Machine learning: An overview of the landscape

B. Artsybasheff (image source)

Lecture overview

- The statistical learning viewpoint
- Supervised classification
- Beyond supervised classification: A taxonomy of prediction problems and types of supervision

How can we build an agent to…

Play chess?

Translate between languages?

Recognize object categories?

Fly a drone?

How can we build an agent to achieve expertise?

- Good old-fashioned AI (GOFAI) answer:
  Program expertise into the agent
  - Never worked (in general)…

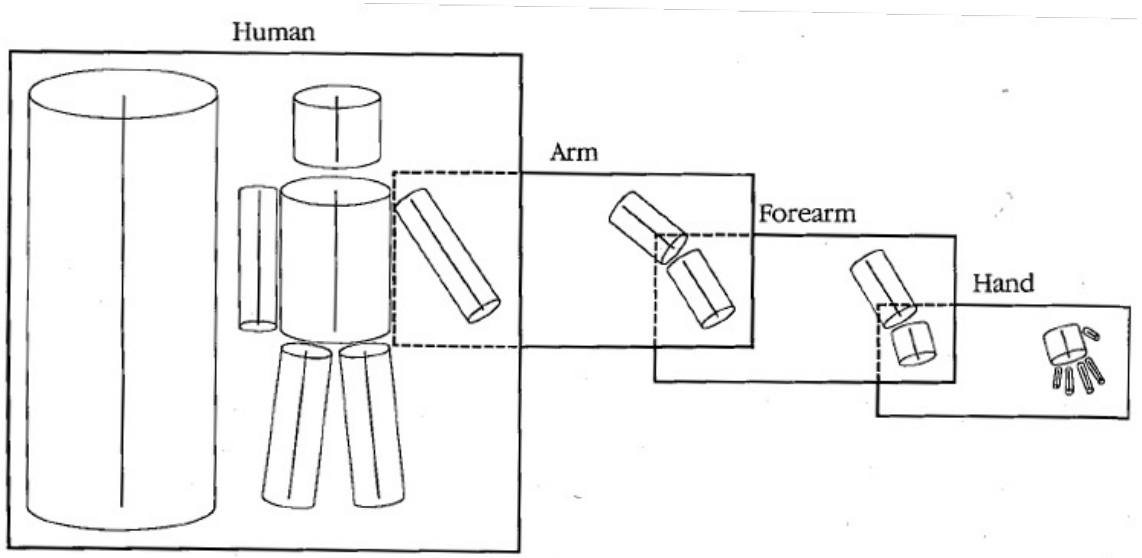

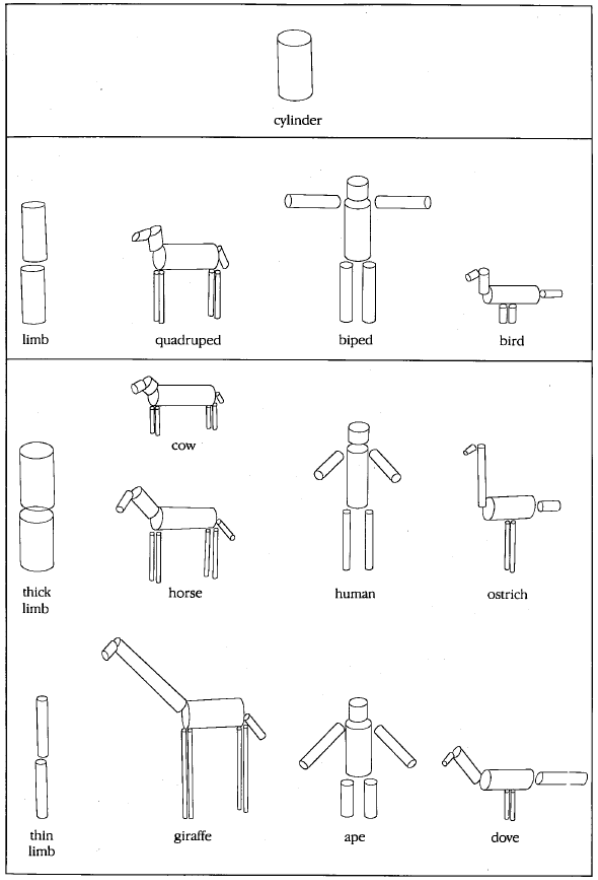

Figures from Marr's Vision (1982)

How can we build an agent to achieve expertise?

- Good old-fashioned AI (GOFAI) answer:
  Program expertise into the agent
  - Never worked (in general)…
- Modern answer: Program into the agent the *ability to improve performance based on experience*
  - Experience should come from *training data* or *demonstrations*
  - Learning is optimizing performance of the agent on the training data, with the hope that it will *generalize* to unseen inputs
  - This is the *statistical learning* viewpoint

