based: Processing Latency

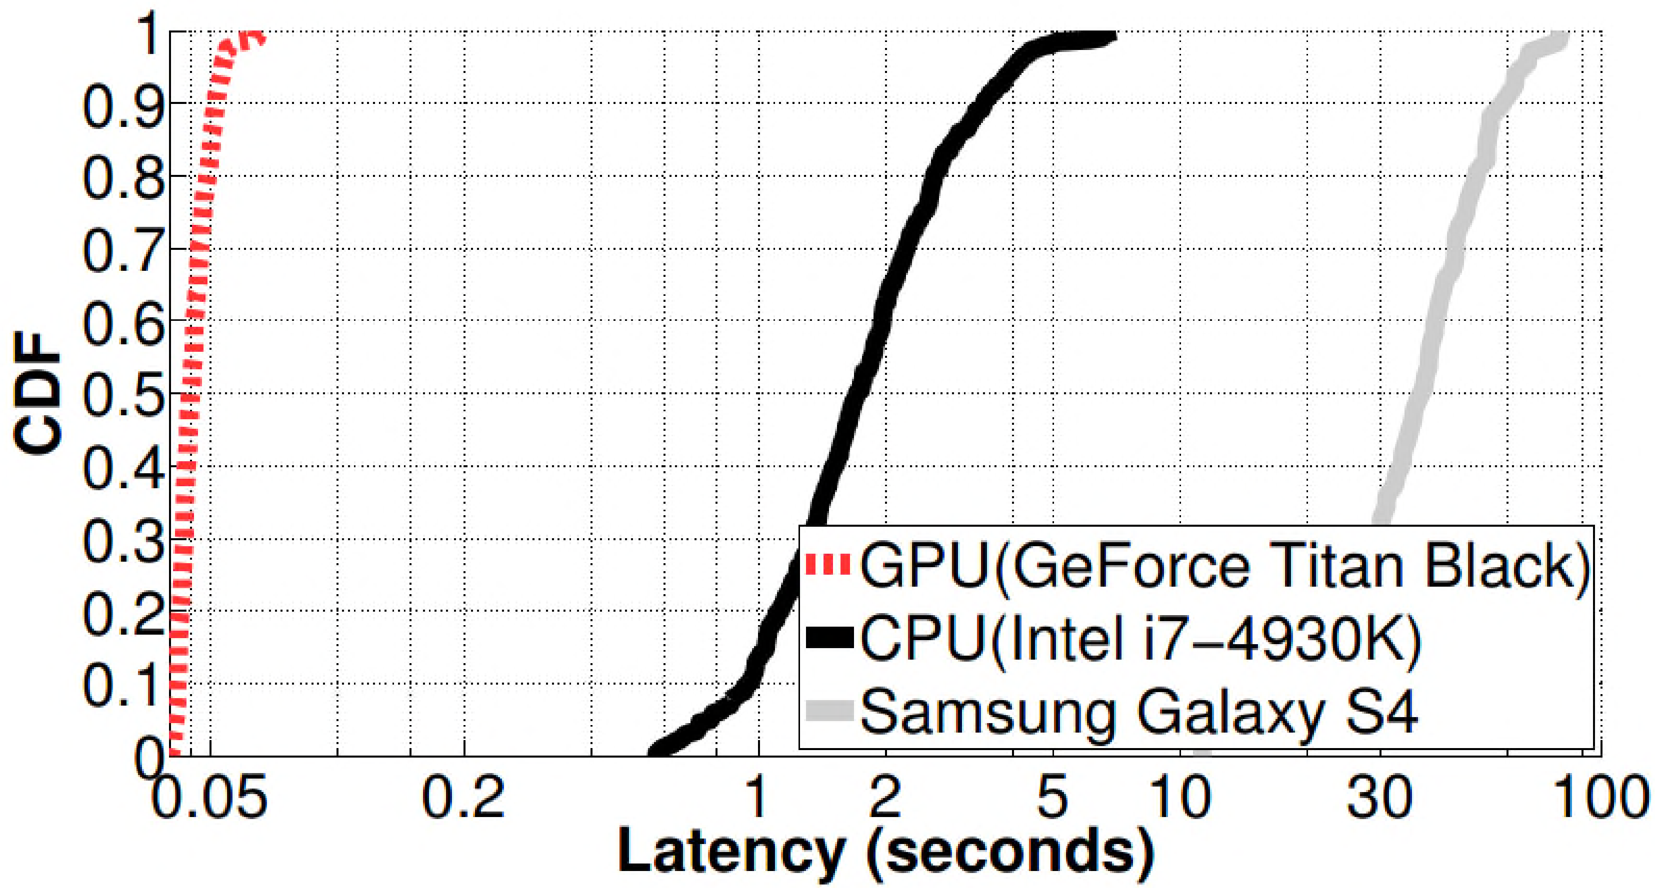

Vision-based: Upload Latency

- Latency uploading a 1920*1080 frame

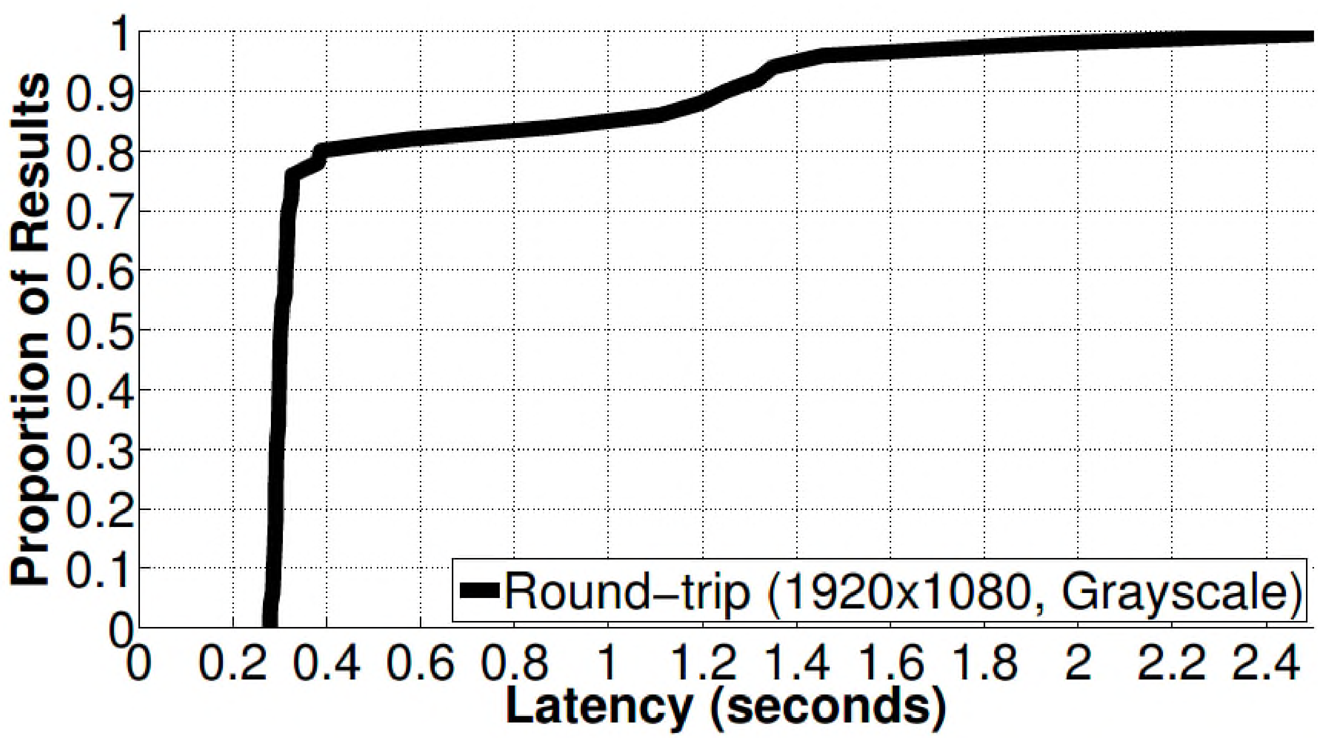

Vision-based: Accuracy and Latency of Object Matching
(Against 100 Objects)

