On Representation Learning

Alexander G. Ororbia II

Introduction to Machine Learning

CSCI-736

2/4/2025

Companion reading:
Chapter 15 of Deep Learning textbook

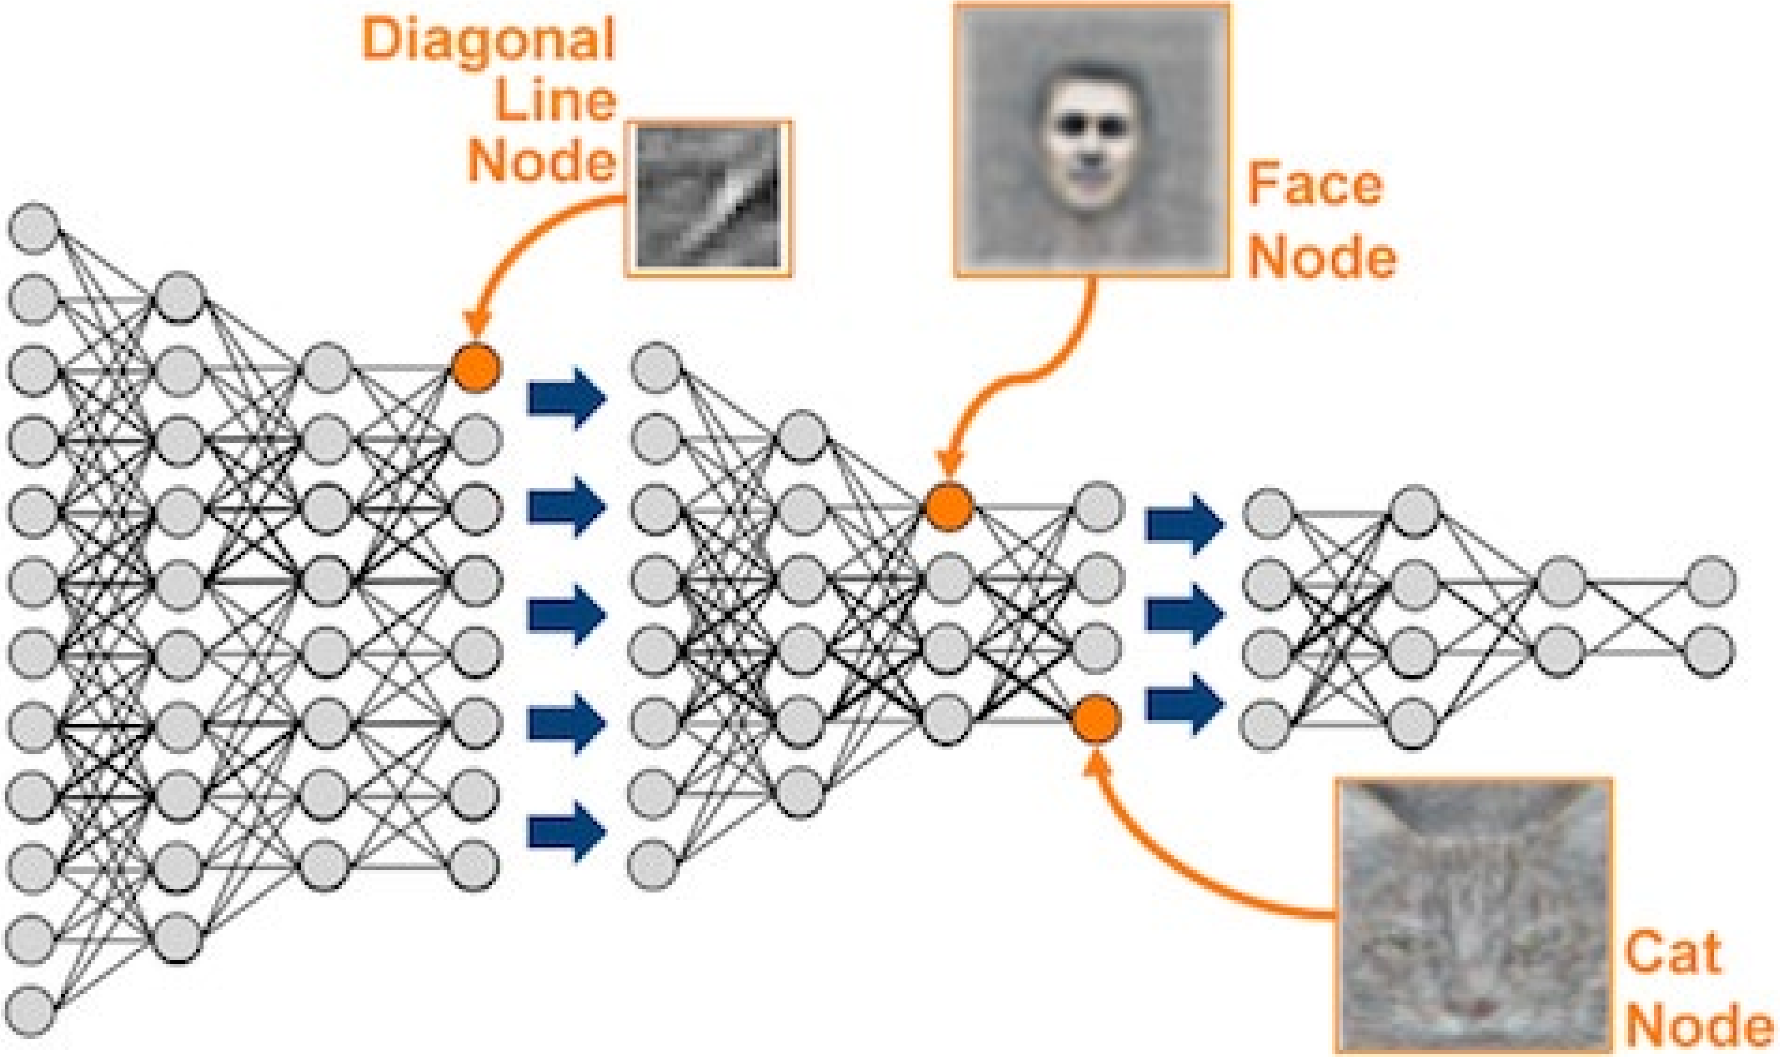
Diagonal Line Node
Face Node
Cat Node

On the plethora of model structures…

THE SPACE OF NEURAL ARCHITECTURES

Deep zoo!