Example: Image classification

Input	Desired output
	apple
	pear
	tomato
	cow
	dog
	horse

Example: Image classification

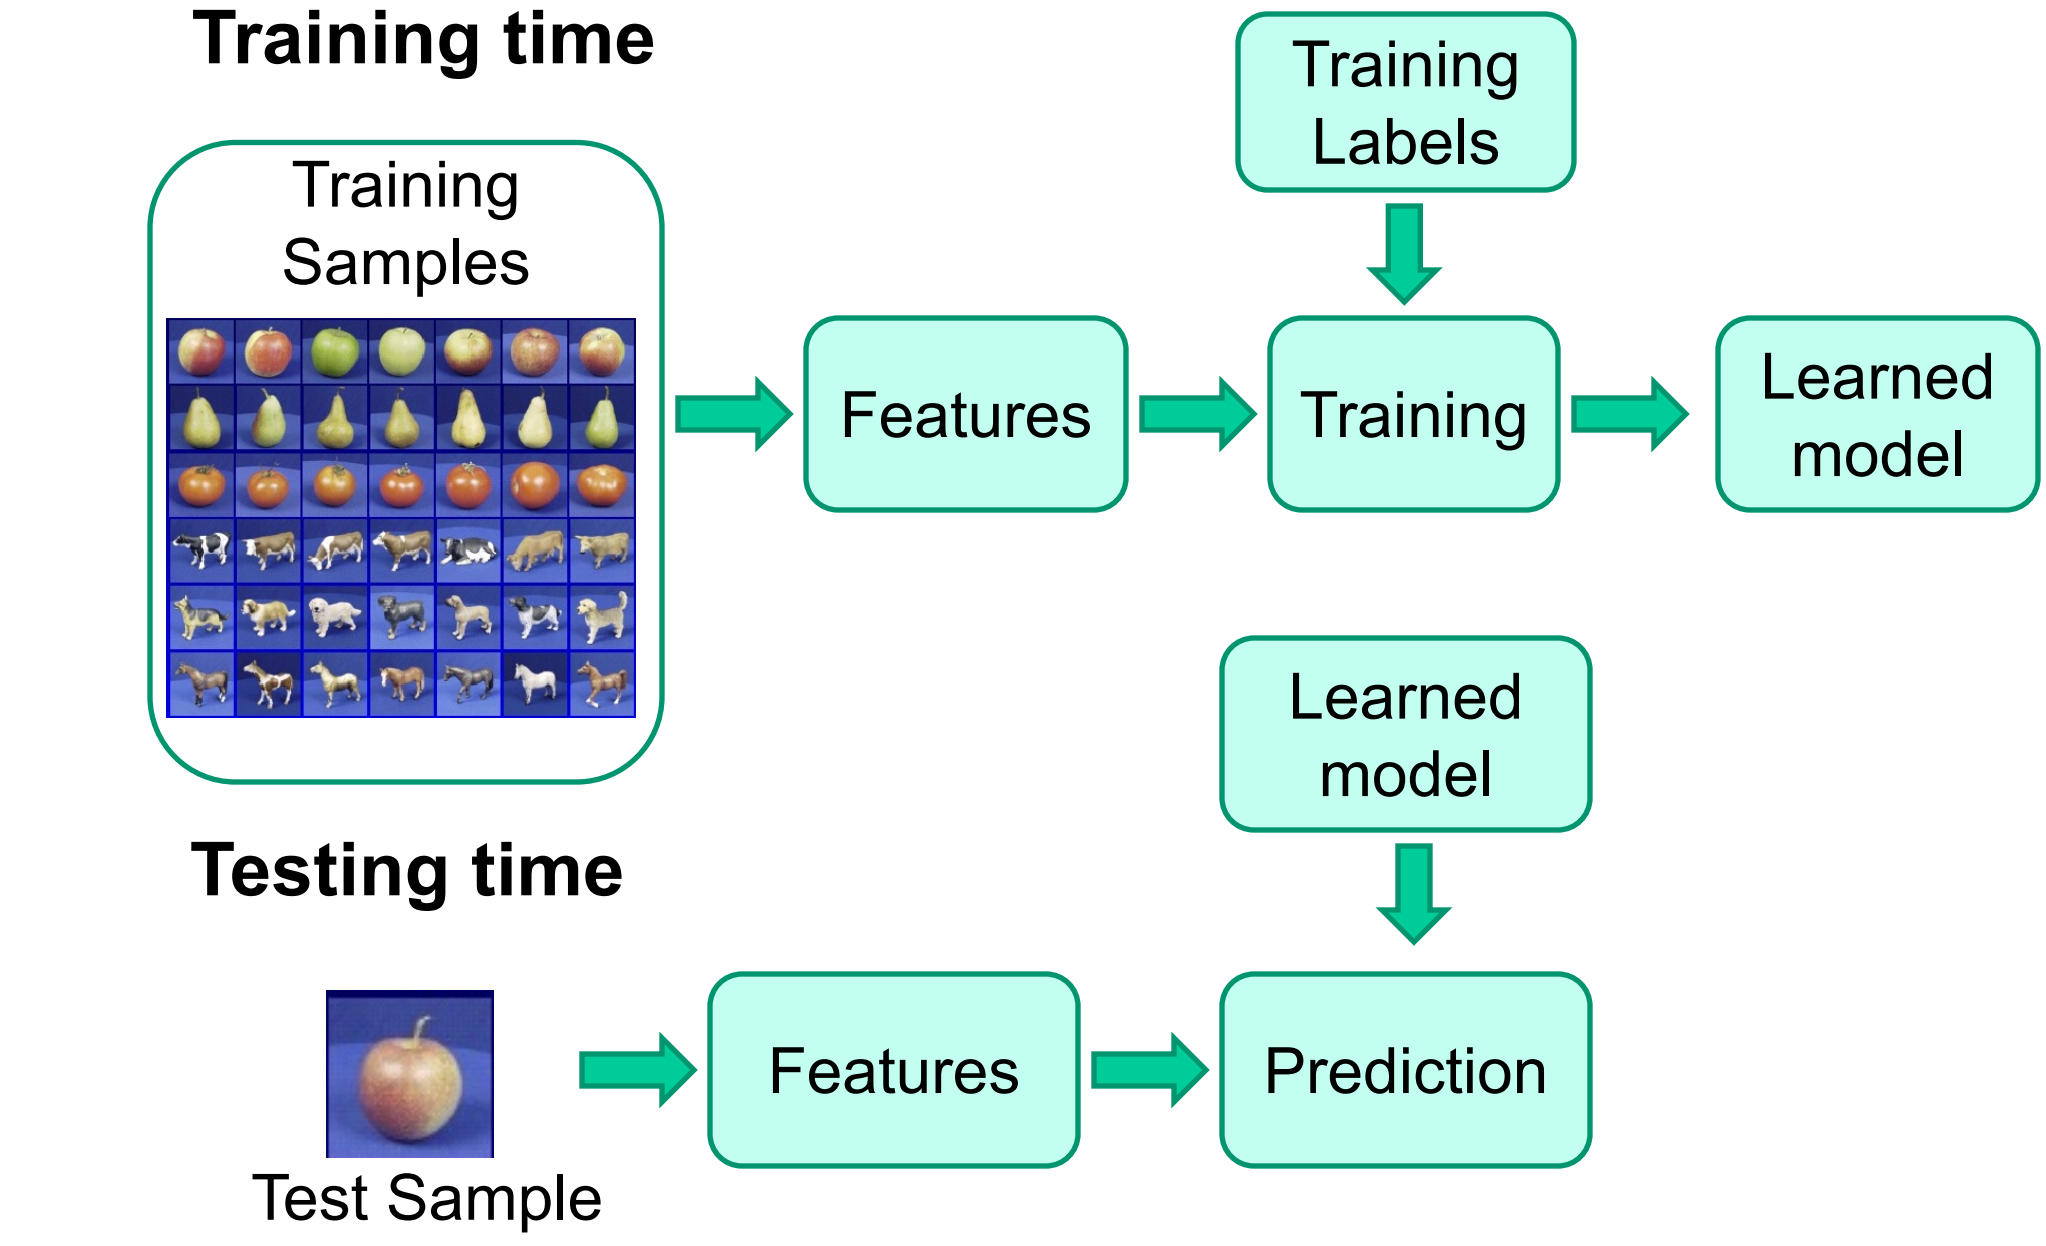

The basic *supervised learning* framework

$$y = f(x)$$

output

prediction function

input

- **Training** (or **learning**): given a *training set* of labeled examples $\{(x_1, y_1), \ldots, (x_N, y_N)\}$, instantiate a predictor f
- **Testing** (or **inference**): apply f to a new *test example* x and output the predicted value $y = f(x)$

More supervised learning examples: Text classification

Spam classification

Ok, Iknow this is blatantly OT but I'm beginning to go insane. Had an old Dell Dimension XPS sitting in the corner and decided to put it to use, I know it was working pre being stuck in the corner, but when I plugged it in, hit the power nothing happened.

Dear Sir.

First, I must solicit your confidence in this transaction, this is by virture of its nature as being utterly confidencial and top secret. …

Sentiment classification

"I love this movie.
I've seen it many times
and it's still awesome."

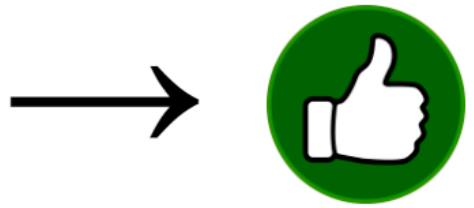

"This movie is bad.
I don't like it it all.
It's terrible."