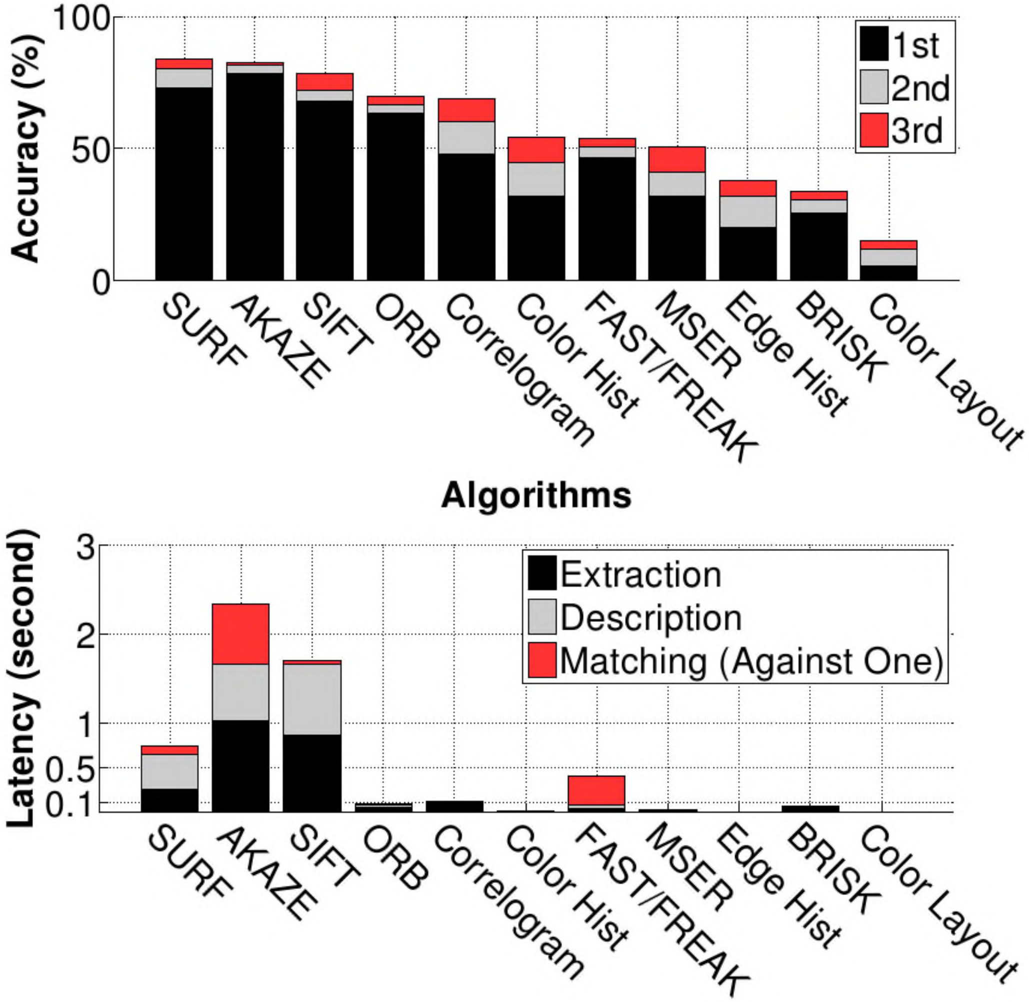

- Must minimize number of objects matched against

OverLay System

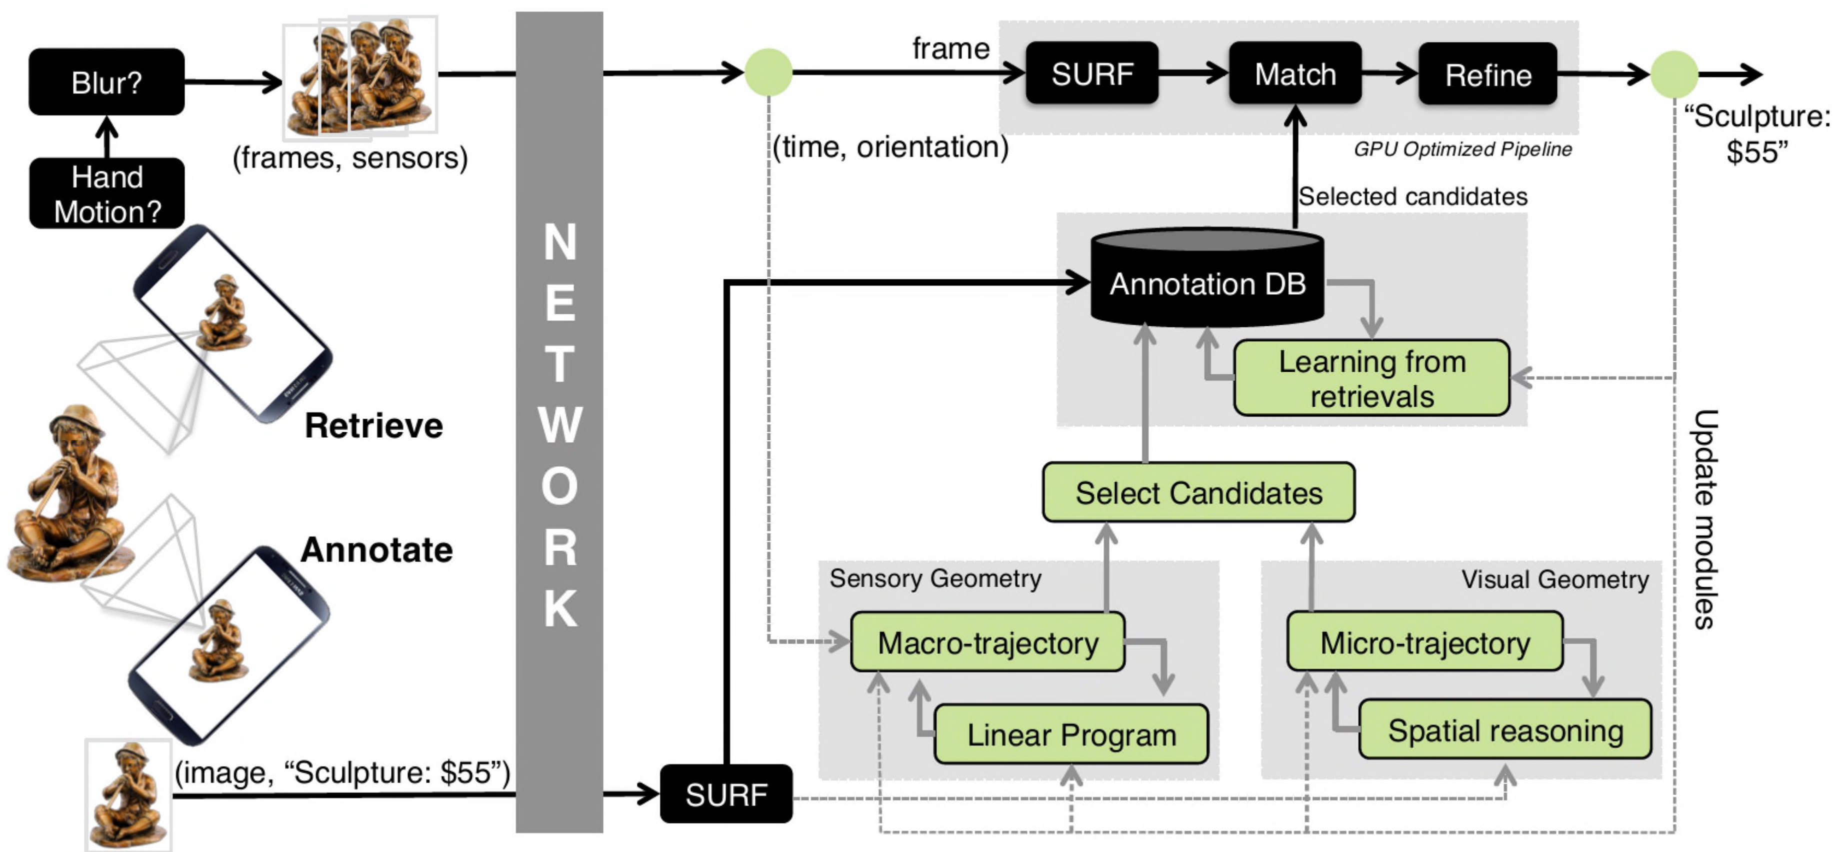

Main Idea: Geometric Neighborhoods

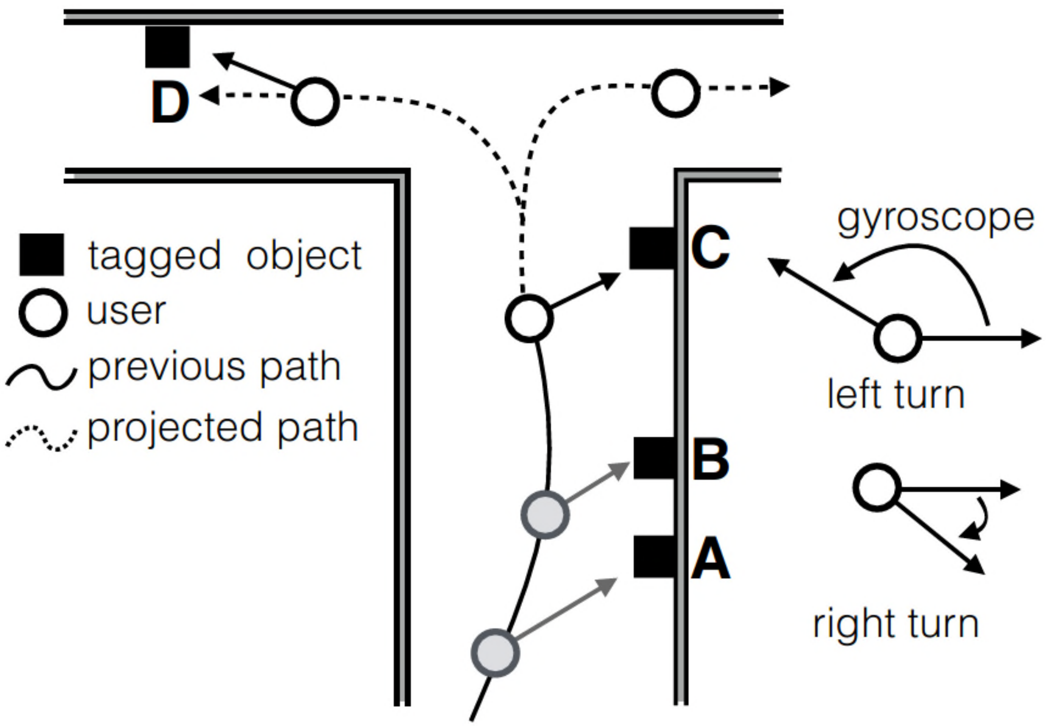