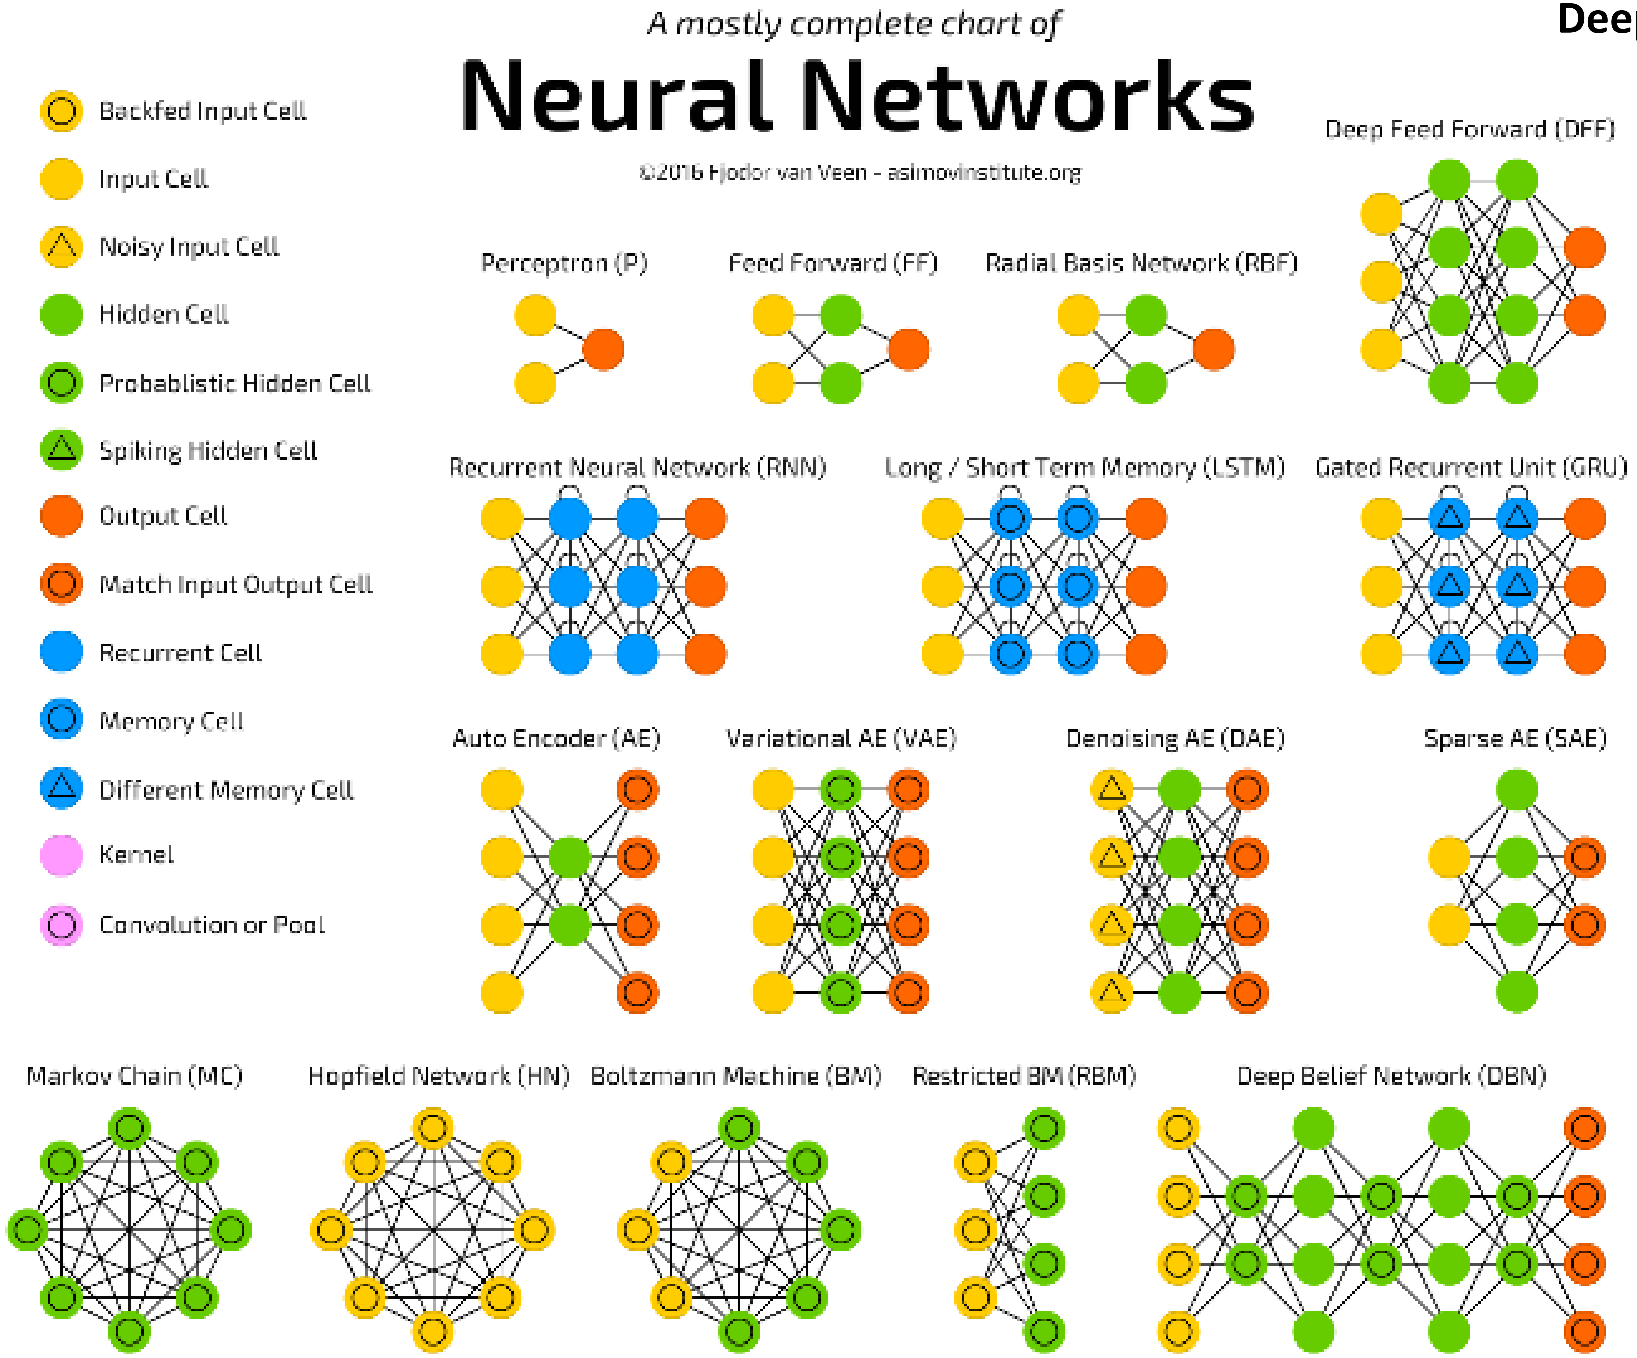

http://www.asimovinstitute.org/neural-network-zoo/

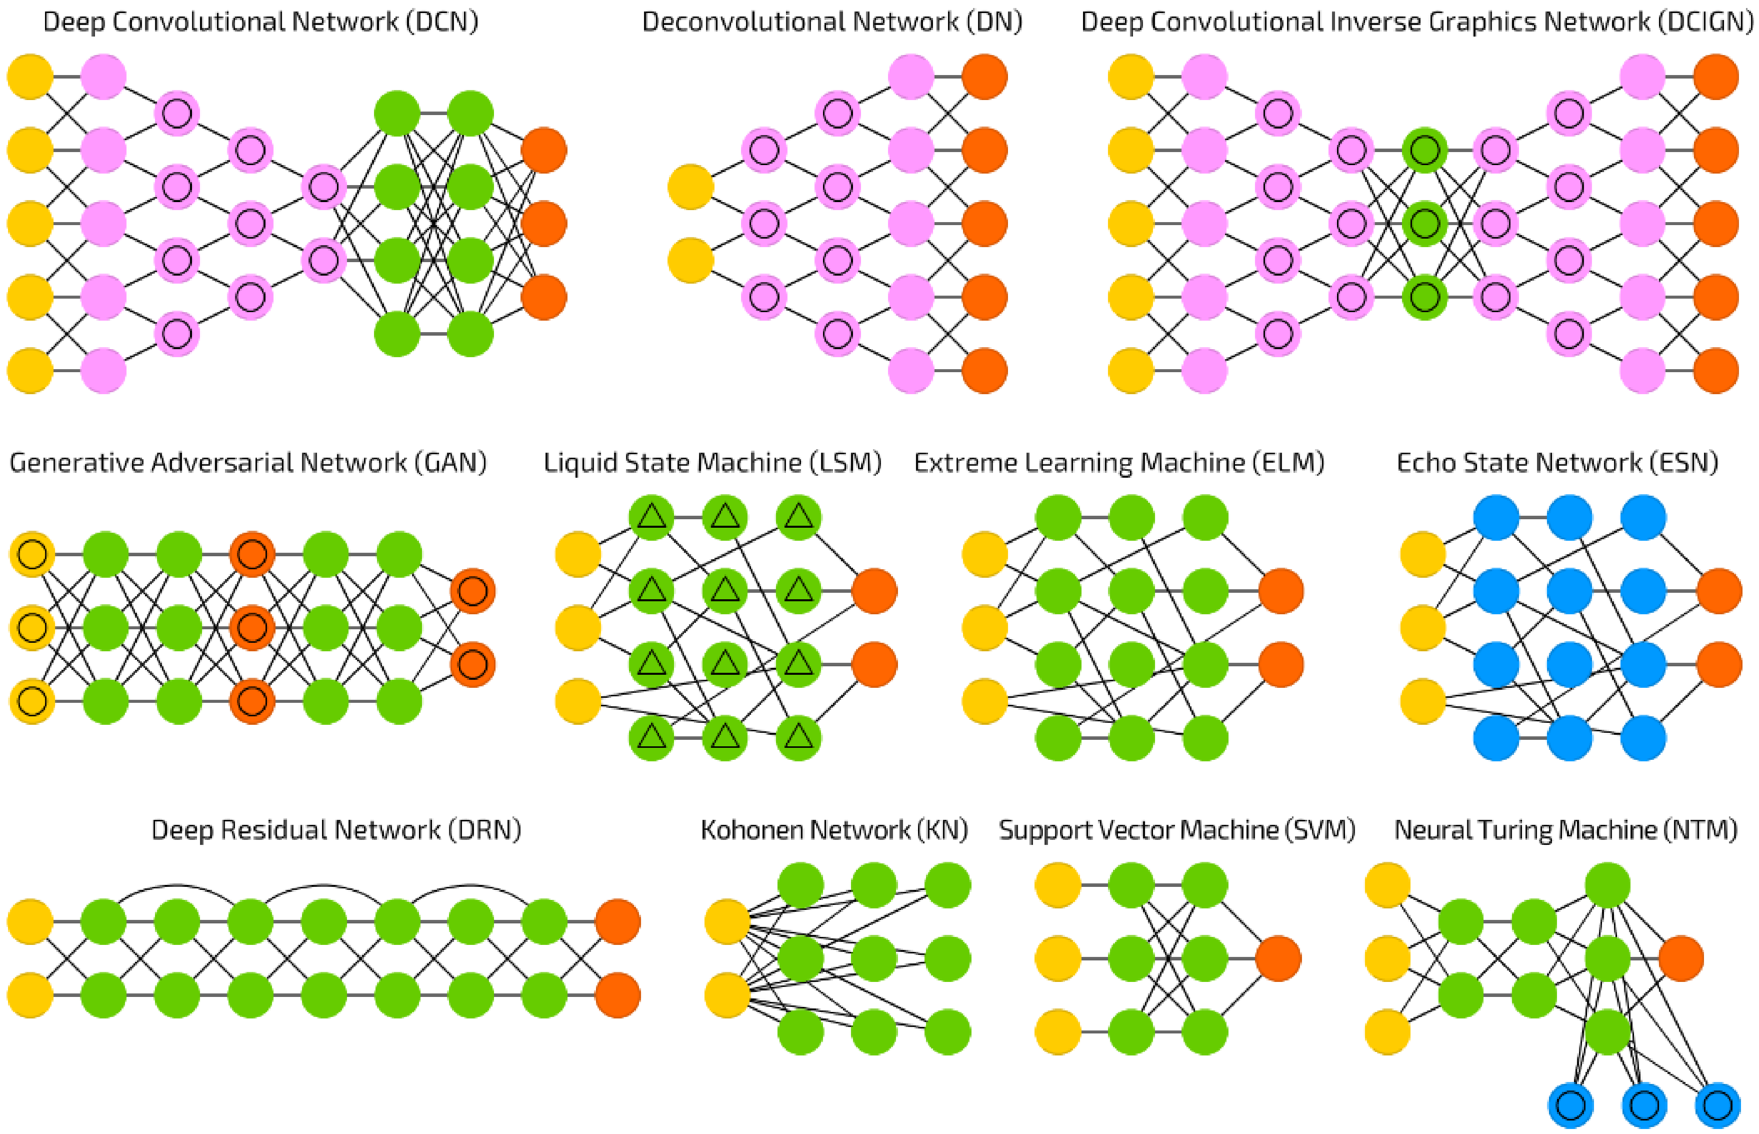

http://www.asimovinstitute.org/neural-network-zoo/

What is Representation Learning?

