SIGCOMM 2018 Topic Preview: Applications and IoT

Hamed Haddadi
Imperial College London

Imperial College London

Session papers

11:00 am - 12:40 am Main-Conference Session 5: Applications and IoT

Session Chair: Mohammad Alizadeh *(MIT, USA)*
Location: Vigadó, 2nd-Floor Ceremonial Hall

11:00 am - 11:25 am	**AWStream: Adaptive Wide-Area Streaming Analytics** Ben Zhang *(Berkeley, USA)*, Xin Jin *(JHU, USA)*, Sylvia Ratnasamy, John Wawrzynek, Edward A. Lee *(Berkeley, USA)*	
11:25 am - 11:50 am	**Chameleon: Video Analytics at Scale via Adaptive Configurations and Cross-Camera Correlations** Junchen Jiang *(Microsoft and UChicago, USA)*, Ganesh Ananthanarayanan, Peter Bodik, Siddhartha Sen *(Microsoft, USA)*, Ion Stoica *(Berkeley, USA)*	
11:50 am - 12:15 pm	**RF-Based 3D Skeletons** Mingmin Zhao, Dina Katabi, Yonglong Tian, Hang Zhao, Mohammad Abu Alsheikh, Tianhong Li, Antonio Torralba, Rumen Hristov *(MIT, USA)*	
12:15 pm - 12:40 pm	**MUTE: Bringing IoT to Noise Cancellation** Sheng Shen, Nirupam Roy, Junfeng Guan, Haitham Hassanieh, Romit Roy Choudhury *(UIUC, USA)*	

Deep Neural Nets

coz' supervised learning is hard..

How the brain works.