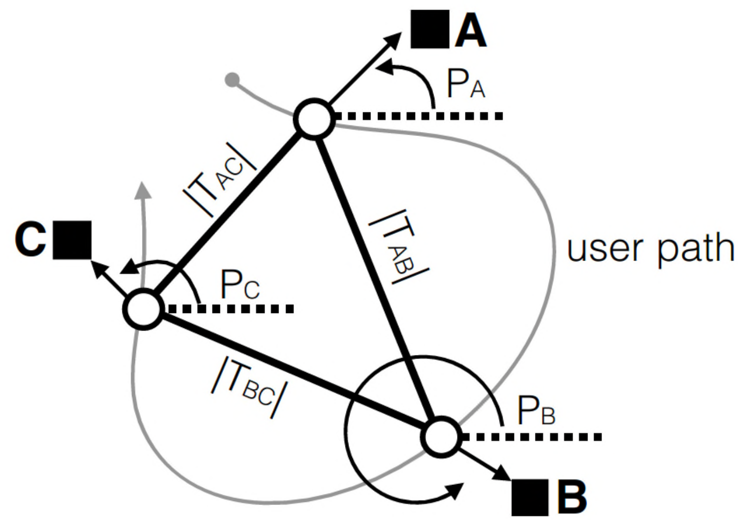

Question: How many images to keep of each object?

- Too few → accuracy problems
- Too many → latency problems
- What is the best way to reduce number of views without reducing accuracy?