- Learn features automatically
  - Find a transformation of raw data input to a representation / space that can be effectively exploited in machine learning tasks
- Can be viewed as complementary to machine learning
  - “Automatic” pre-processing (or automated feature engineering)
- What makes one representation better than another?
  - ***Good representation*** = one that makes a subsequent learning task easier (choice of representation depends on the choice of subsequent learning task)

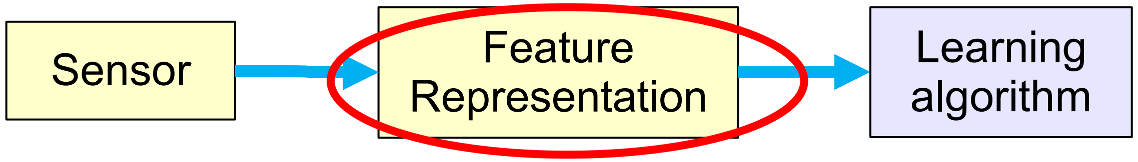

Why?

- *Intelligent systems need:*
  - Knowledge (constraints/priors)
  - Learning (optimization/search)
  - Generalization
    (guessing where probability mass concentrates)
  - Ways to fight off curse of dimensionality
    (exponentially many configurations of variables to consider)
  - Needs to disentangle underlying explanatory factors
    (making sense of data)

- Good **features** essential for successful ML: 90% of effort

raw input data → represented by **useful** features → MACHINE LEARNING

- Handcrafting features vs learning them

- Good representation?
- guesses

  the features / factors / causes

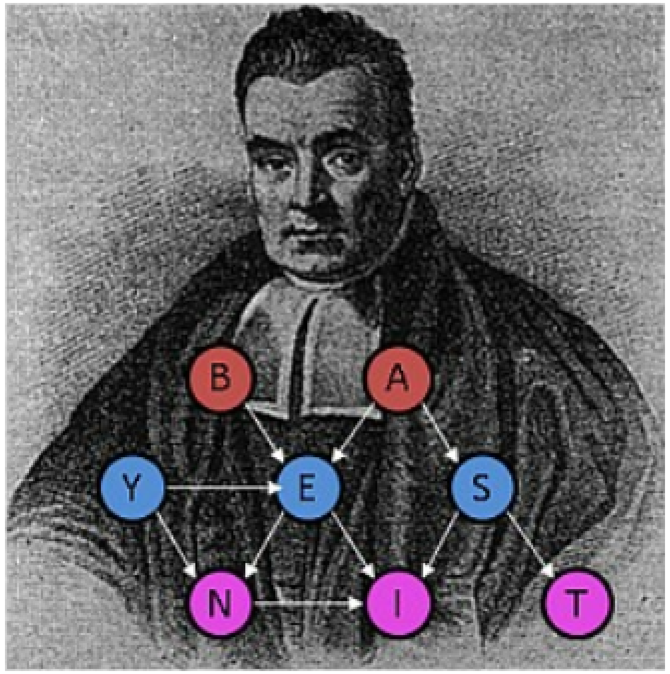

(Courtesy of Yoshua Bengio)

How do humans generalize from very few examples?

- They **transfer** knowledge from previous learning:
  - **Abstract** (i.e. deep) representations
  - Explanatory factors
- Previous learning from: unlabeled data

  + labels for other tasks

(Courtesy of Yoshua Bengio)

Computer Vision is Hard

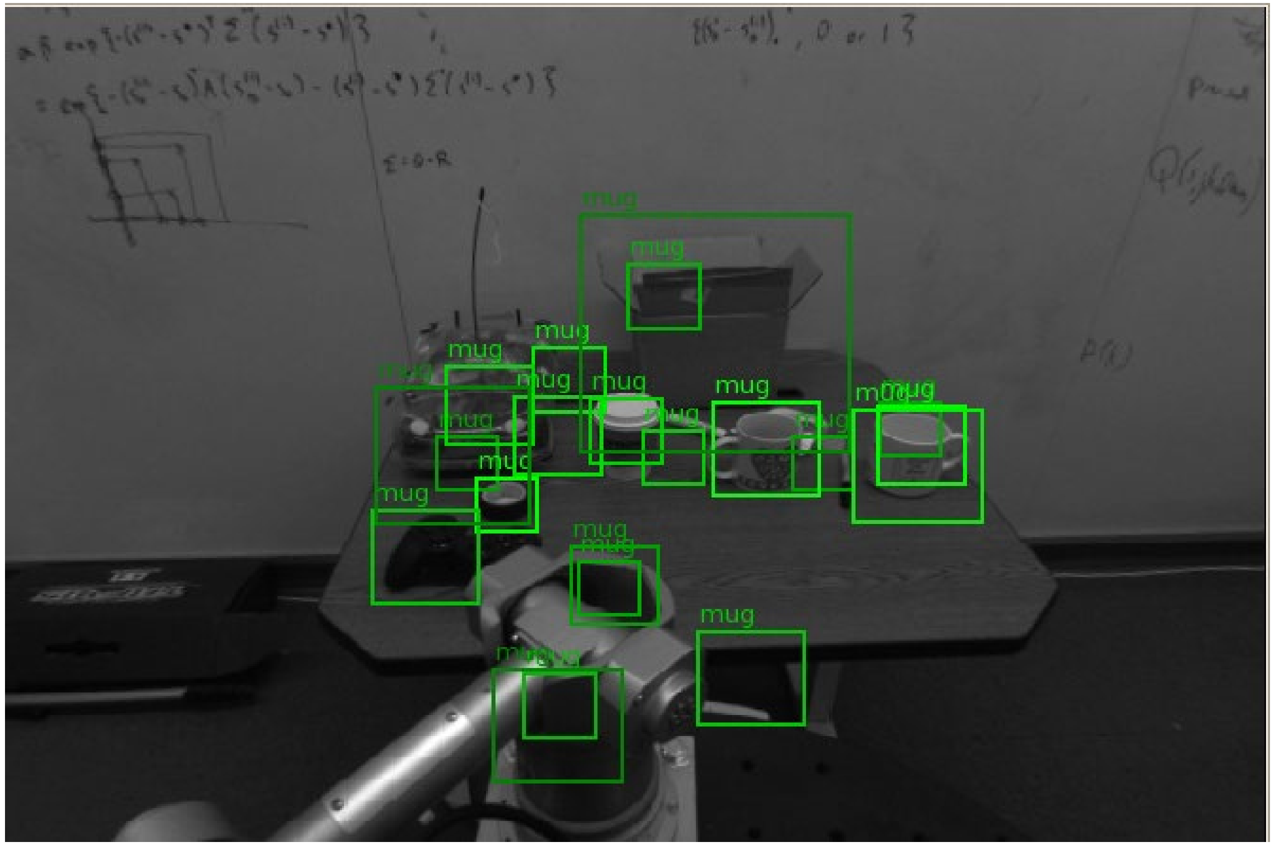

How is Computer Perception Done?