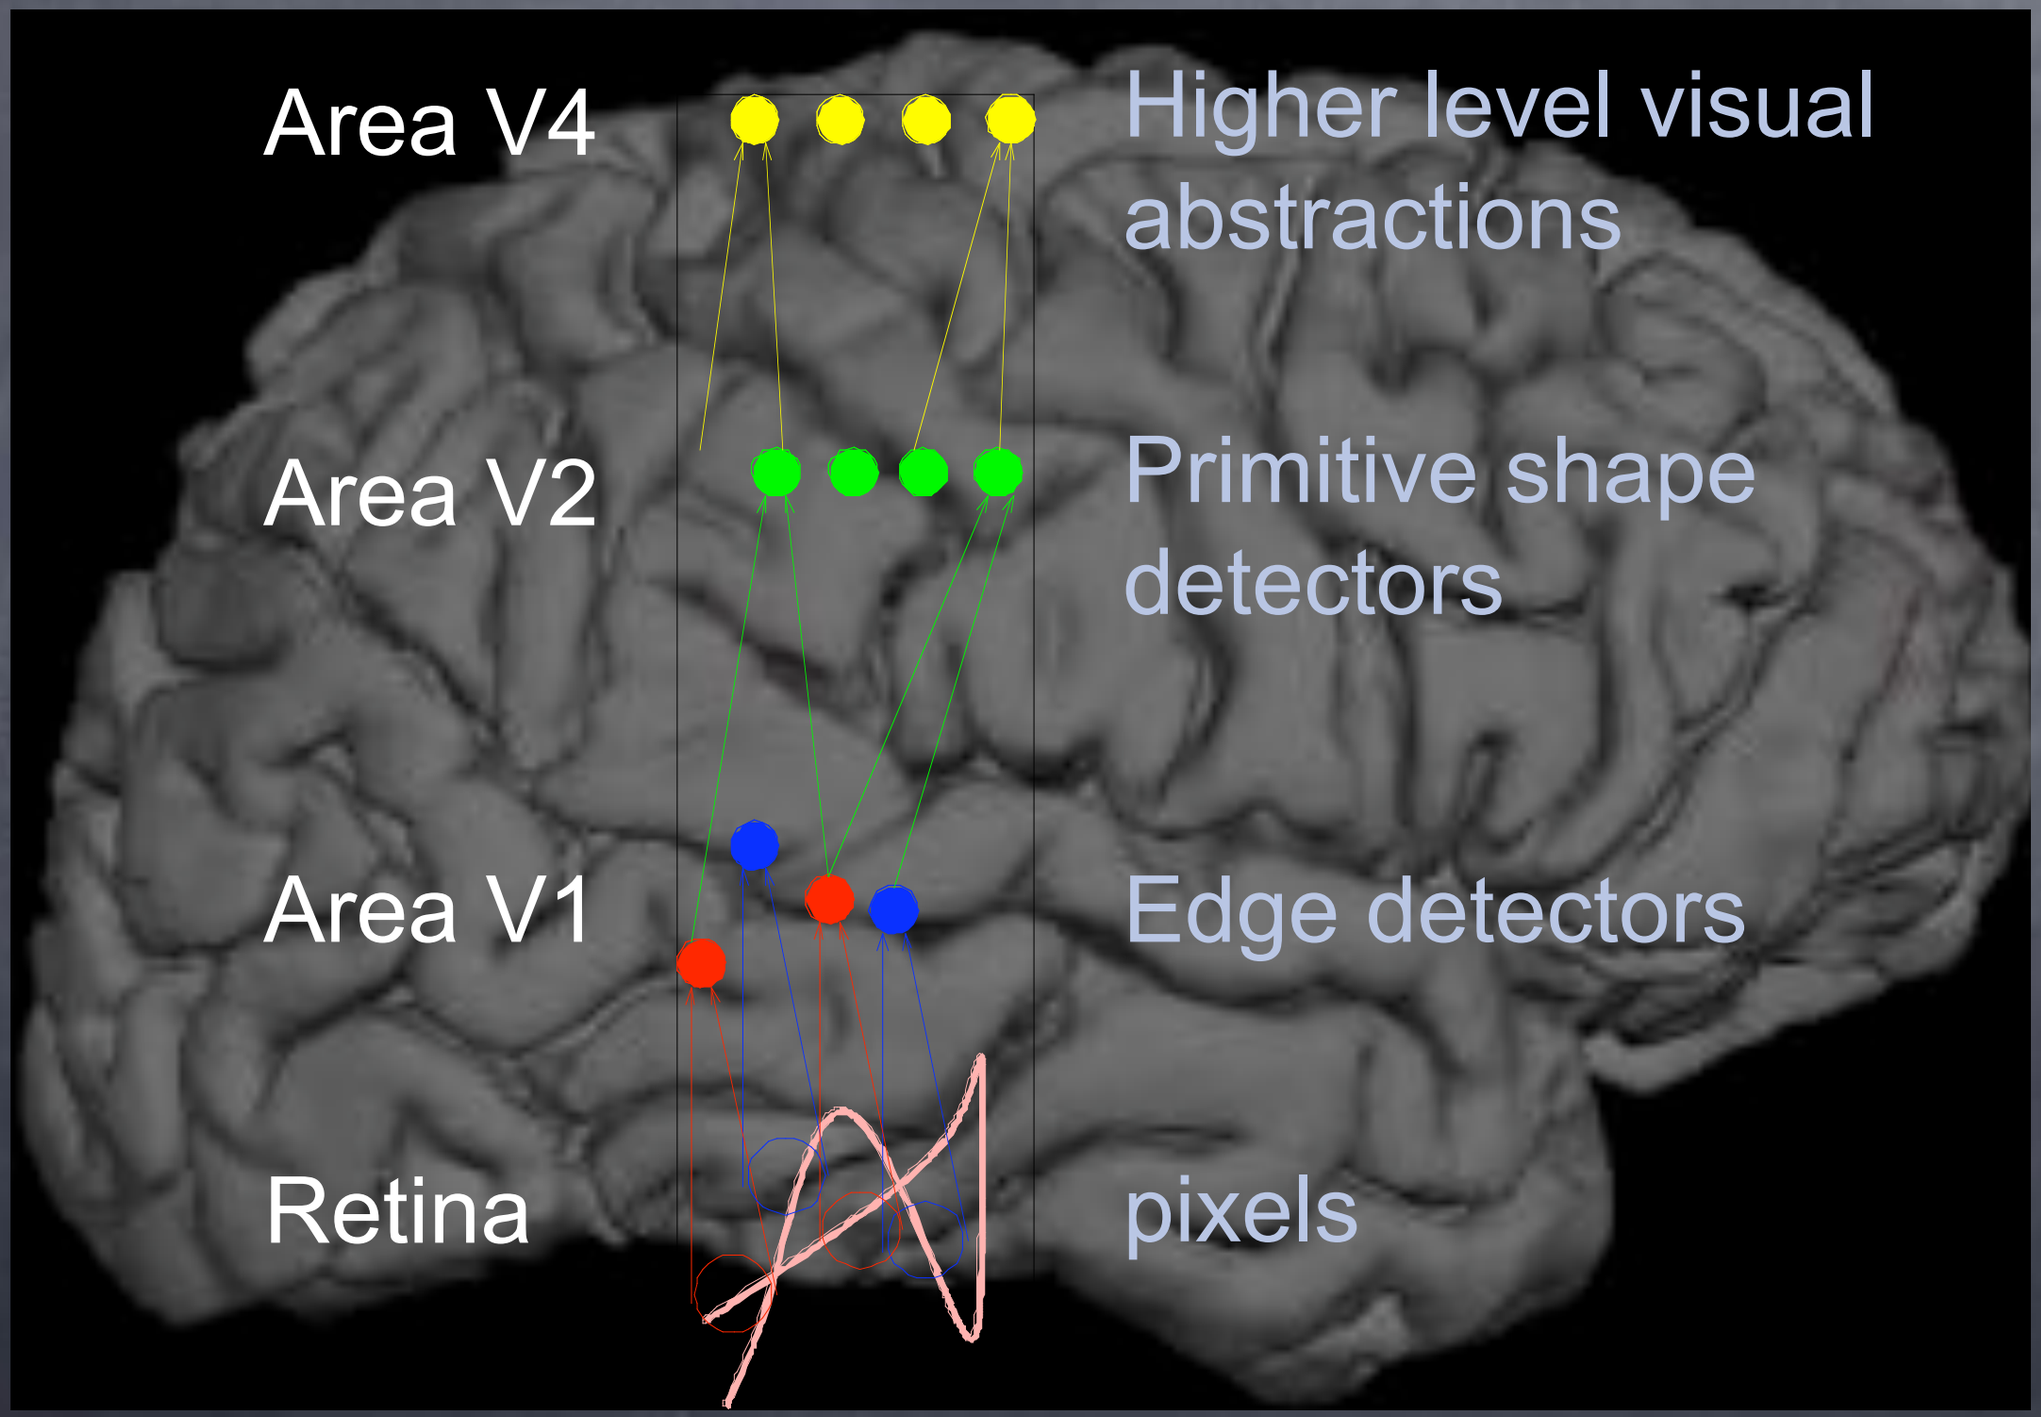

The brain activation

- Many neurons active simultaneously in the brain: around 1%
- The input is represented by the activation of a set of features that are not mutually exclusive.