Evaluation

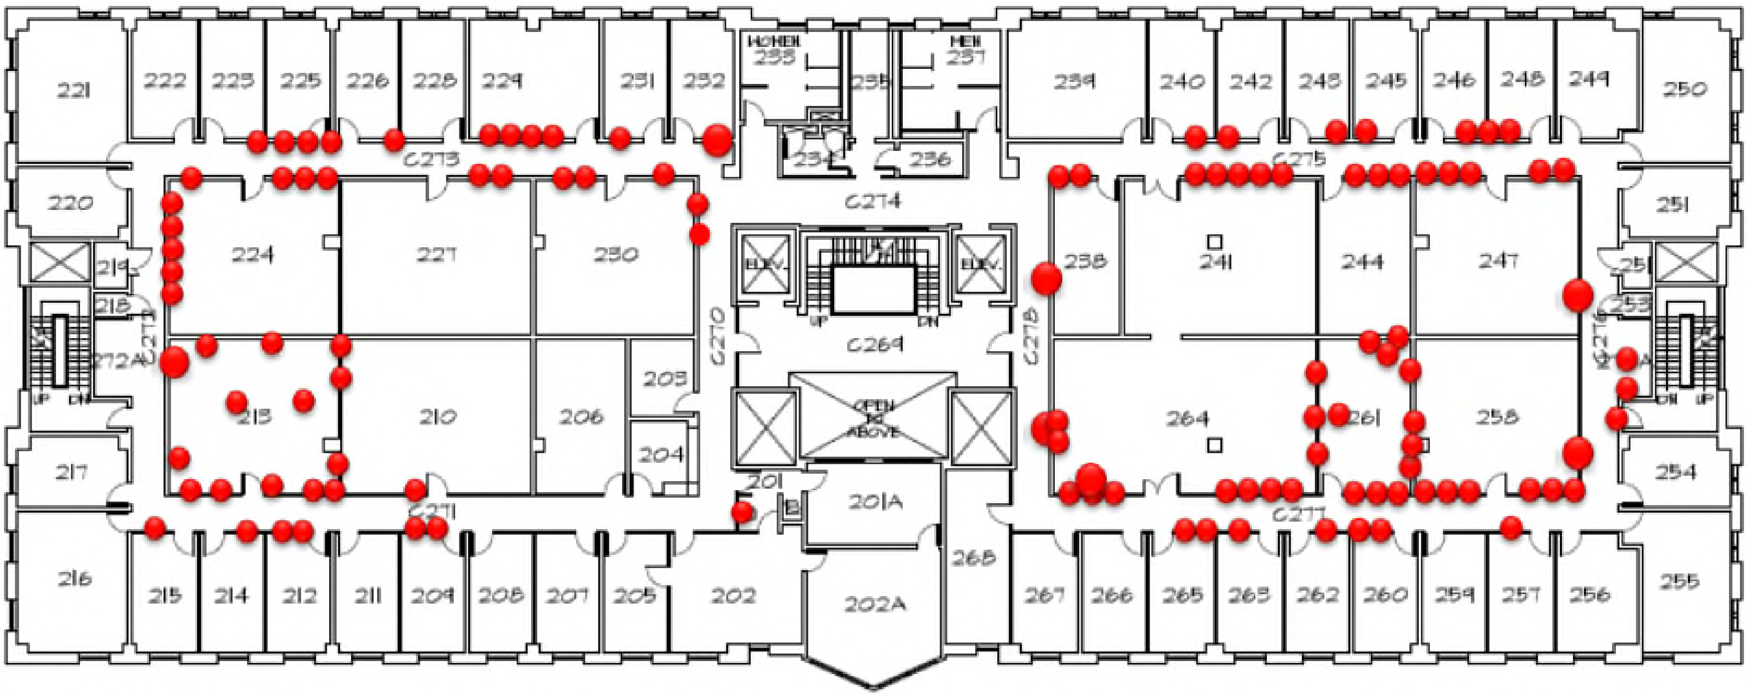

221
222
224
226
228
250
220
251
224
227
241
244
247
216
210
206
204
264
258
217
254
201
201A
268
202
202A
255
WOMEN
MEN
ELEV.
ELEV.
OPEN TO ABOVE
UP
DN