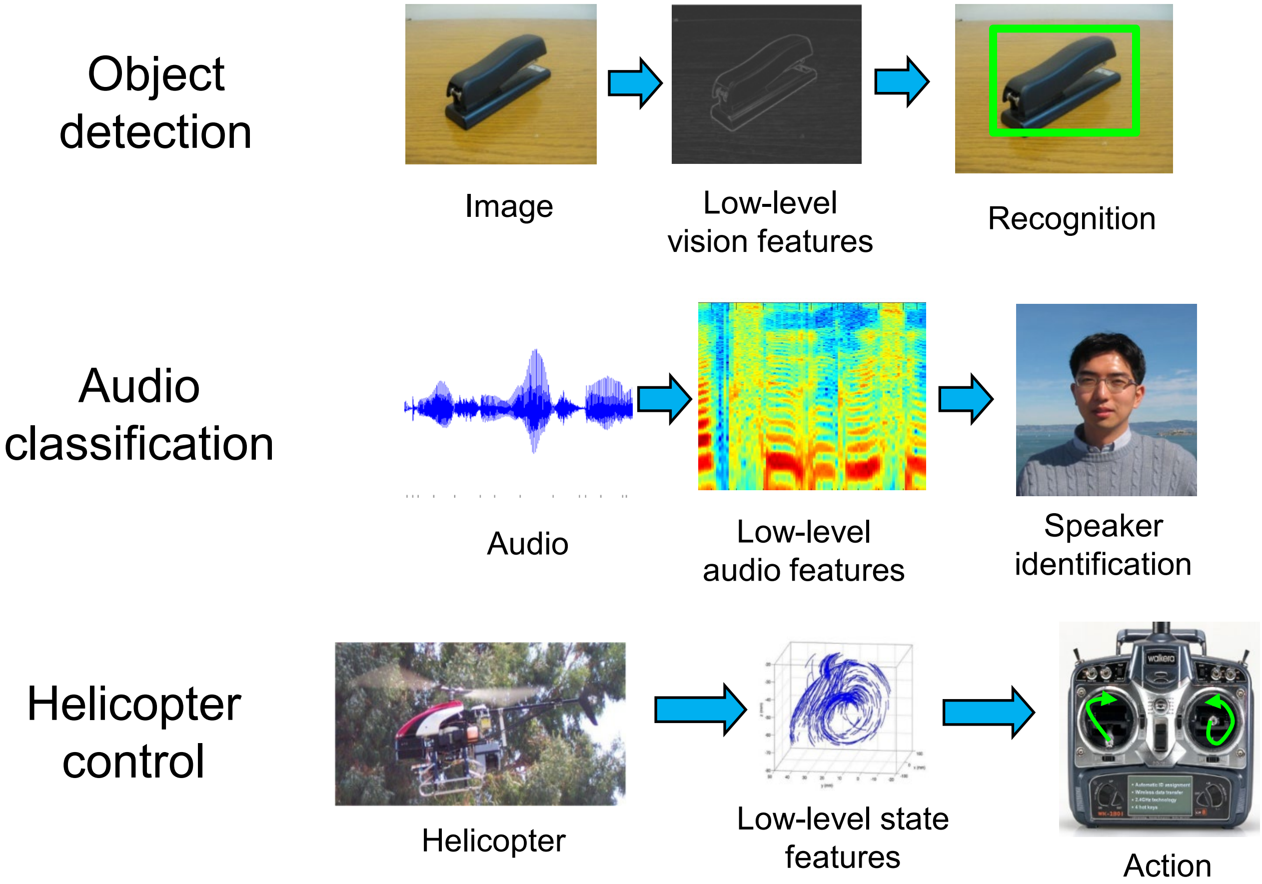

How is Computer Perception Done?

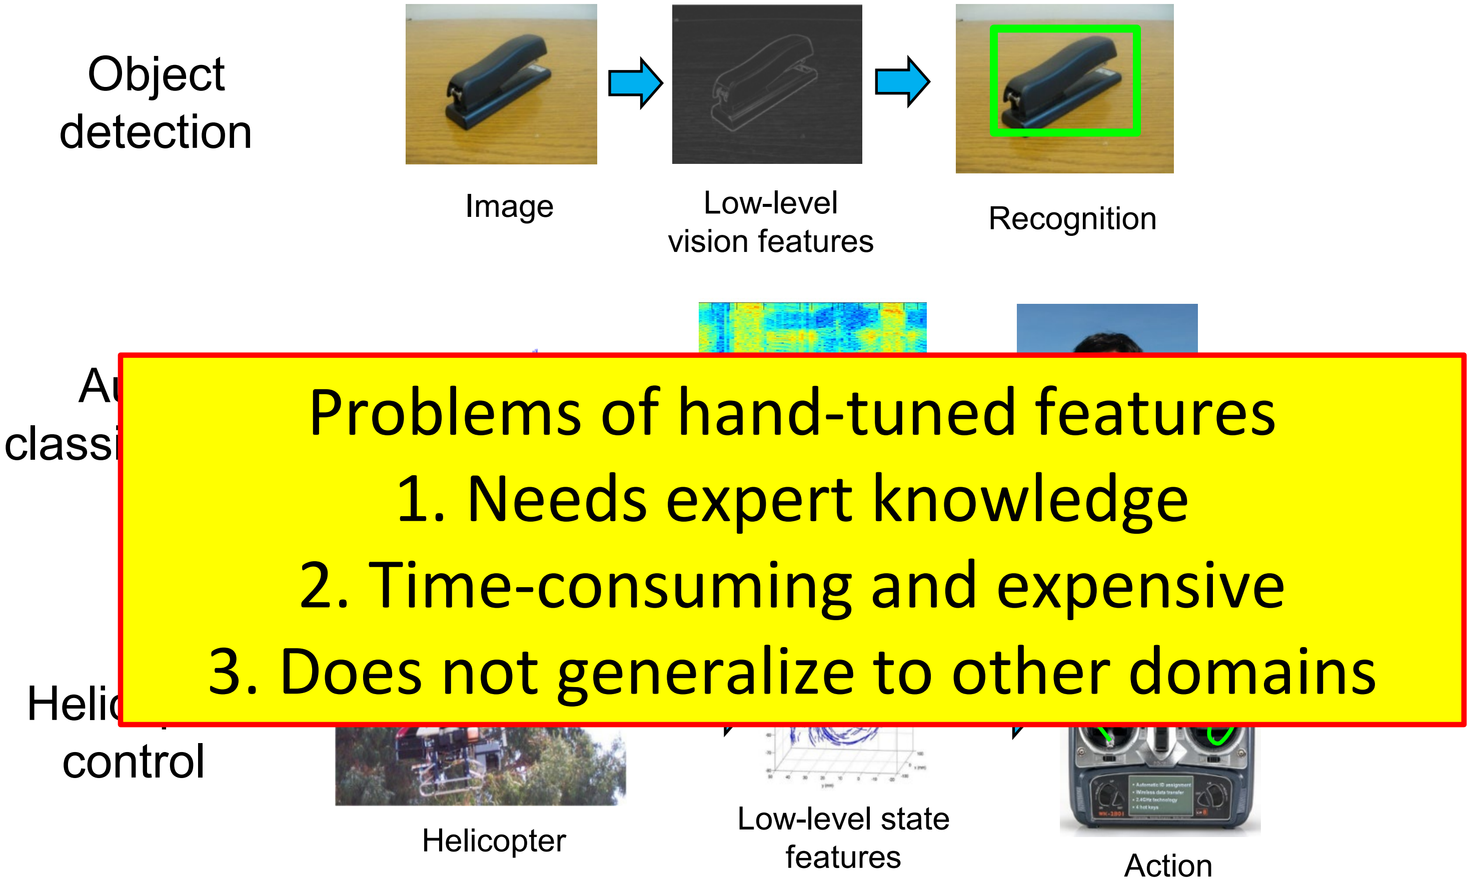

Problems of hand-tuned features

1. Needs expert knowledge
2. Time-consuming and expensive
3. Does not generalize to other domains

<u>In the context of a deep ANN:</u>
Training w/ supervised criterion naturally leads to representation at every hidden layer (moreso near top hidden layer) taking on properties that make discriminative/task-centric learning easier