- Can be exponentially more efficient than local representations

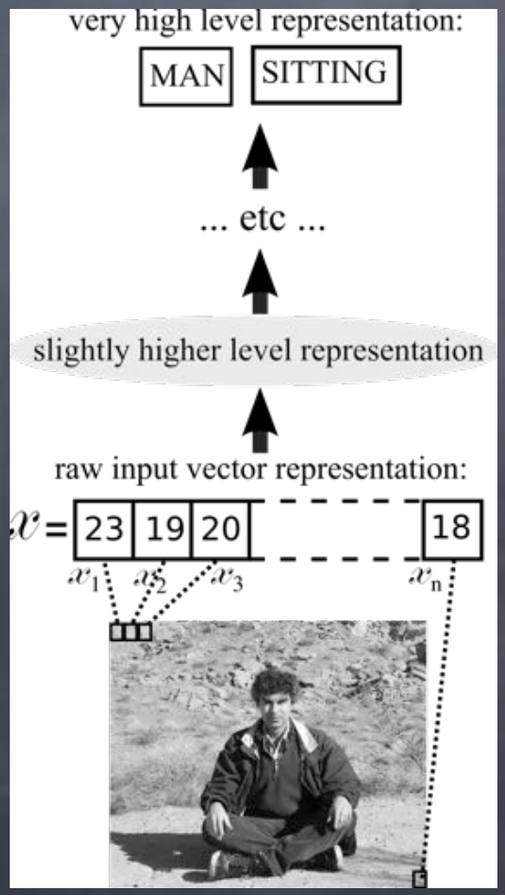

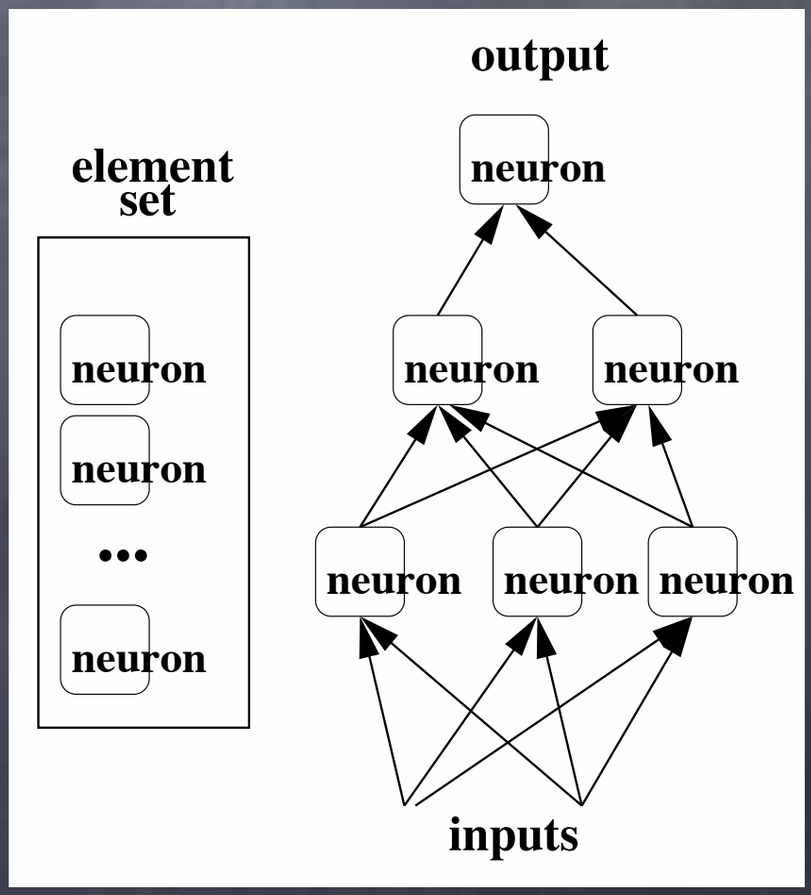

Training Neural Nets

- Unsupervised Learning for Deep Architectures
  - Pretty hard to train large RBMs, DBNs, etc
- **Convolutional Neural Networks**
  - **Easier to train!** (few inputs per neuron) helps gradients to propagate through so many layers (**backpropagation**)