Evaluation

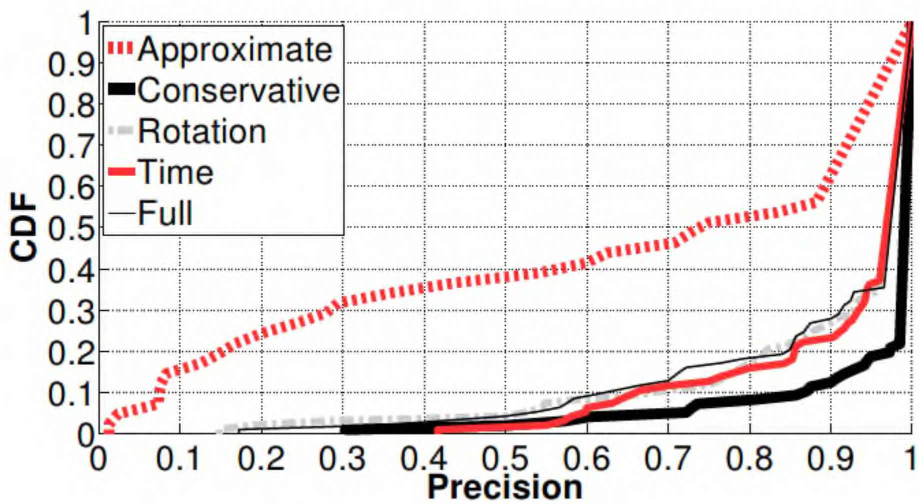

Approximate
Conservative
Rotation
Time
Full
CDF
Precision
0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1

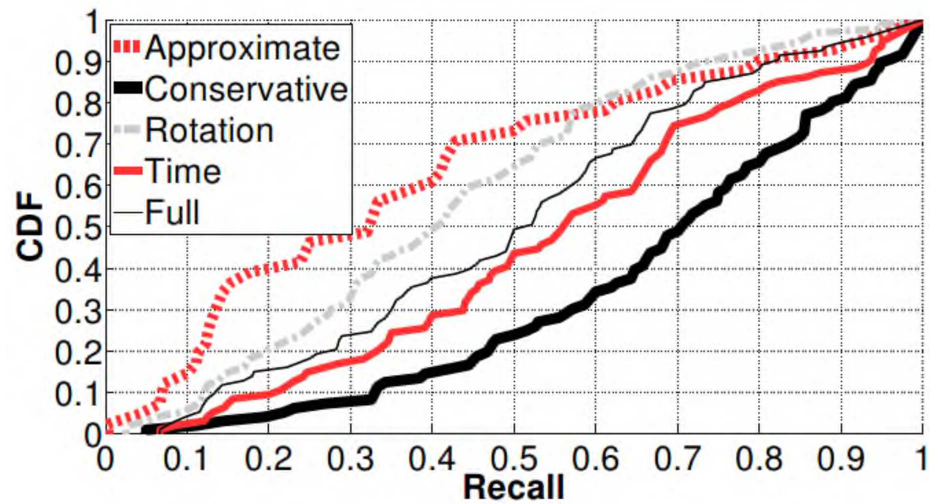

Approximate
Conservative
Rotation
Time
Full
CDF
Recall
0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1

Evaluation

- Optimizations reduce the amount of matching that needs to happen

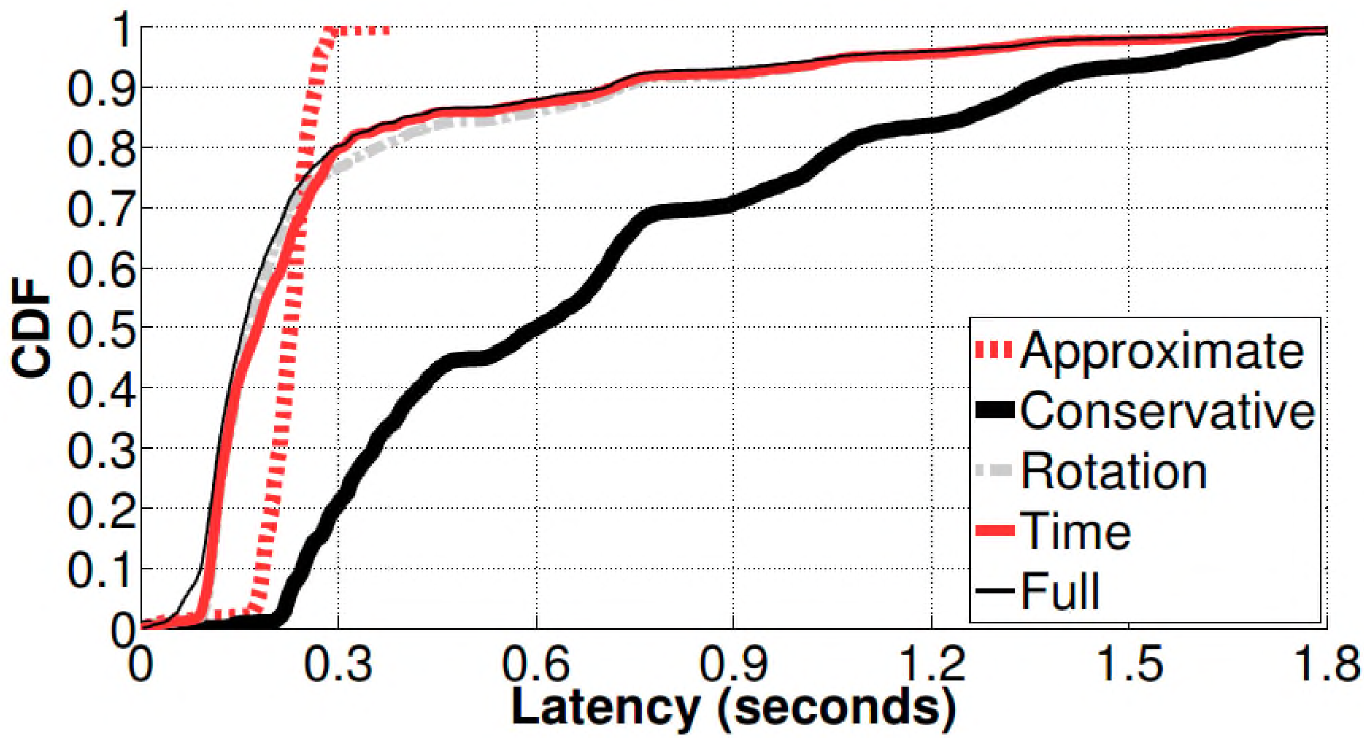

Ultra Power Eye-Gaze Tracking

- What was the main idea in reducing power of eye-gaze trackers?
  - Do (what) instead of (what)?
- What main challenges did the authors list in implementing the above low-power gaze-tracking idea?
  - How did they argue that this was a difficult problem?