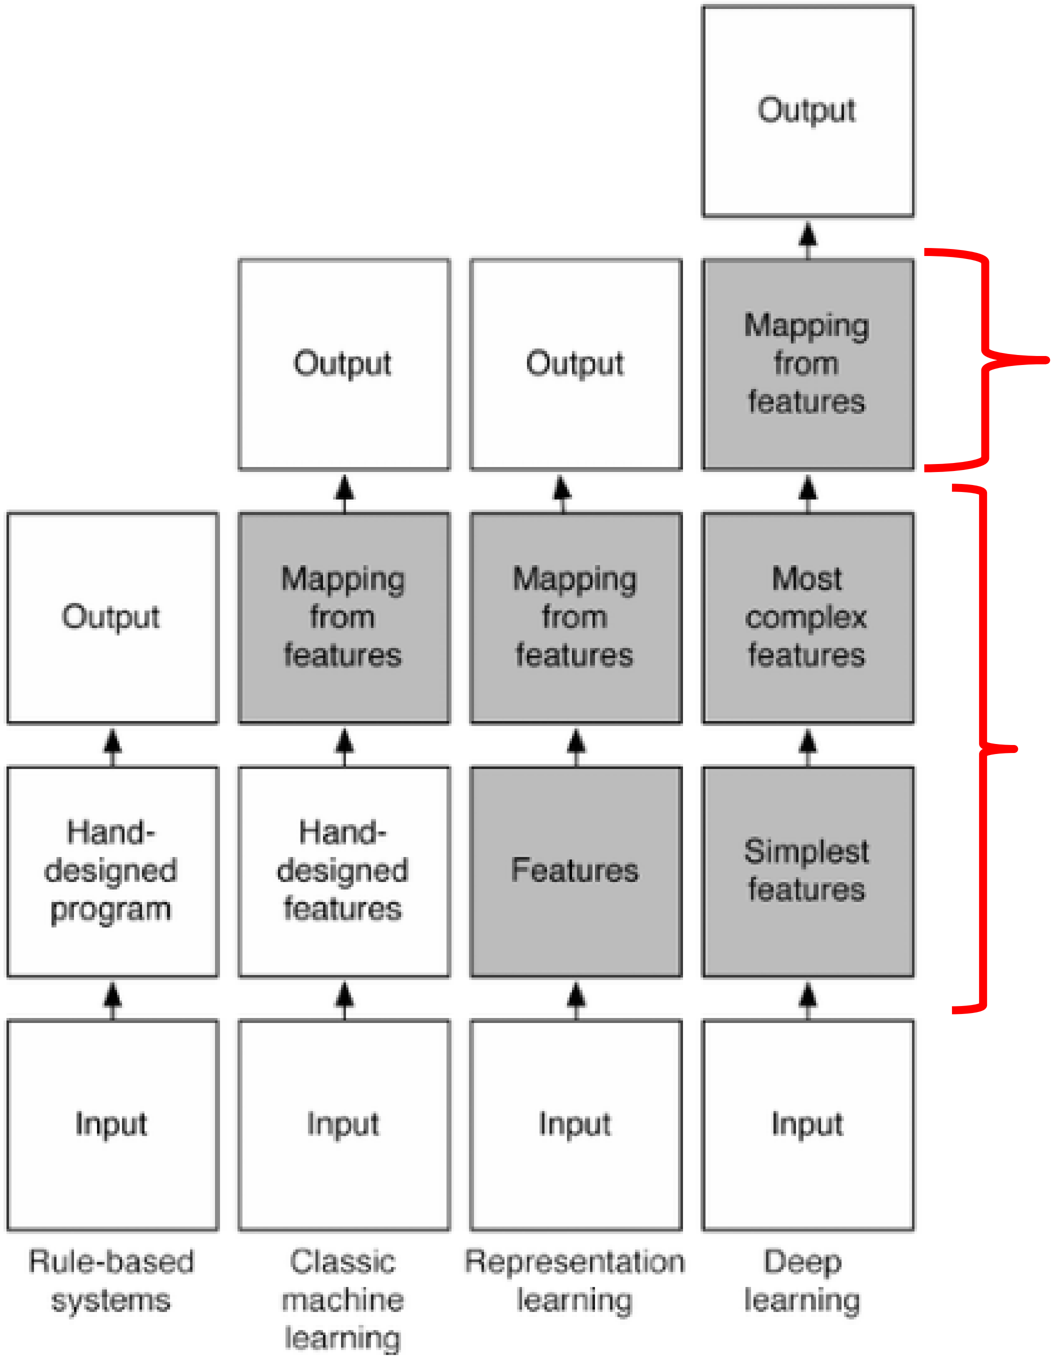

Linear classifier/regressor

Representation = set of transforms w/in ANN

Deep Architectures are More Expressive

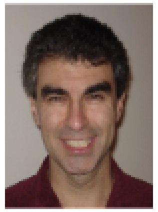
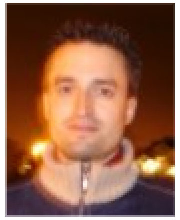


Theoretical arguments:

2 layers of { Logic gates / Formal neurons / RBF units } = universal approximator

RBMs & auto-encoders = universal approximator

Theorems on advantage of depth:
(Hastad et al 86 & 91, Bengio et al 2007, Bengio & Delalleau 2011, Braverman 2011)

Some functions compactly represented with k layers may require exponential size with 2 layers

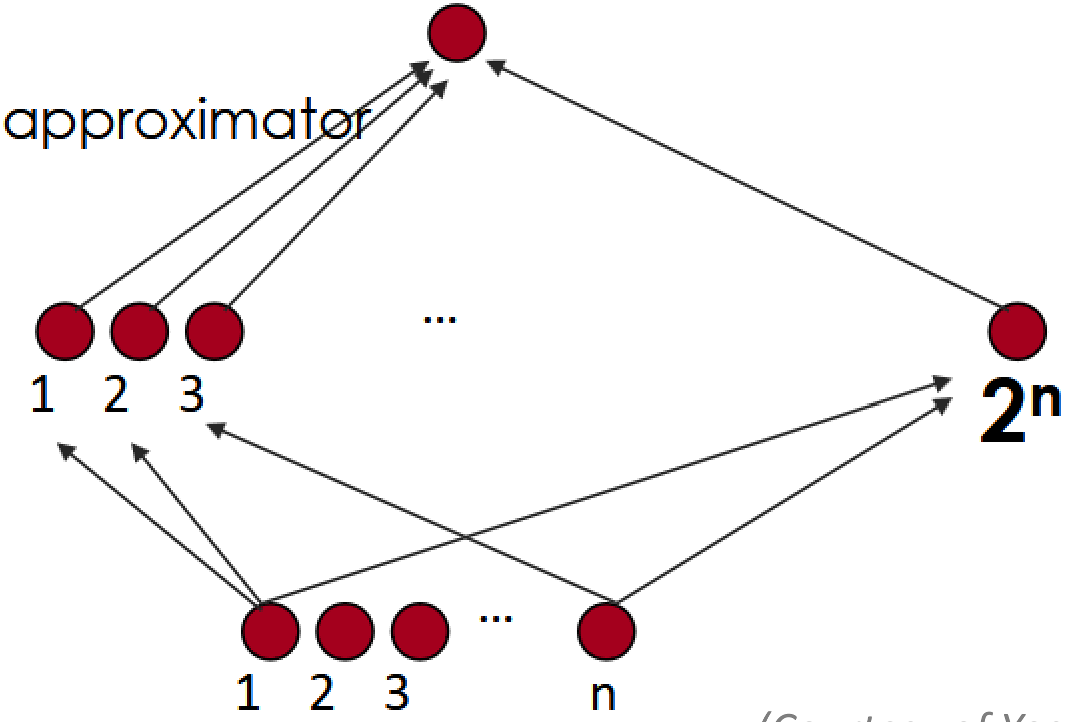

(Courtesy of Yoshua Bengio)