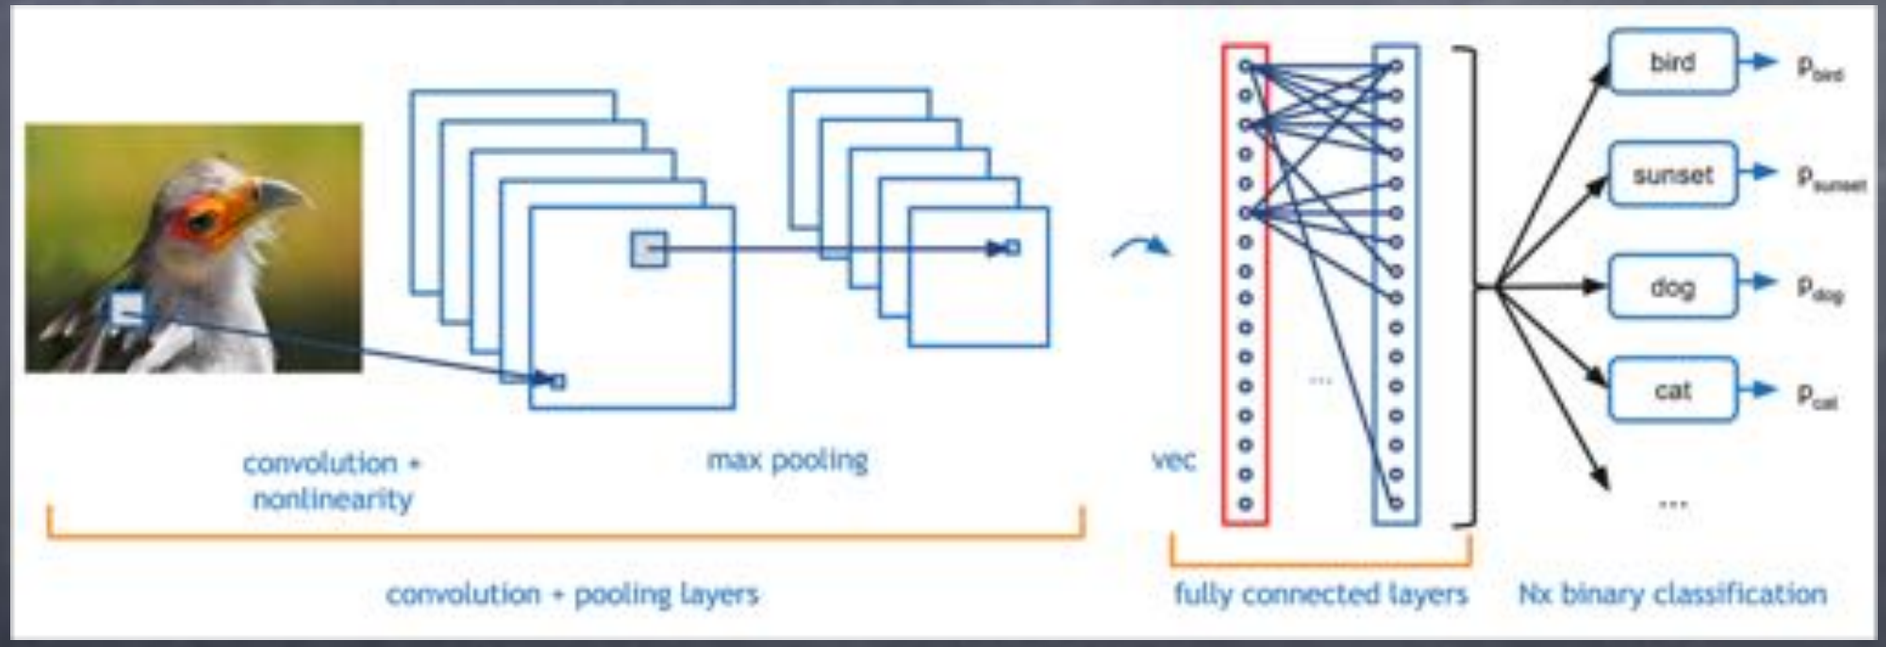

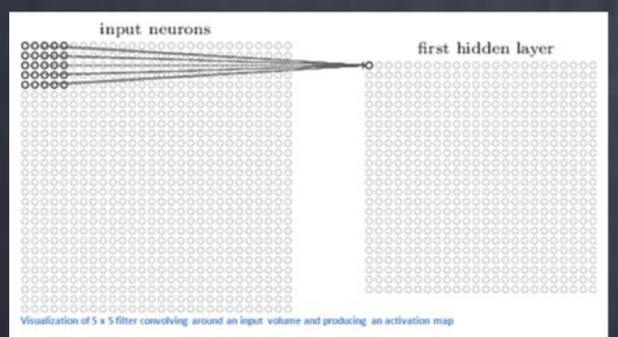

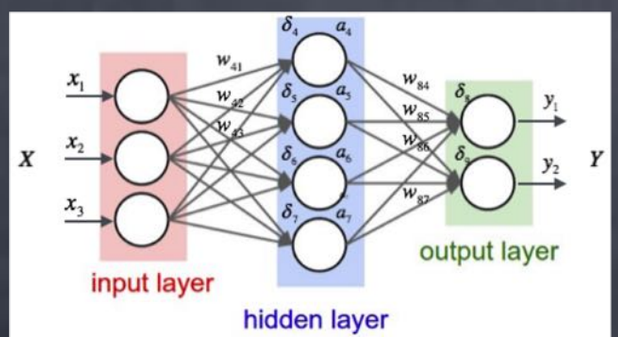

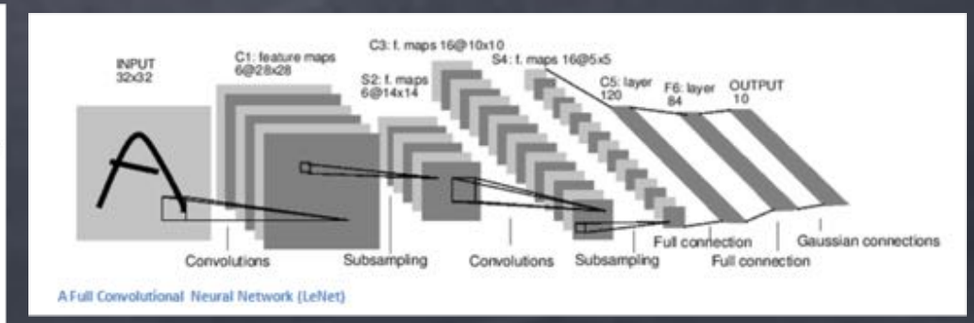

Success metrics

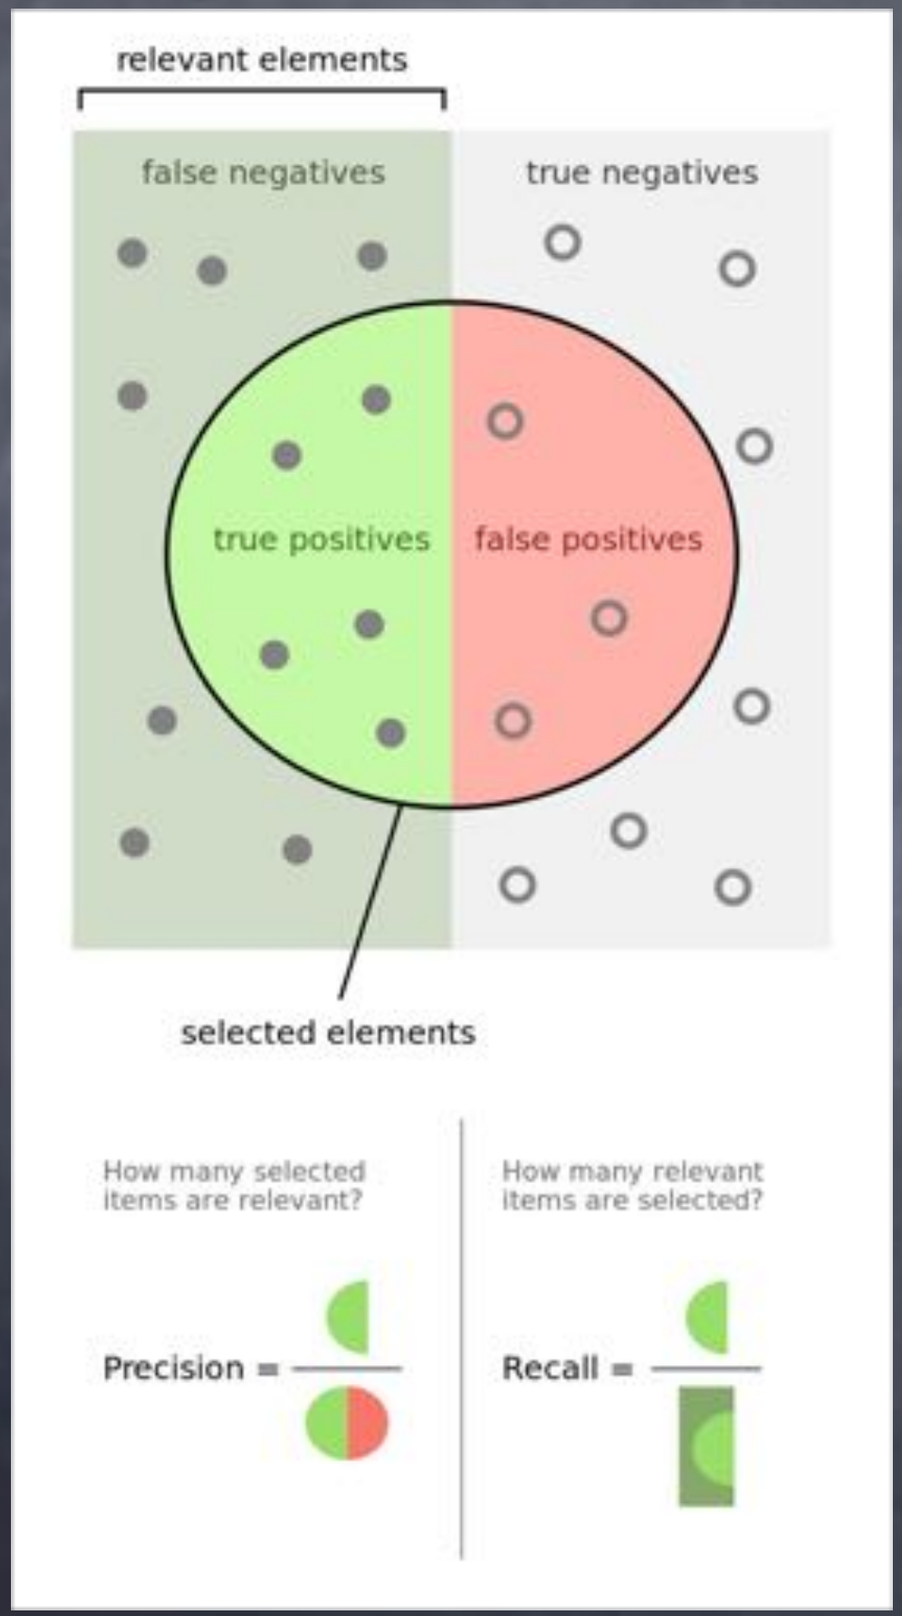

$$F_1 = \frac{2}{\frac{1}{\text{recall}} + \frac{1}{\text{precision}}} = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

Success metrics

- Classification Accuracy. (the number of correct predictions made as a ratio of all predictions made.)
- Logarithmic Loss. (a performance metric for evaluating the predictions of probabilities of membership to a given class)
- Area Under ROC Curve. (model's ability to discriminate between positive and negative classes)
- Confusion Matrix. (a table with predictions on the x-axis and accuracy outcomes on the y-axis. The cells of the table are the number of predictions made by a machine learning algorithm.)

IoT

It's *thing*, and it's connected to the *internet!*

It was just a dumb thing. Then we put a chip in it.
Now it's a ***smart*** thing.