Ultra Power Eye-Gaze Tracking

- Photo-diodes measure reflected light

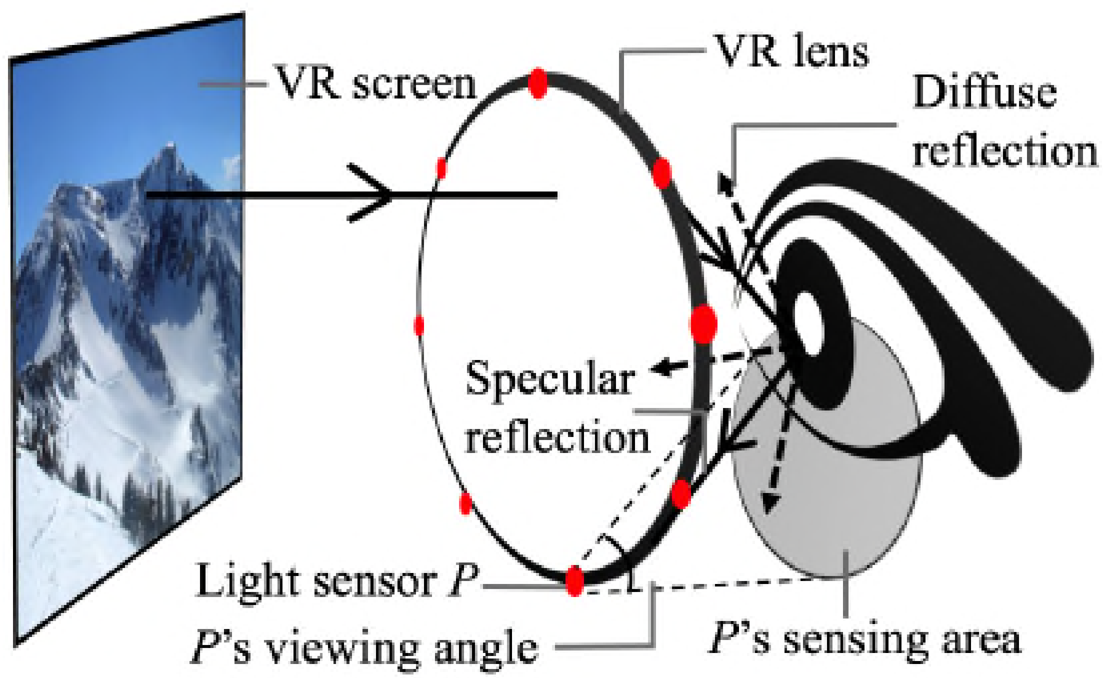

Ultra Power Eye-Gaze Tracking

- Reflected light is non-uniform and depends on scene

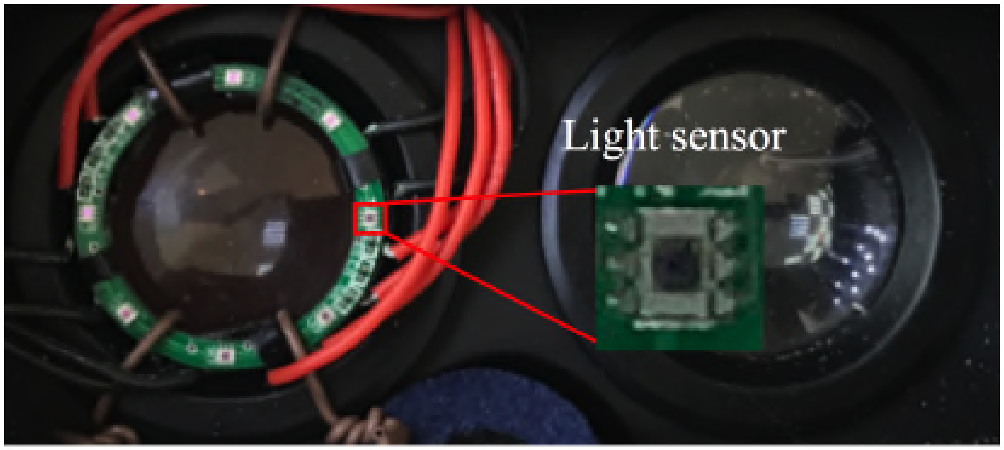

(a) A ring-shaped PCB on a VR lens

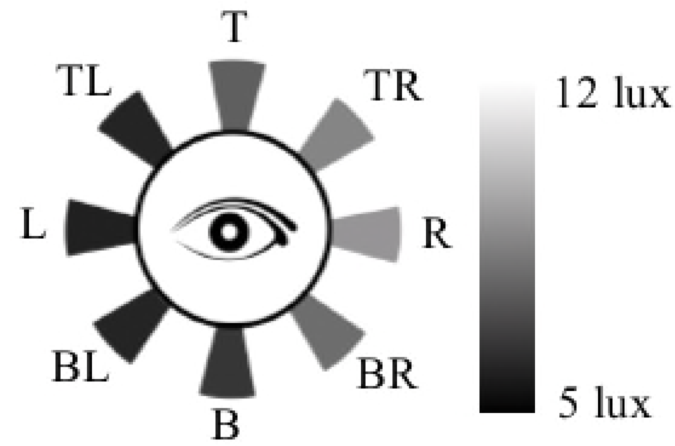

(b) Reflected light w/ center pupil

Ultra Power Eye-Gaze Tracking

- Reflected light also depends on pupil position (why?)