Successive model layers learn deeper intermediate representations

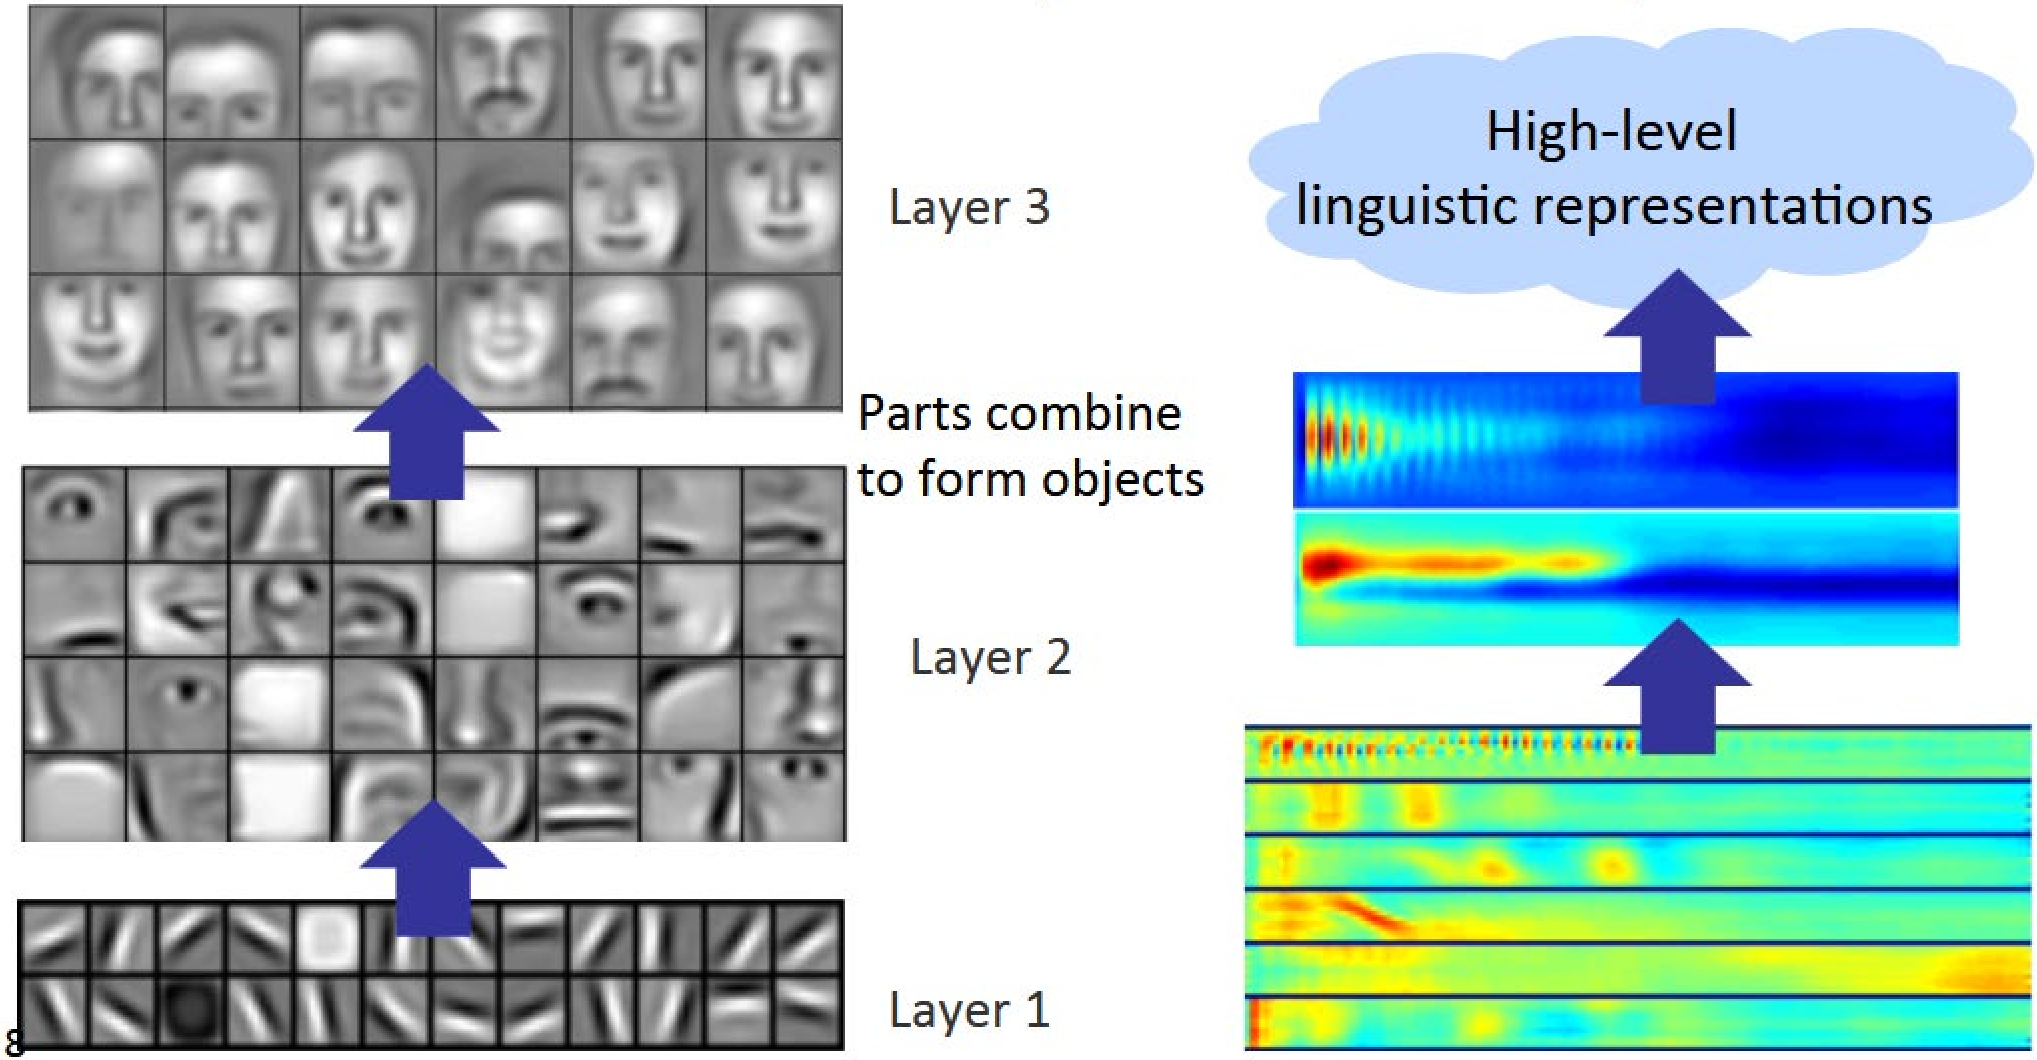

Prior: underlying factors & concepts compactly expressed w/ multiple levels of abstraction

(Courtesy of Yoshua Bengio)

The Problem of Representation Learning

- Representation learning problems face trade-off between preserving as much information about input (as possible) and attaining nice properties
  - Models (supervised or unsupervised) have main training objective but learn a representation as a "side effect"
- Offers pathway to facilitate semi-supervised learning
  - ***Hypothesis***: unlabeled data can be used to learn a good representation

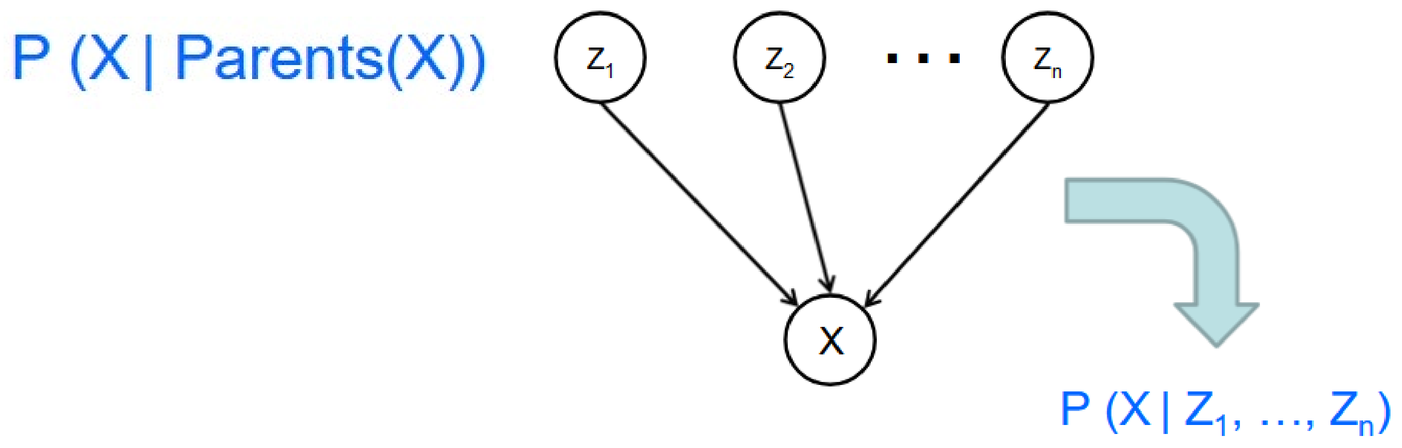

The Problem of Representation Learning

- Representation learning problems face trade-off between preserving as much information about input (as possible) and attaining nice properties, e.g., might want things like independence of detectors
- Often add constraints to shape representation in some cases
  - Density estimation – encourage elements of representation/latent vector **z** to be independent (distributions w/ more independencies are easier to model)

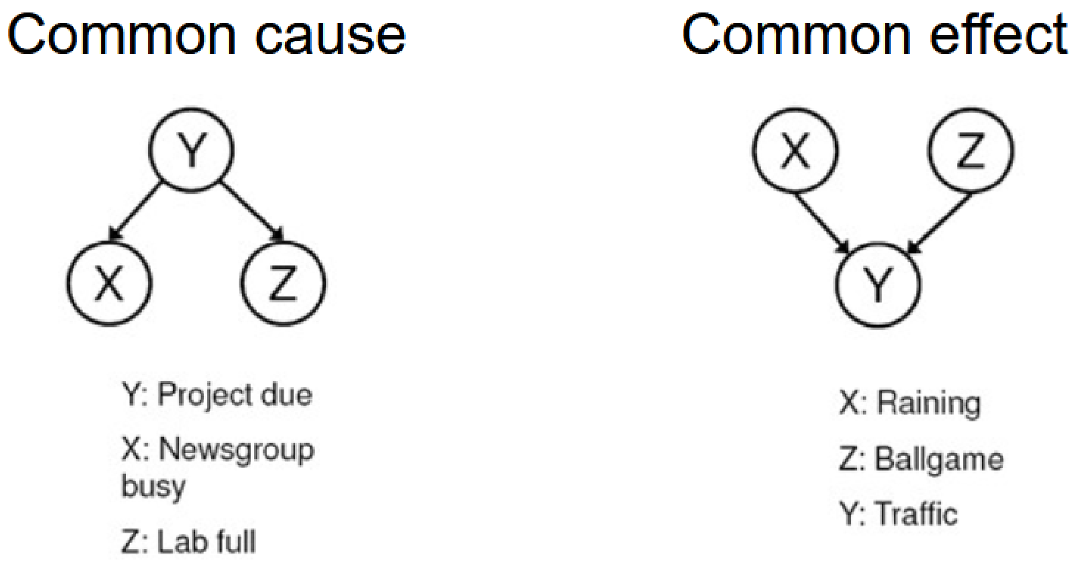

Representations Can Be Acquired through Multi-Task Learning

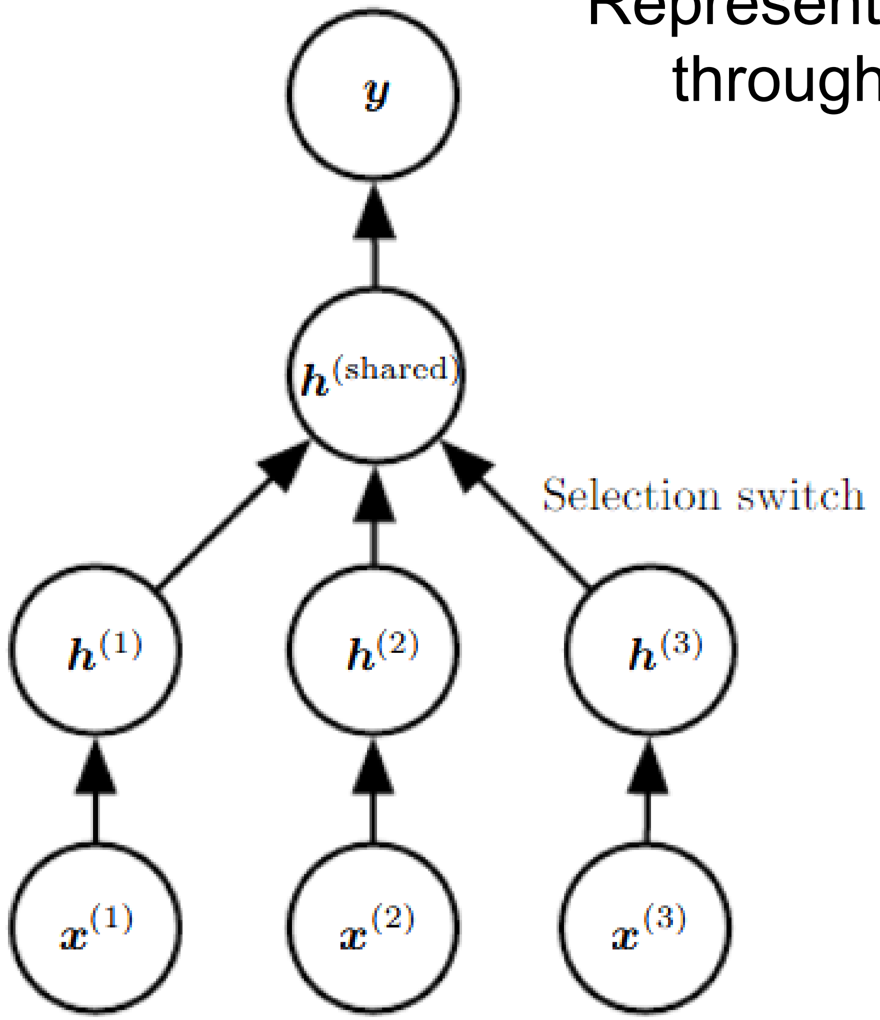

Representations Can Be Acquired by Mapping Across Encodings of Modalities/Variable Sets