weputachipinit.tumblr.com

Data around us

- IoT devices
- Cyber Physical Systems

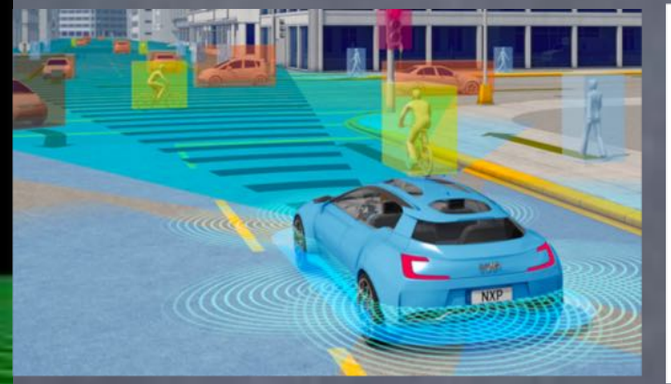

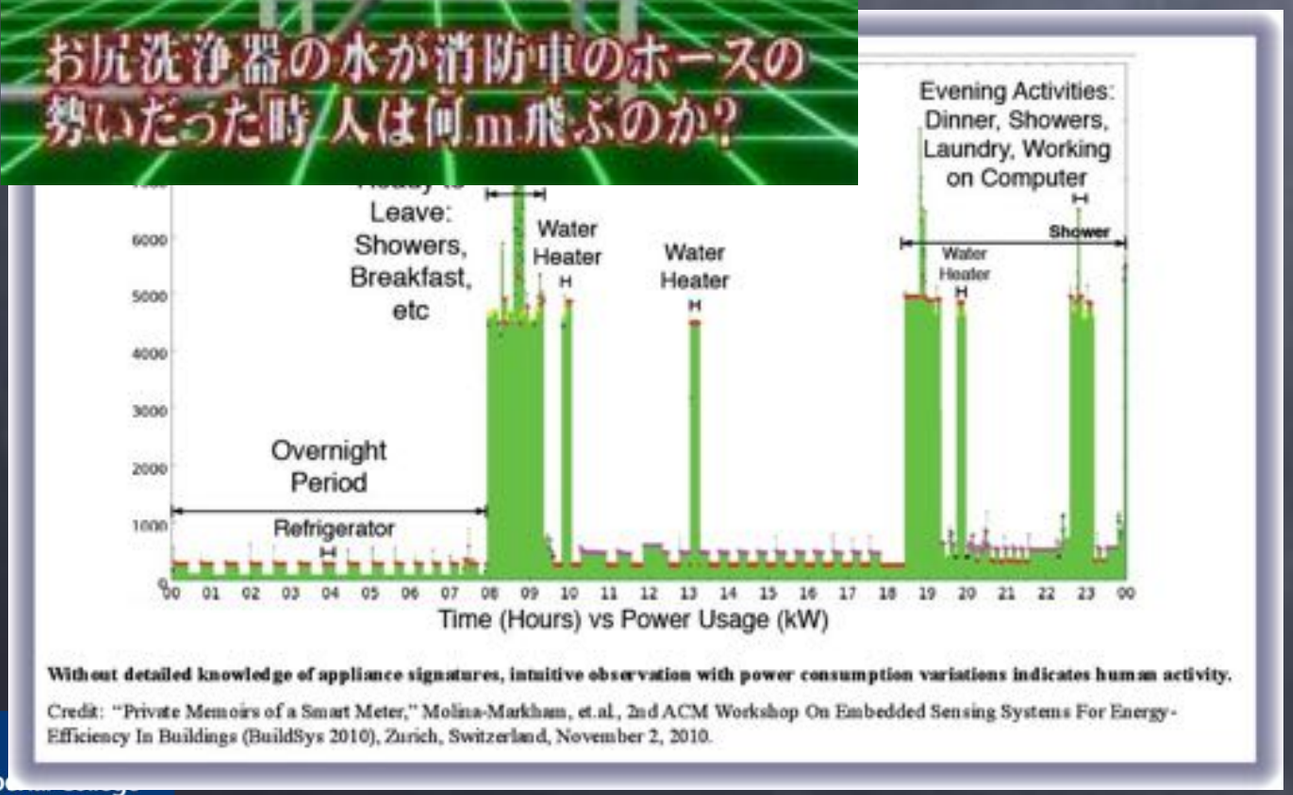

Without detailed knowledge of appliance signatures, intuitive observation with power consumption variations indicates human activity.

Credit: "Private Memoirs of a Smart Meter," Molina-Markham, et.al., 2nd ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings (BuildSys 2010), Zurich, Switzerland, November 2, 2010.