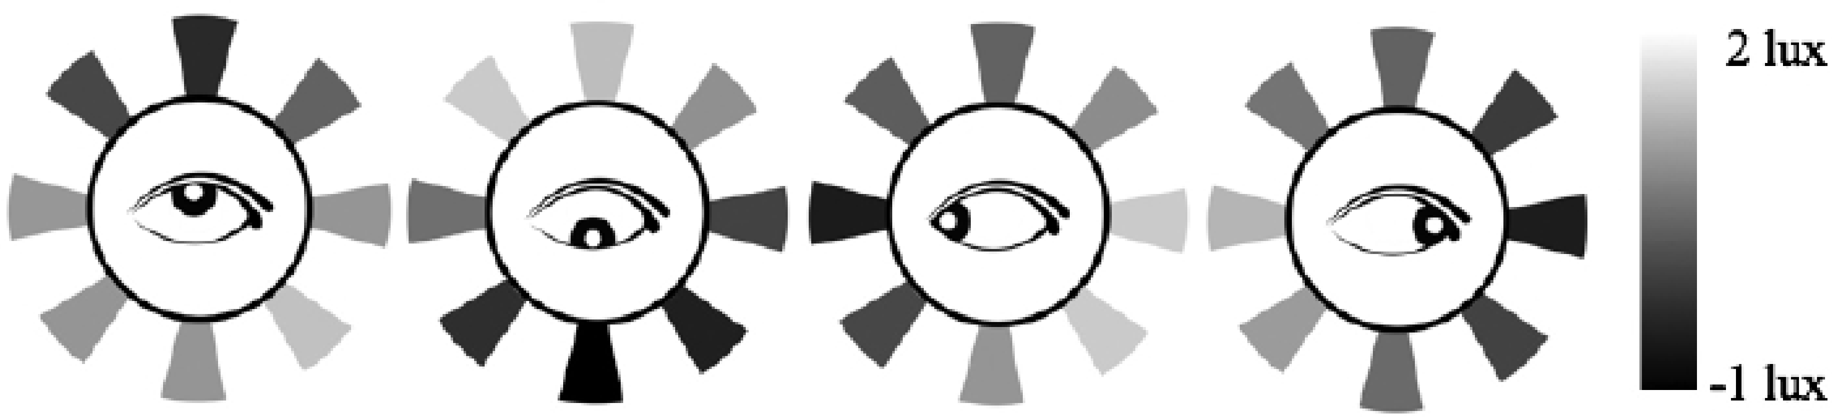

Ultra Power Eye-Gaze Tracking

- Reflected light also depends on user

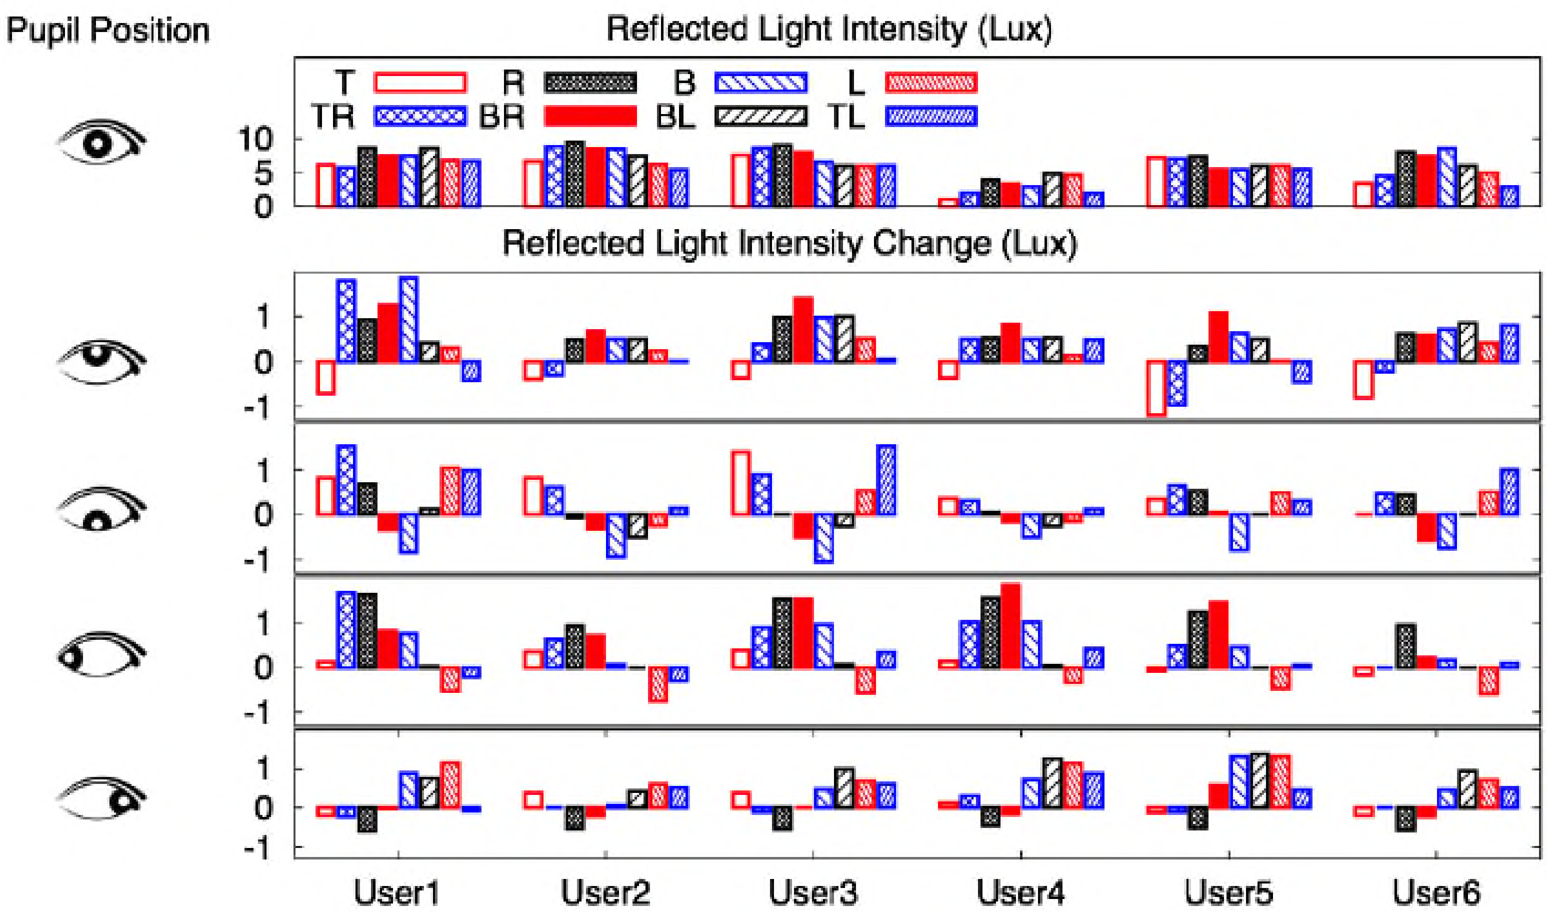

General Idea

- Understand correlations between front and back facing sensors for different eye positions and use it to train a boosted decision tree classifier to estimate eye position vector