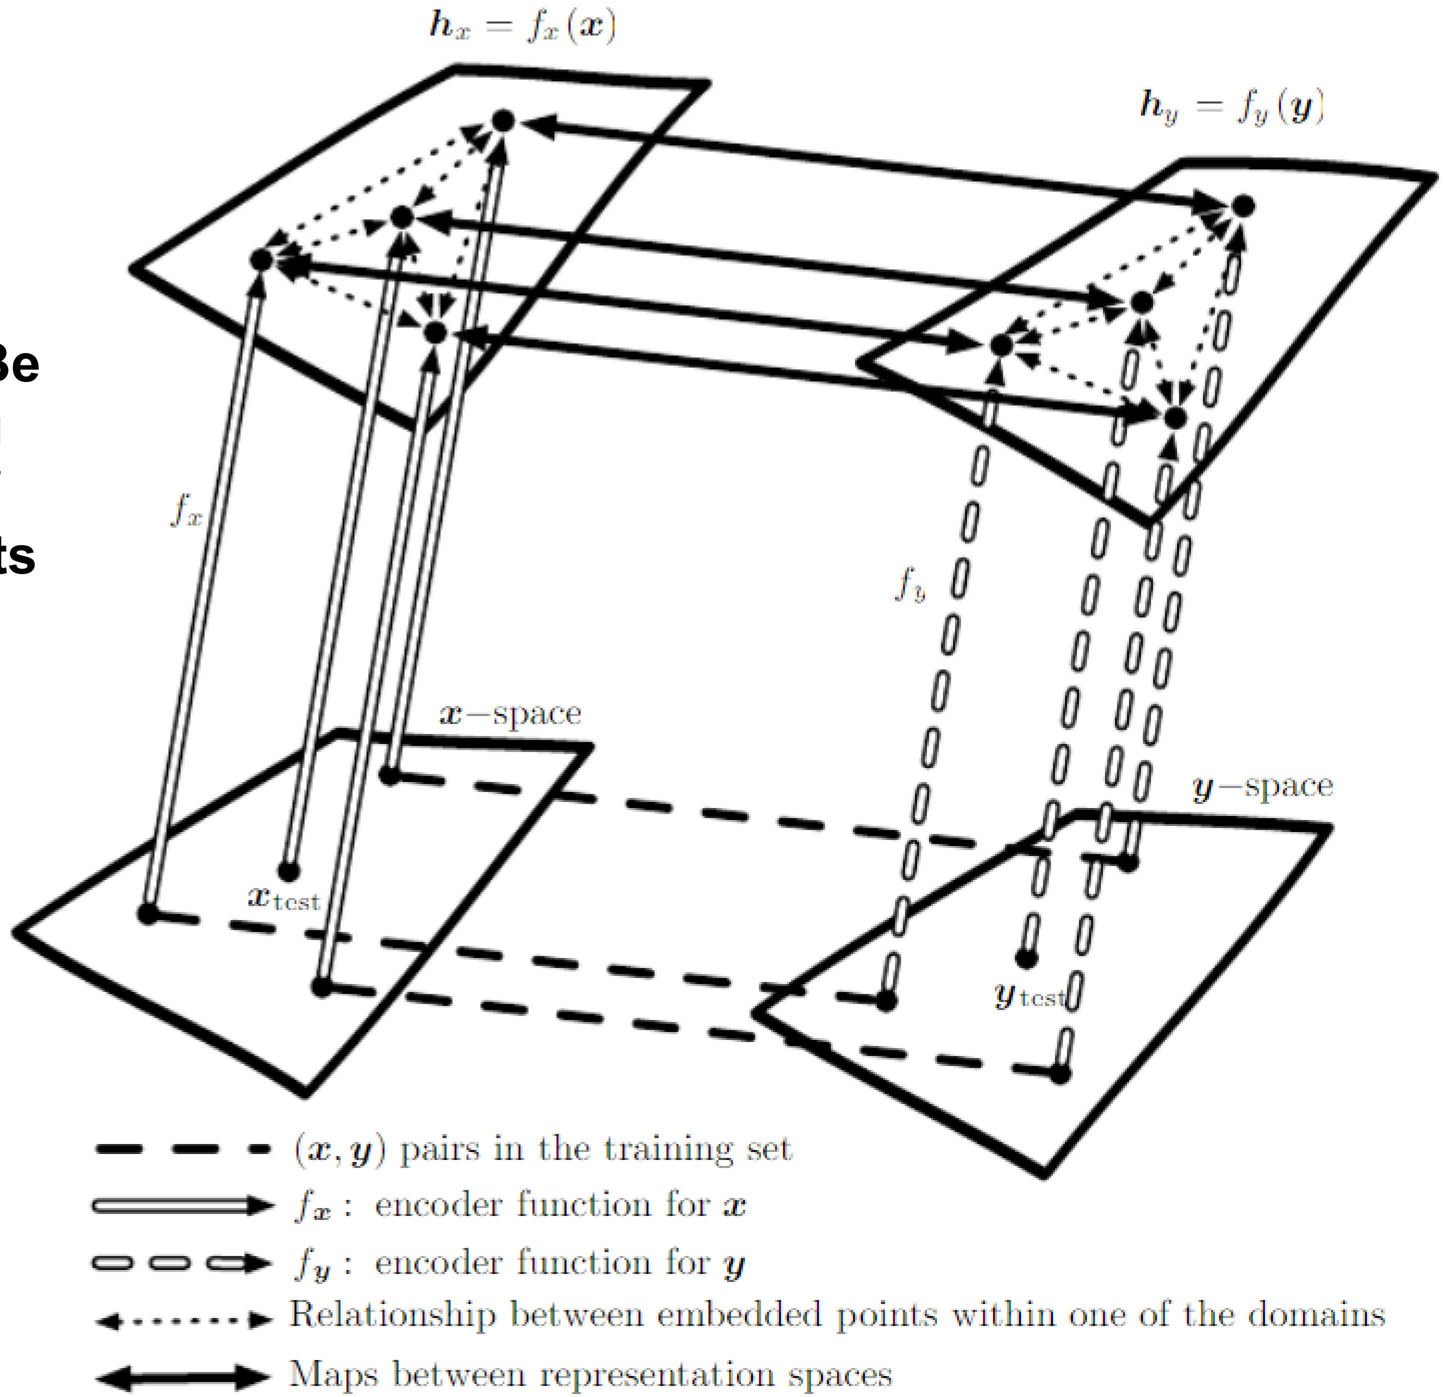

QUESTIONS?