Trade-offs are hard

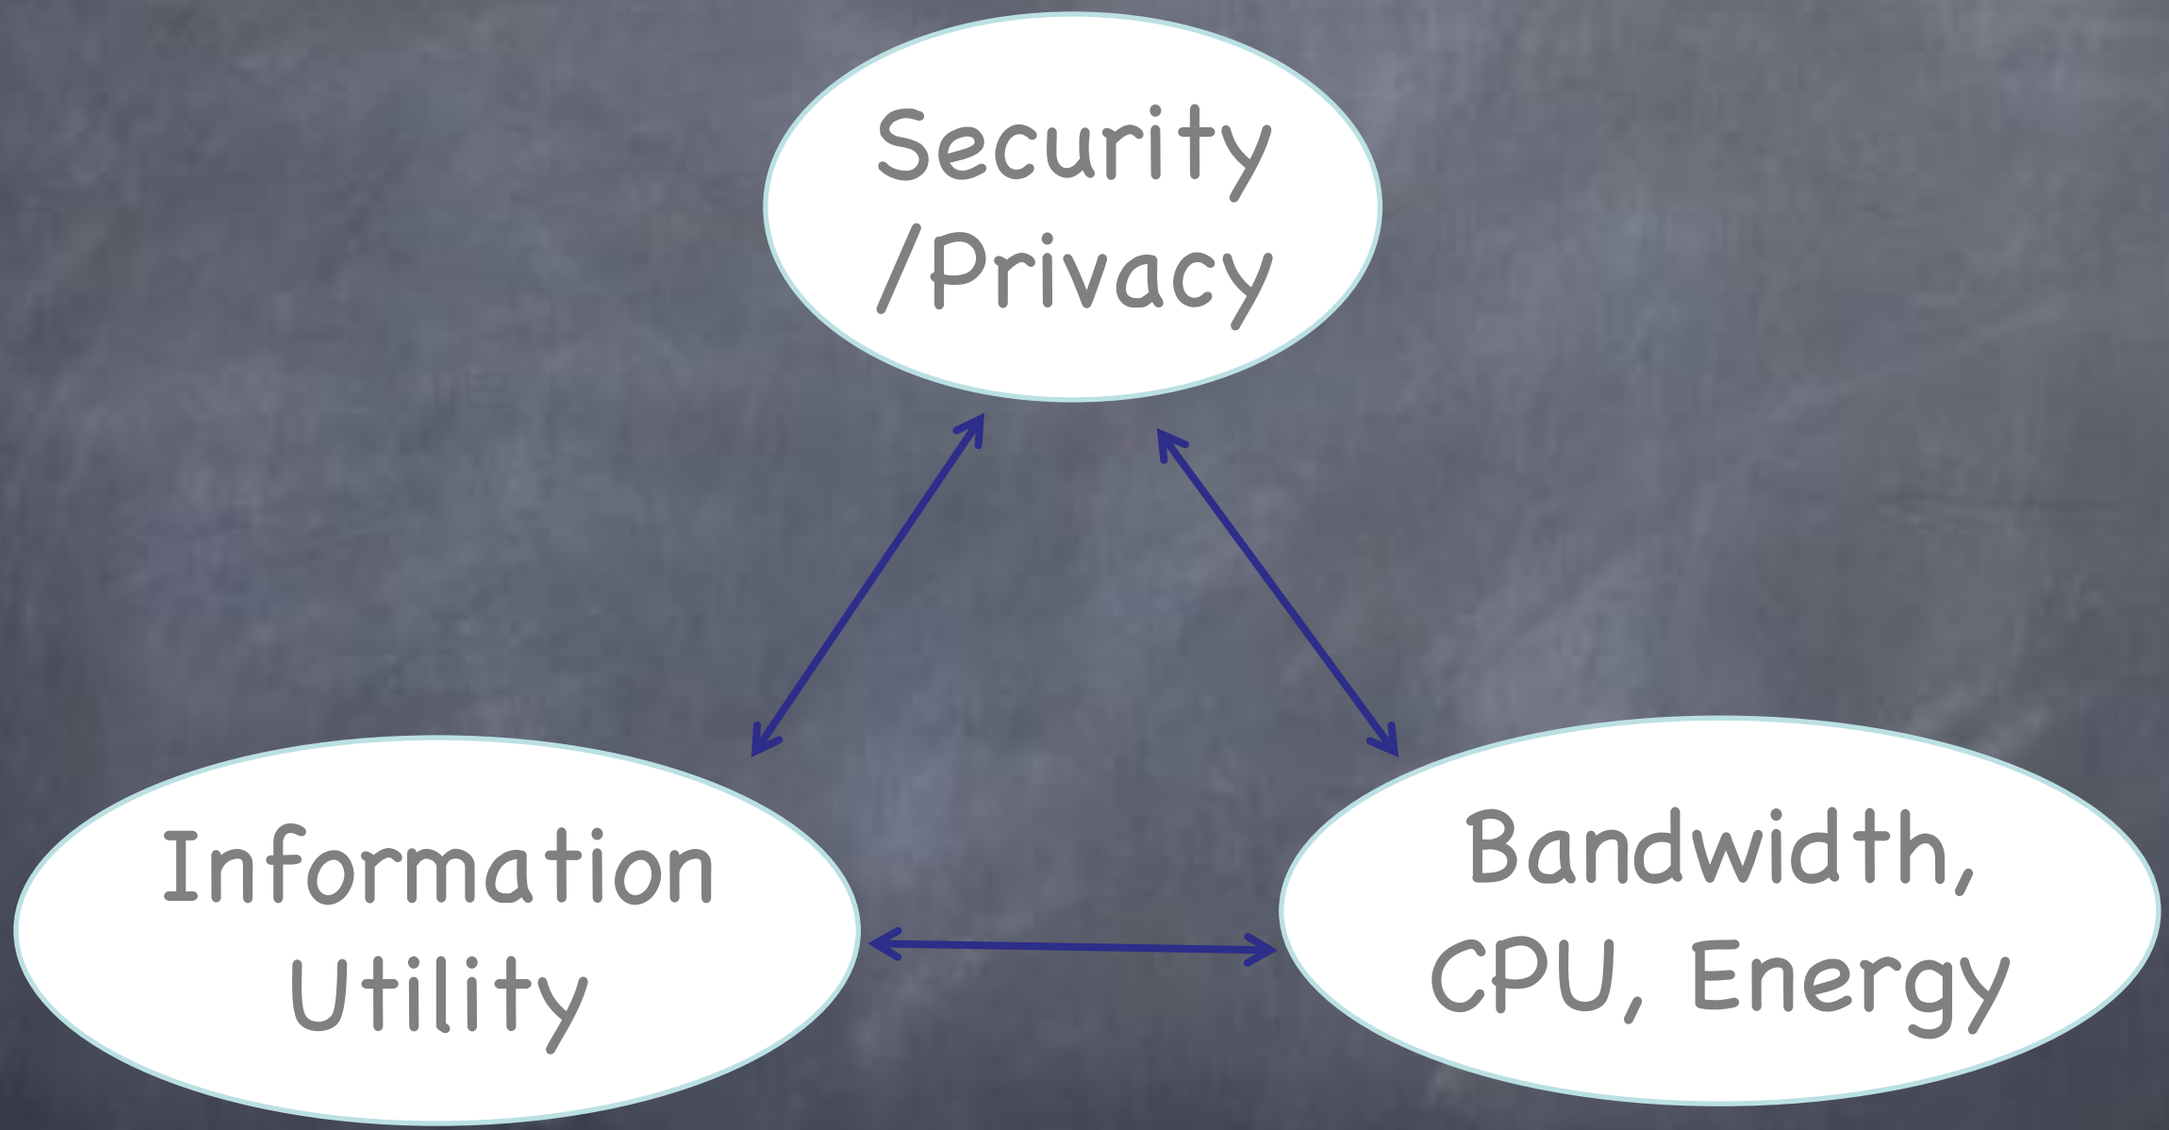

Lots of Hype, but also plenty of opportunities for networking and system to have impact.

That's it folks

For more information, software, and papers:

haddadi.github.io

We are hiring at all levels!

www.imperial.ac.uk/SysAI

Imperial College London