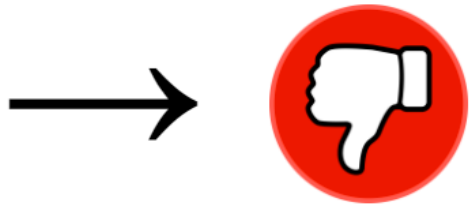

Image source

Another example: Grasp classification

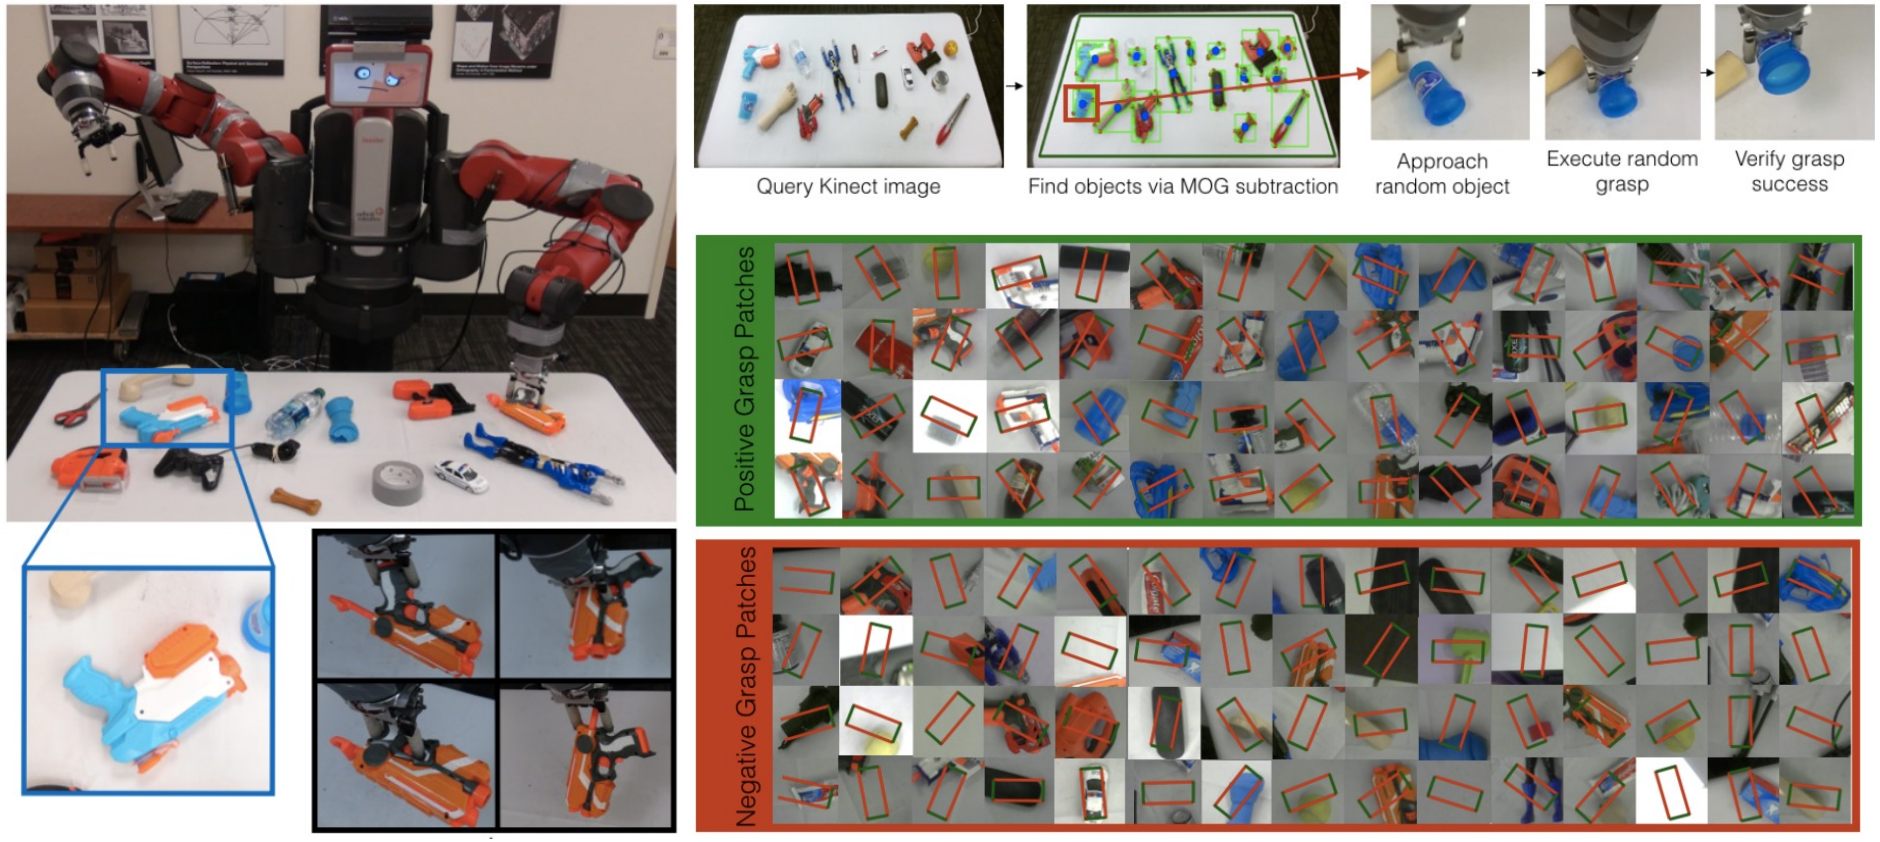

L. Pinto and A. Gupta. Supersizing self-supervision: Learning to grasp from 50K tries and 700 robot hours. ICRA 2016

YouTube video

Lecture overview

- The statistical learning viewpoint
- Supervised classification
- Beyond supervised classification: A brief taxonomy

Beyond classification: Regression

Date prediction

Vittayakorn et al. (2017)

Location prediction

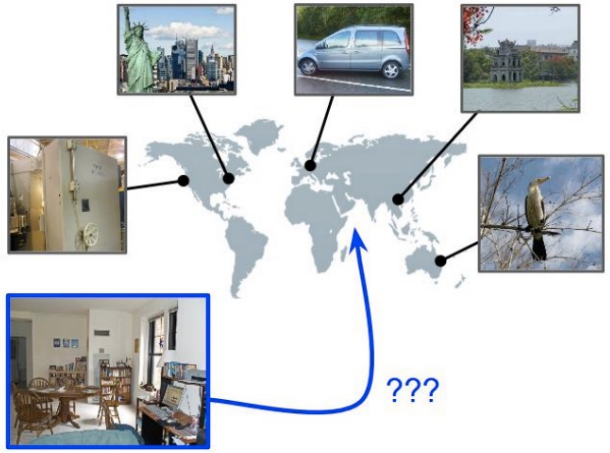

Vo et al. (2017)

Image colorization

Zhang et al. (2016)

Depth prediction

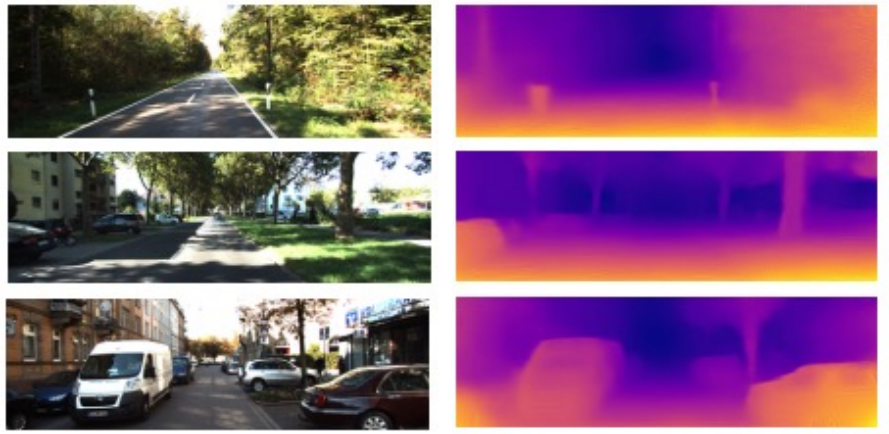

Wang et al. (2017)