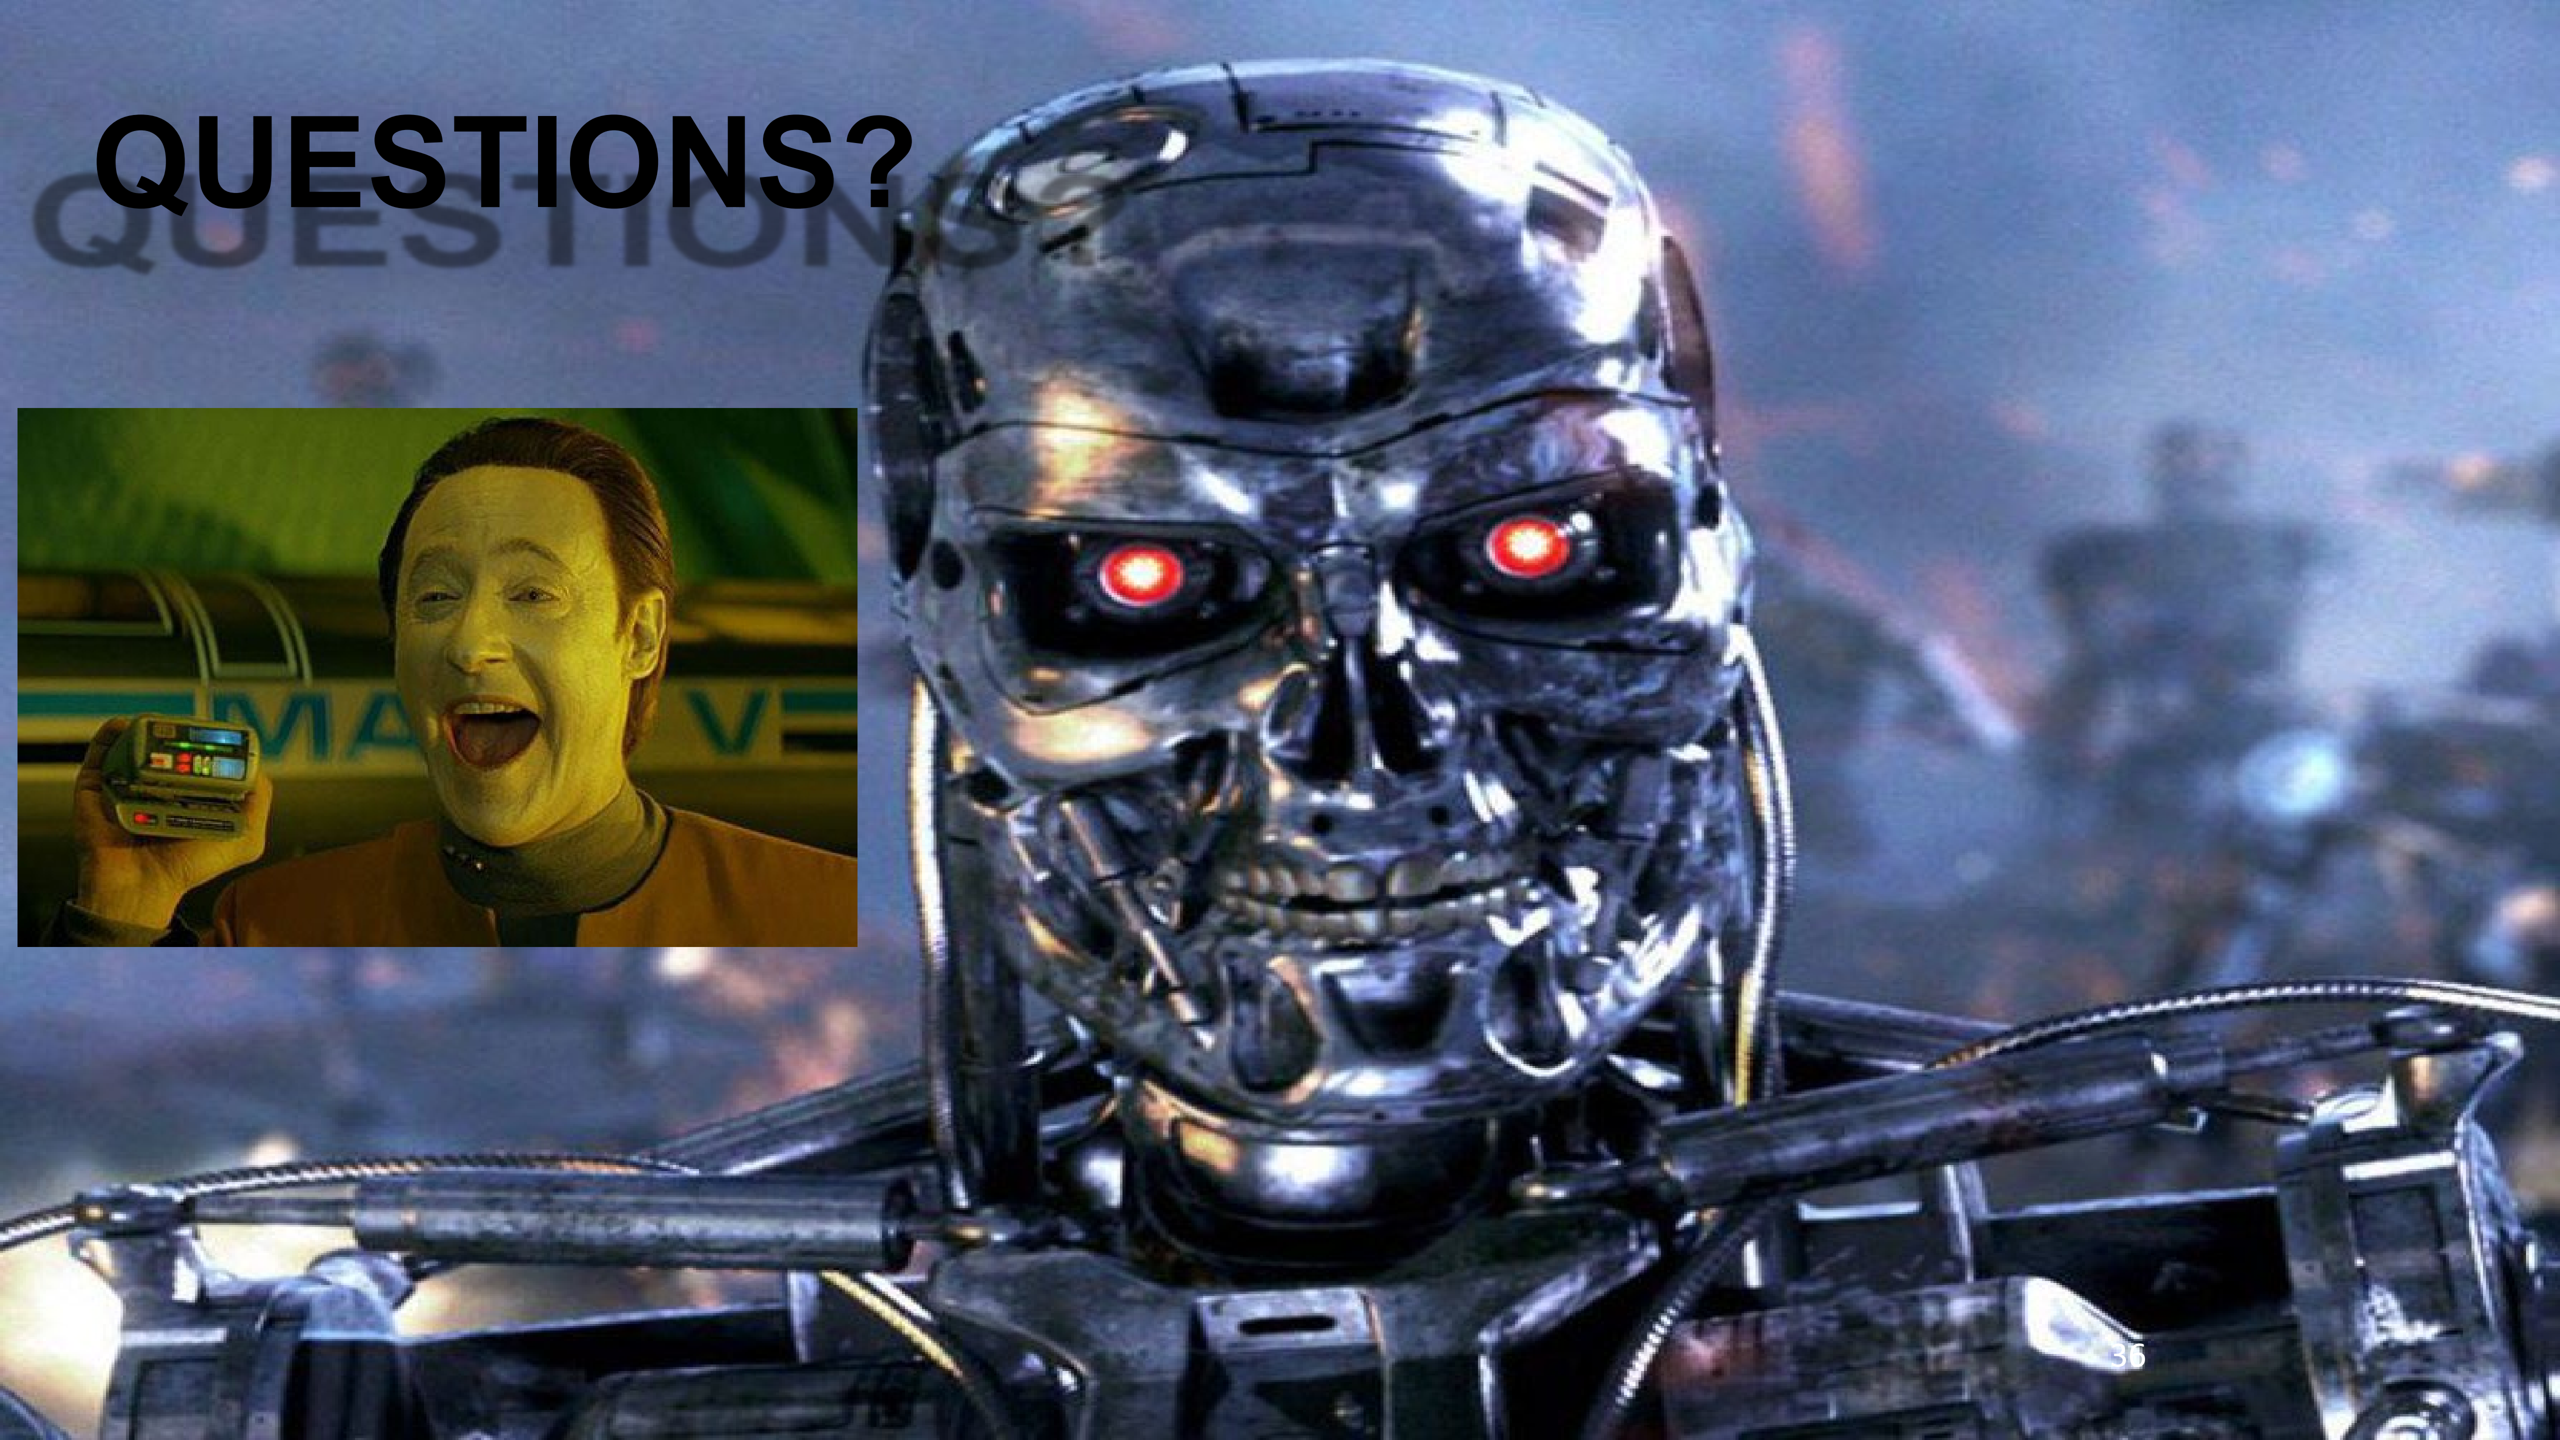

Unsupervised Feature Learning






Find a better way to represent images (or low-level data in general) than pixels (or low-level/raw sensory features).