Examples

- Correlations with eye staring straight ahead (pupil in the middle)

Back sensor (facing eye)	Front sensor (facing display): T	TR	R	BR	B	BL	L	TL
T	0.48	0.75	0.87	0.87	$\boxed{0.96}$	0.81	0.59	0.54
TR	0.69	0.55	0.76	0.76	0.79	$\boxed{0.91}$	0.88	0.78
R	0.66	0.64	0.55	0.67	0.79	0.83	$\boxed{0.93}$	0.89
BR	$\boxed{0.86}$	0.79	0.62	0.51	0.64	0.74	0.76	0.84
B	$\boxed{0.97}$	0.83	0.68	0.55	0.50	0.51	0.78	0.79
BL	0.84	$\boxed{0.90}$	0.89	0.86	0.85	0.61	0.60	0.58
L	0.80	0.91	$\boxed{0.92}$	0.78	0.65	0.49	0.52	0.58
TL	0.66	0.88	0.89	$\boxed{0.91}$	0.86	0.78	0.52	0.47

Blink Detection

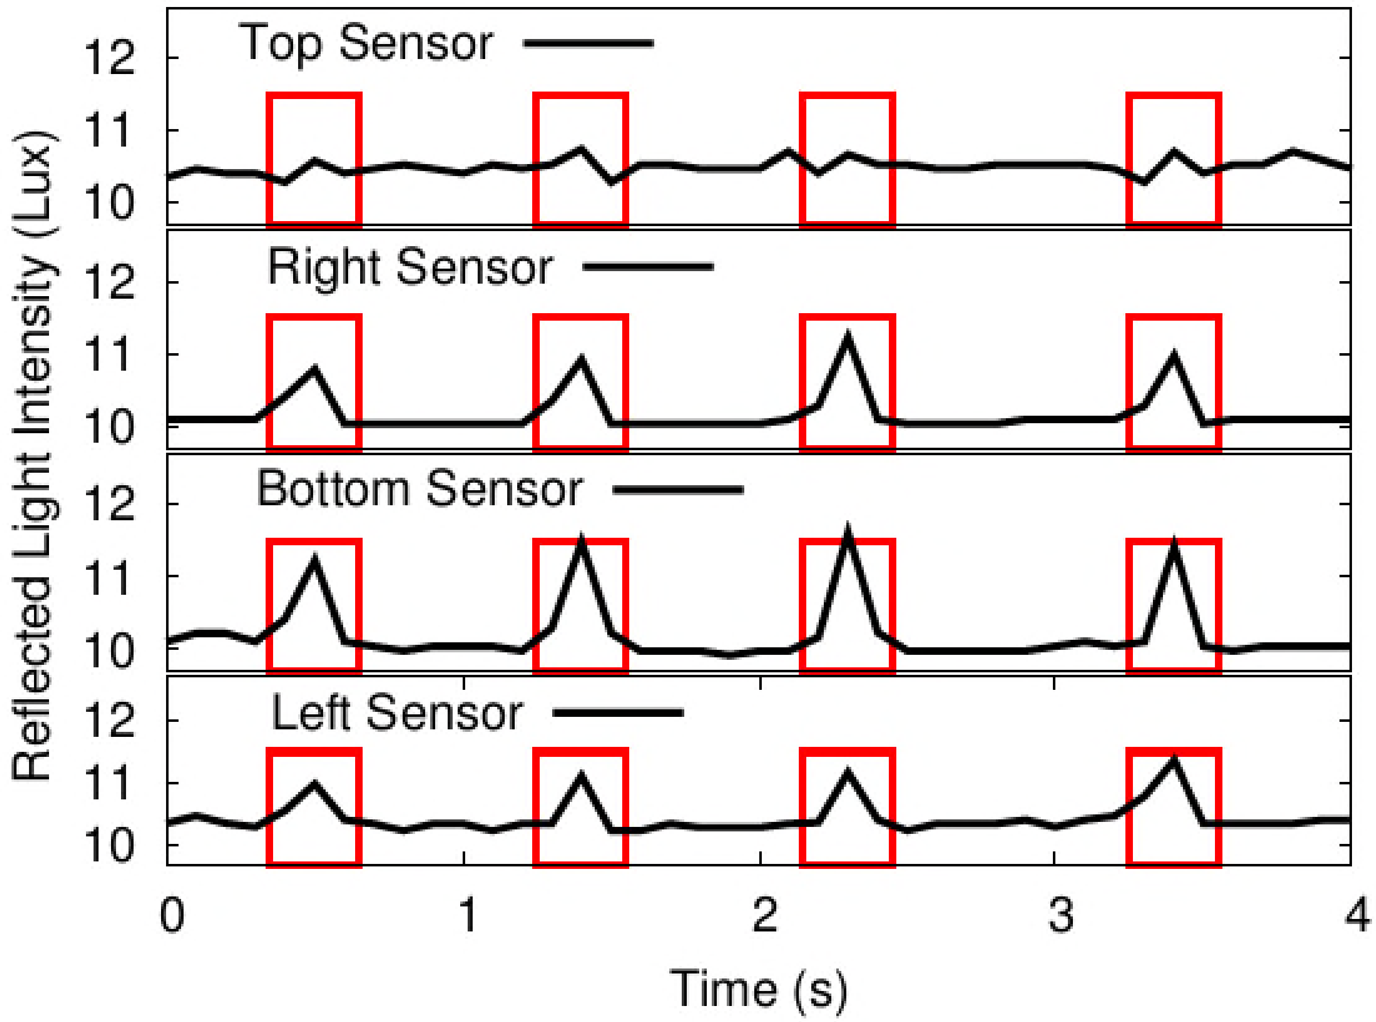

Evaluation: Tracking Error

- Smaller error when training and testing with the same user