Beyond classification: Structured prediction

Image

brace

Word

The screen was a sea of red

Sentence

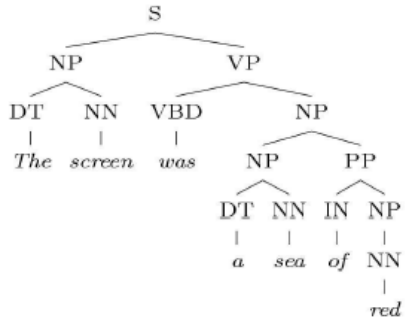

Parse tree

RSCCPCYWGGCPW
GQNCYPEGCSGPKV

Amino-acid sequence

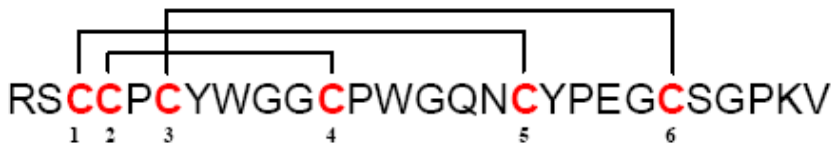

Bond structure

Source: B. Taskar

Structured and dense prediction for scene understanding

Bounding box prediction, dense prediction

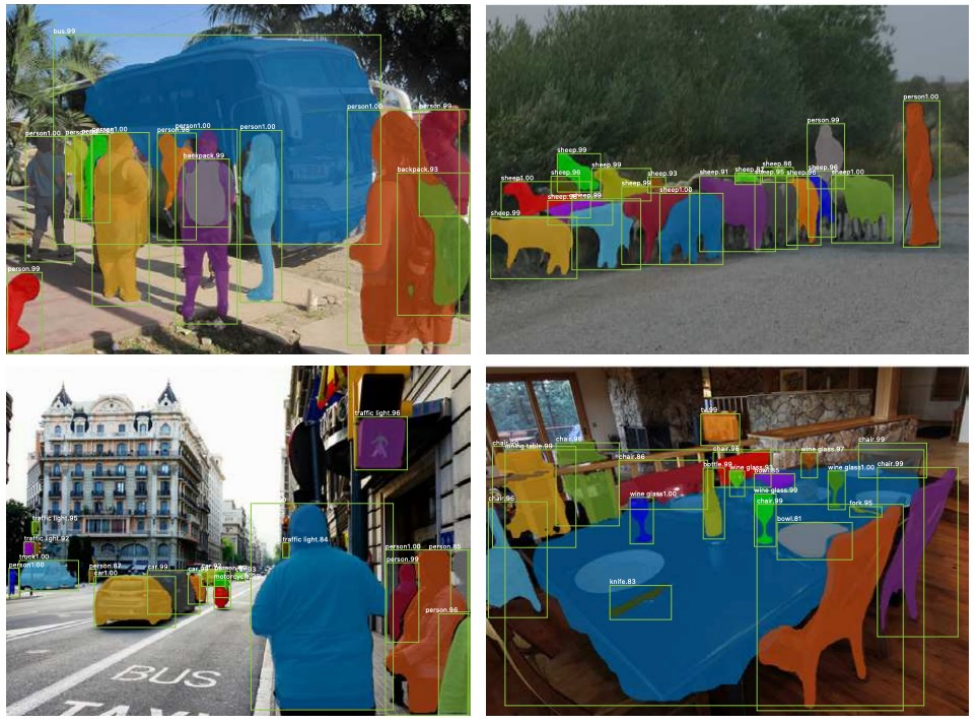

Keypoint prediction

K. He, G. Gkioxari, P. Dollar, and R. Girshick, Mask R-CNN, ICCV 2017

Structured and dense prediction for scene understanding

Image captioning

A. Karpathy, L. Fei-Fei. Deep Visual-Semantic Alignments for Generating Image Descriptions. CVPR 2015

Scene graph generation

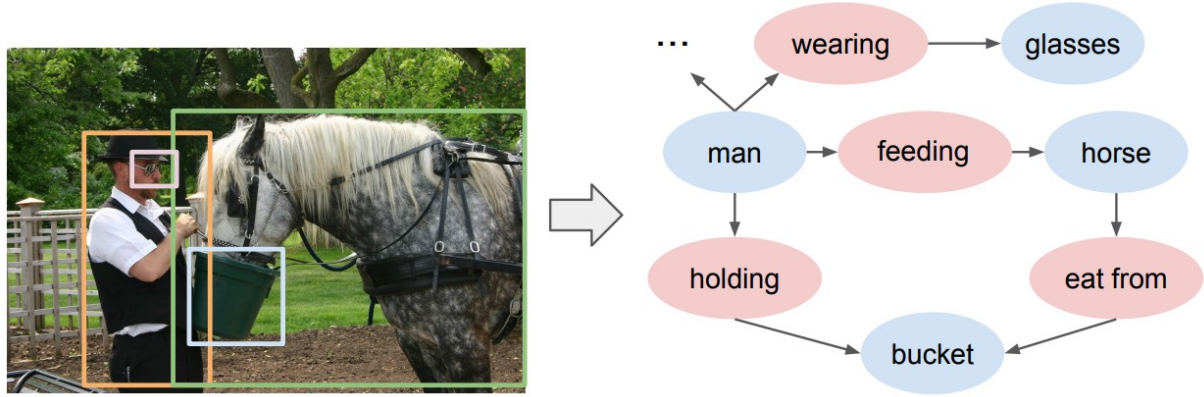

D. Xu, Y. Zhu, C. Choy, and L. Fei-Fei. Scene Graph Generation by Iterative Message Passing. CVPR 2017

Beyond classification and supervised learning

- Other prediction scenarios (output types)
  - Regression
  - Structured prediction
  - Dense prediction
- Other supervision scenarios
  - Unsupervised learning
  - Self-supervised or predictive learning
  - Reinforcement learning
  - Active learning
  - Lifelong learning

Unsupervised learning

- **Clustering**
  - Discover groups of “similar” data points

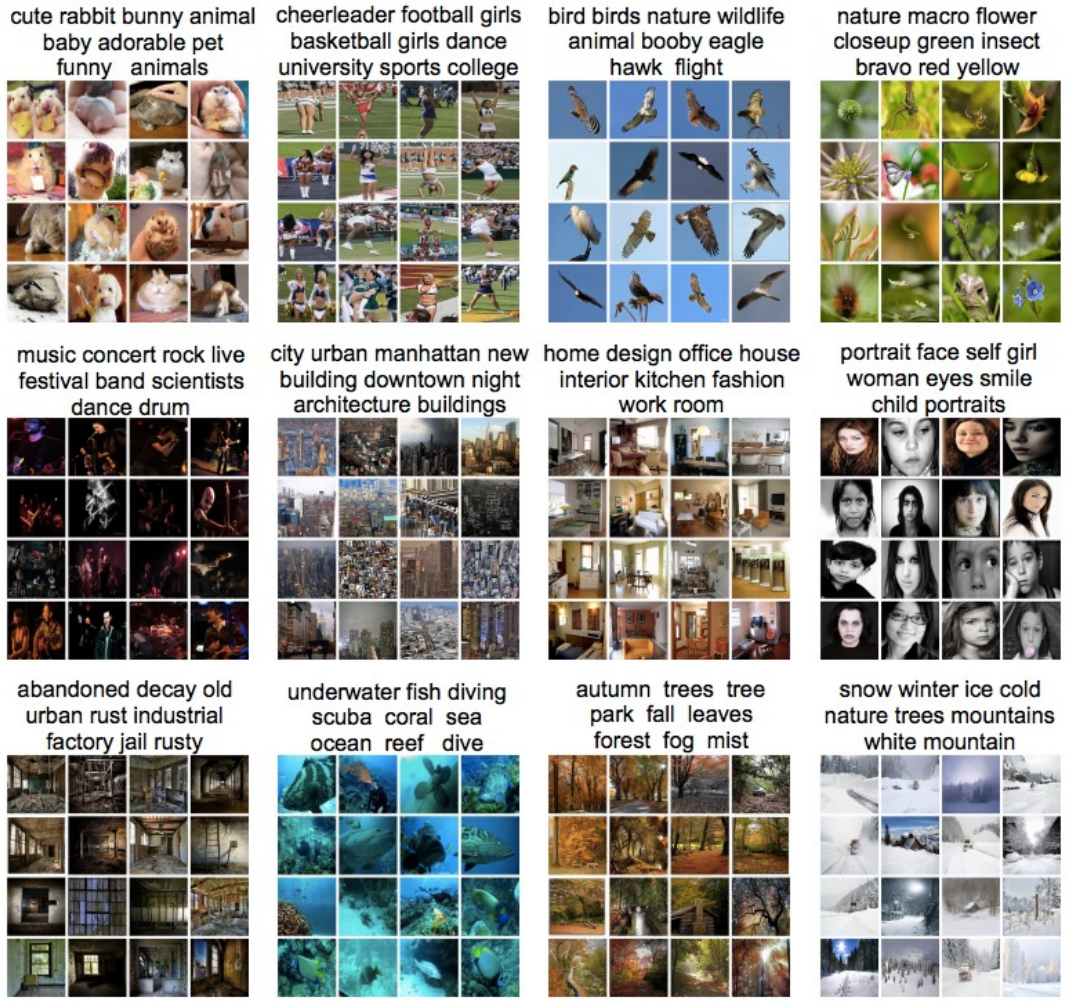

Y. Gong, Q. Ke, M. Isard, and S. Lazebnik. A Multi-View Embedding Space for Modeling Internet Images, Tags, and Their Semantics. IJCV 2014

Unsupervised learning

- **Dimensionality reduction, manifold learning**
  - Discover a lower-dimensional surface on which the data lives

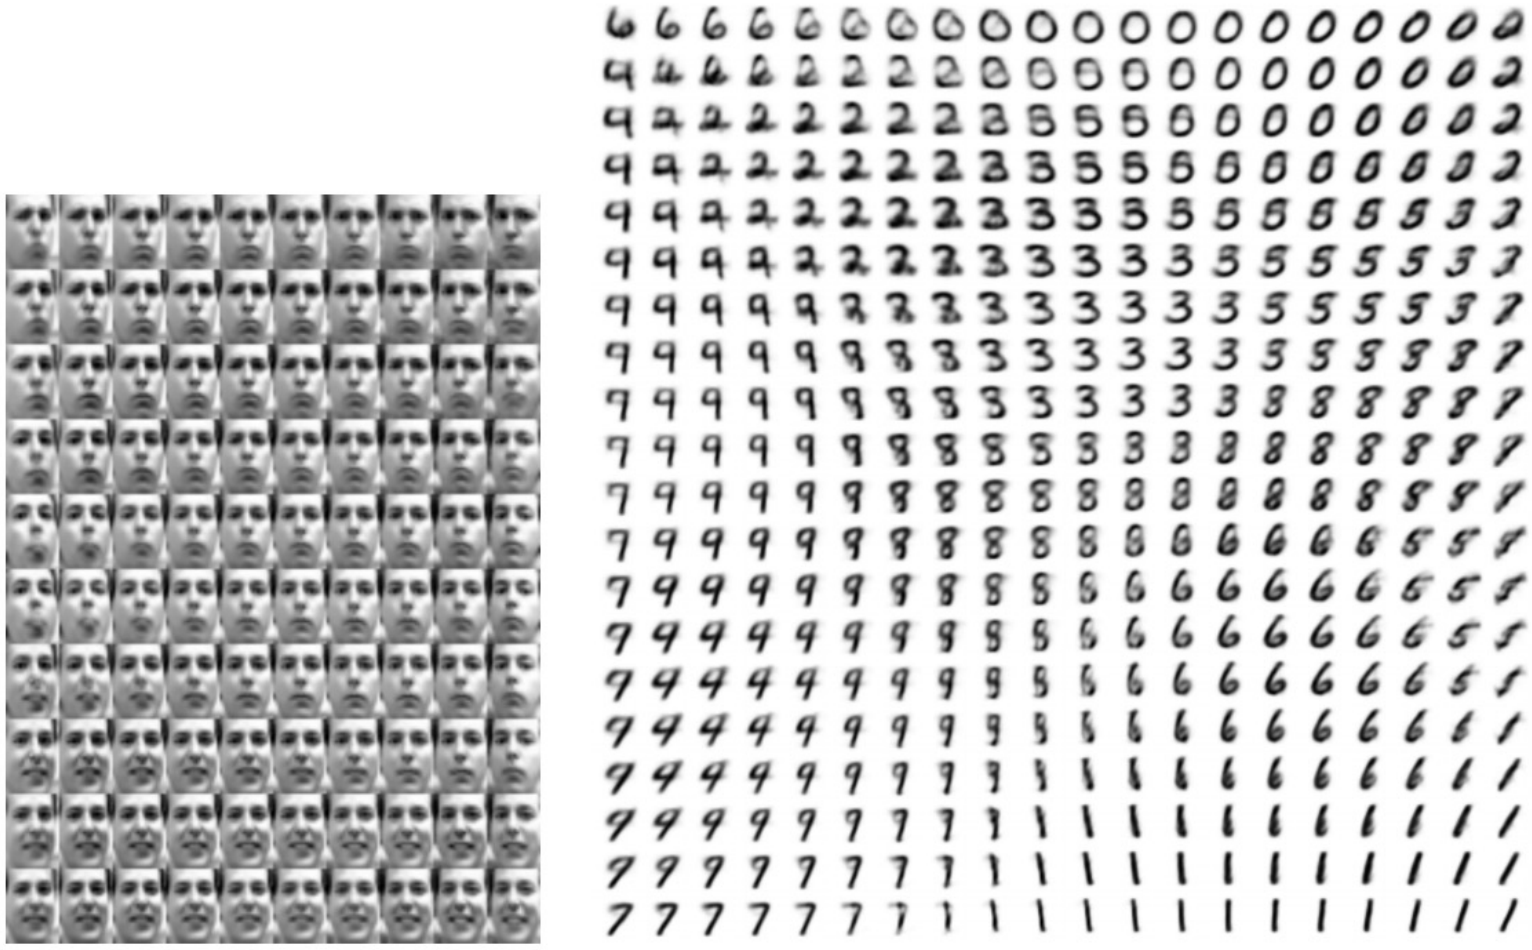

D. Kingma and M. Welling, Auto-Encoding Variational Bayes, ICLR 2014

Unsupervised learning

- Learning the data distribution
  - **Density estimation**: Find a function that approximates the probability density of the data (i.e., value of the function is high for “typical” points and low for “atypical” points)
  - An extremely hard problem for high-dimensional data…

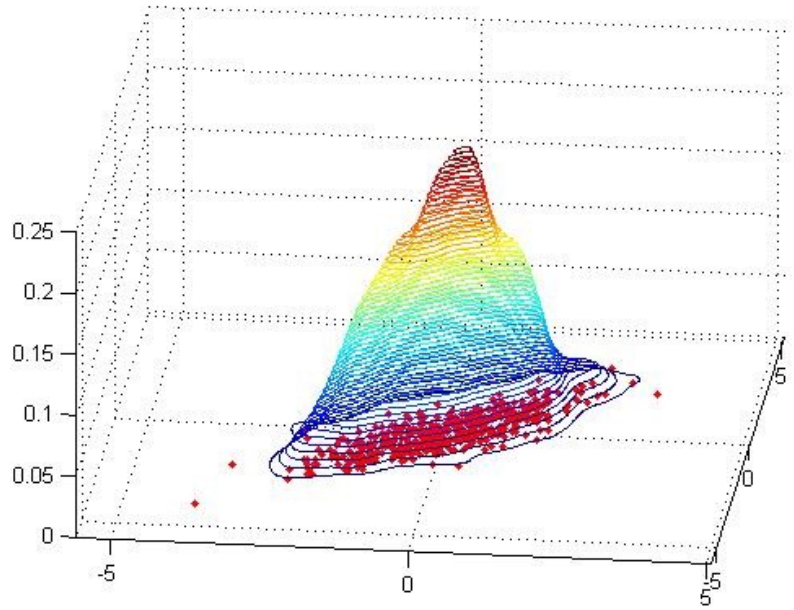

Unsupervised learning

- Learning the data distribution
  - **Learning to sample:** Produce samples from a data distribution that mimics the training set

Generative adversarial networks

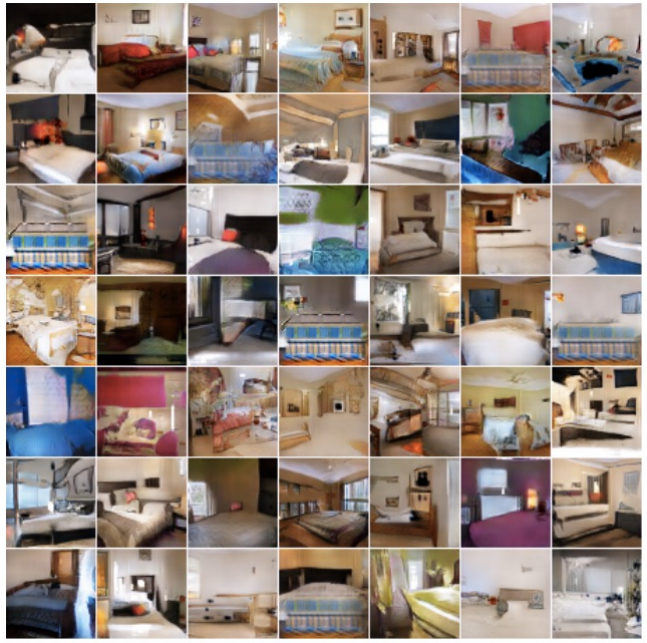

"Bedroom"
(circa 2015)

"Face"
(circa 2015)

Beyond classification and supervised learning

- Other prediction scenarios (output types)
  - Regression
  - Structured prediction
  - Dense prediction
- Other supervision scenarios
  - Unsupervised learning
    - Clustering and quantization
    - Dimensionality reduction, manifold learning
    - Density estimation
    - Learning to sample

Between “unsupervised” and “fully supervised”

Semi-supervised
(labels for a small portion of training data)

Unsupervised
(no labels)

Weakly supervised
(noisy labels, labels not exactly for the task of interest)

Supervised
(clean, complete training labels for the task of interest)

Beyond classification and supervised learning

- Other prediction scenarios (output types)
  - Regression
  - Structured prediction
  - Dense prediction
- Other supervision scenarios
  - Unsupervised learning
  - Self-supervised or predictive learning
  - Reinforcement learning
  - Active learning
  - Lifelong learning

Self-supervised or predictive learning

- Use part of the data to predict other parts of the data
  - Example: Image colorization

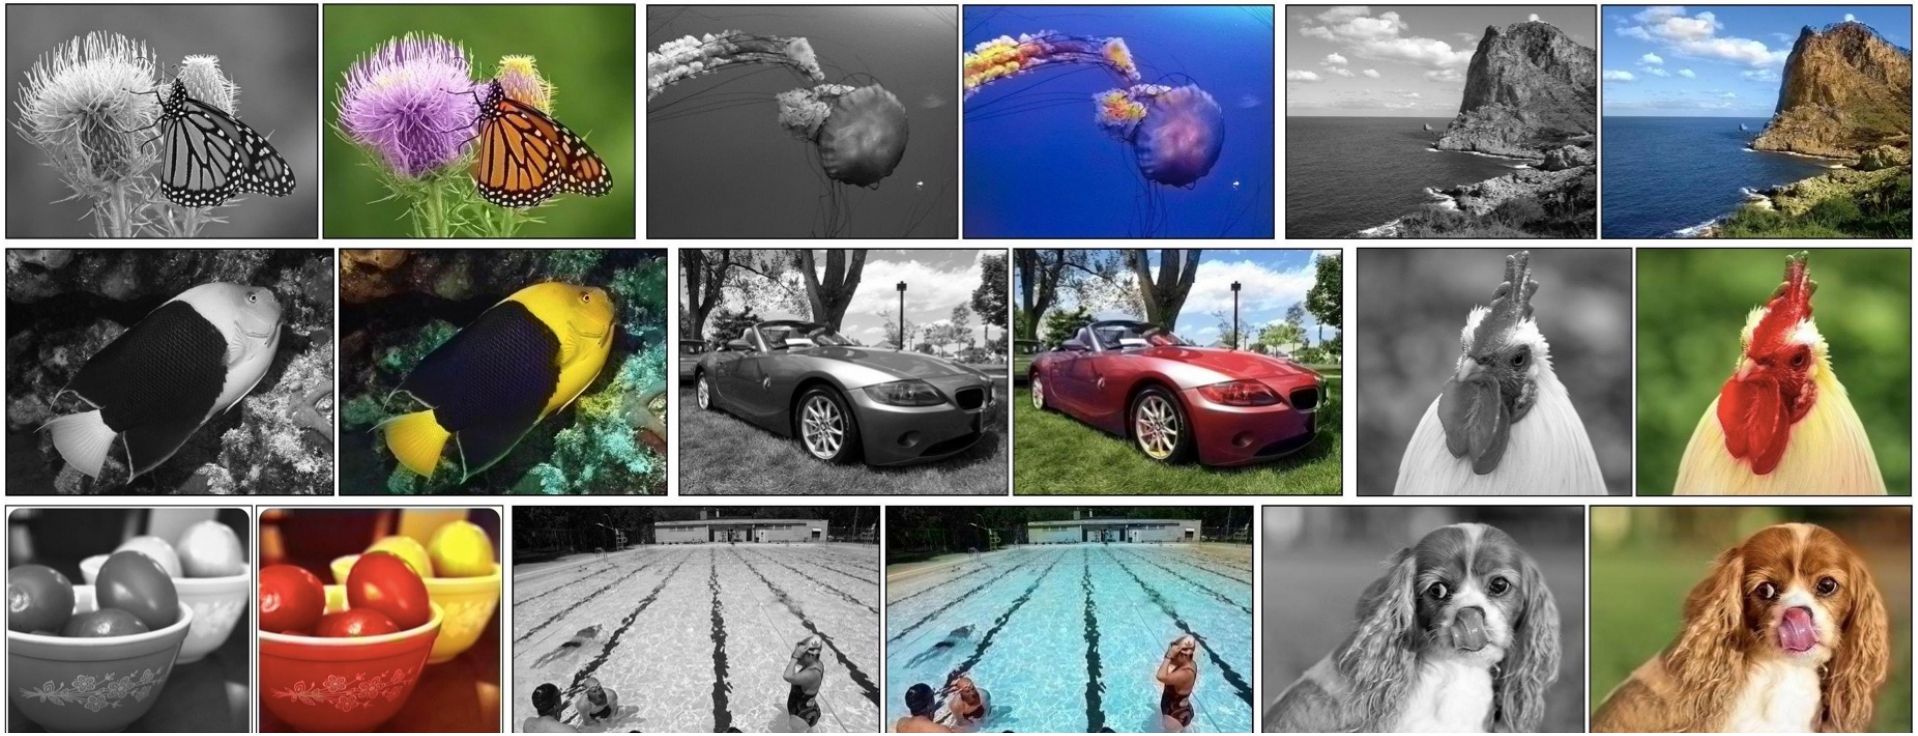

R. Zhang et al., Colorful Image Colorization, ECCV 2016

Self-supervised or predictive learning

- Use part of the data to predict other parts of the data
  - Example: Future prediction

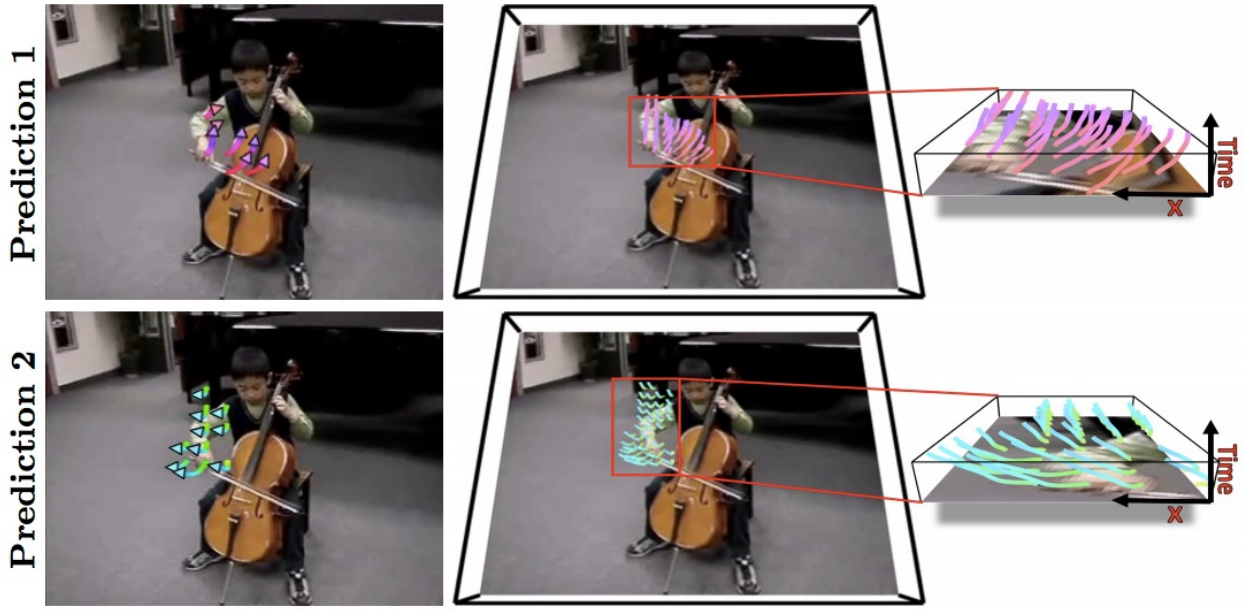

J. Walker et al. An Uncertain Future: Forecasting from Static Images Using Variational Autoencoders. ECCV 2016

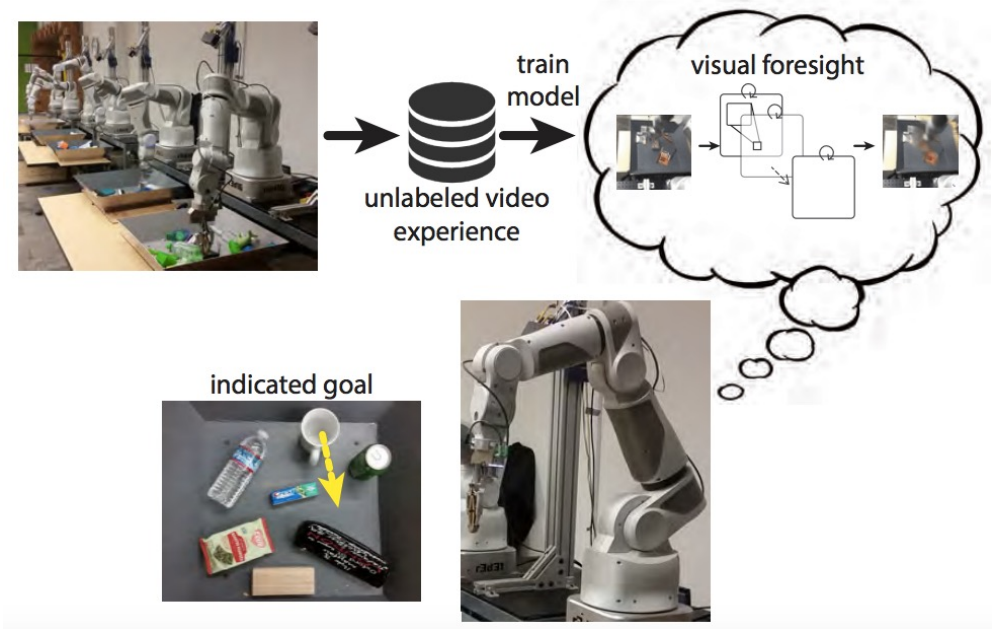

C. Finn and S. Levine. Deep Visual Foresight for Planning Robot Motion. ICRA 2017. YouTube video

Reinforcement learning

- Learn from rewards in a *sequential* environment

Initial gait

Learned gait

N. Kohl and P. Stone. Policy Gradient Reinforcement Learning for Fast Quadrupedal Locomotion. ICRA 2004

Reinforcement learning

Learn from rewards in a *sequential* environment

Sensorimotor learning

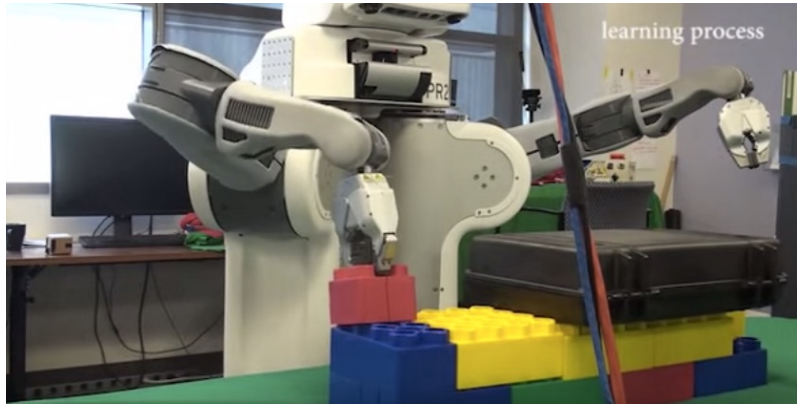

DeepMind's AlphaGo

DeepMind's Atari system

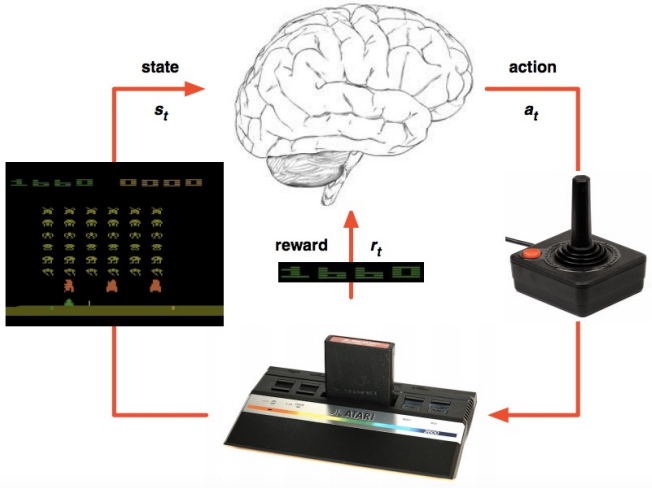

Active learning

- The learning algorithm can choose its own training examples, or ask a “teacher” for an answer on selected inputs

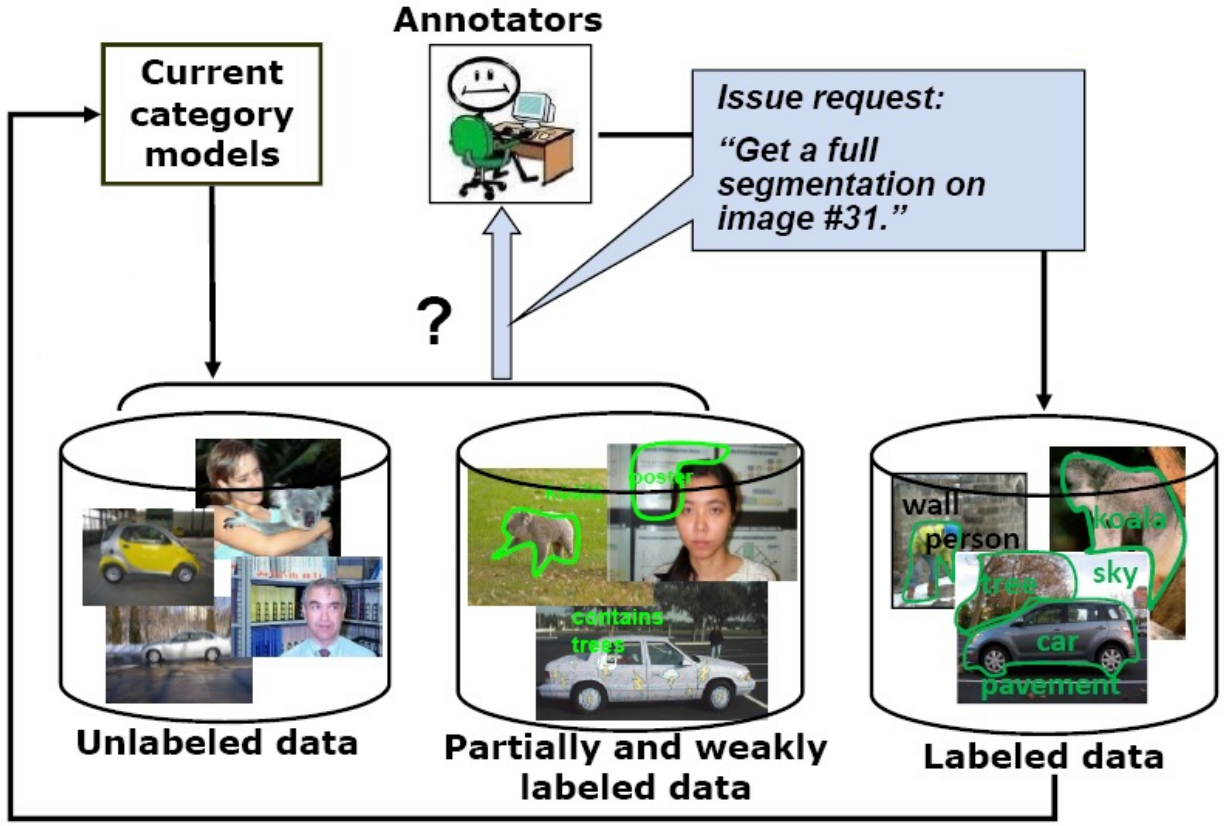

S. Vijayanarasimhan and K. Grauman. Cost-Sensitive Active Visual Category Learning. IJCV 2010

Lifelong or continual learning

Figure 1: **Wanderlust**: Imagine an embodied agent is walking on the street. It may observe new classes and old classes simultaneously. The agent needs to learn fast given only a few samples (red) and recognize the subsequent instances of the class once a label has been provided (green). In this work, we introduce a new online continual object detection benchmark through the eyes of a graduate student to continuously learn emerging tasks in changing environments.

J. Wang et al. Wanderlust: Online Continual Object Detection in the Real World. ICCV 2021

Beyond classification and supervised learning

- Other prediction scenarios (output types)
  - Regression
  - Structured prediction
  - Dense prediction
- Other supervision scenarios
  - Unsupervised learning
  - Self-supervised or predictive learning
  - Reinforcement learning
  - Active learning
  - Lifelong learning

In this class

ML basics, linear classifiers

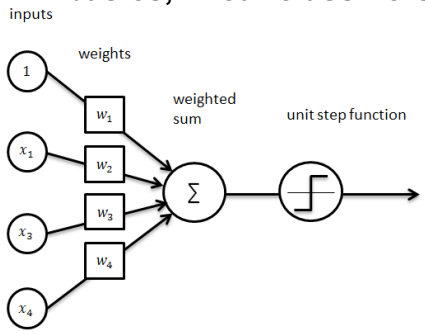

Multilayer neural networks, backpropagation

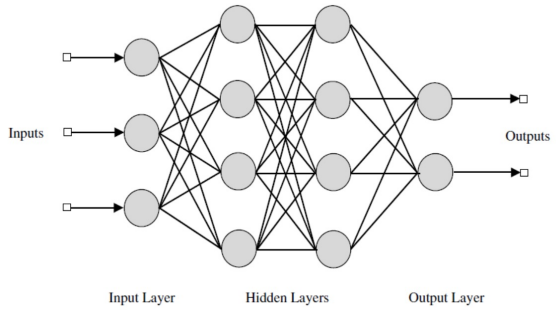

Convolutional networks for classification

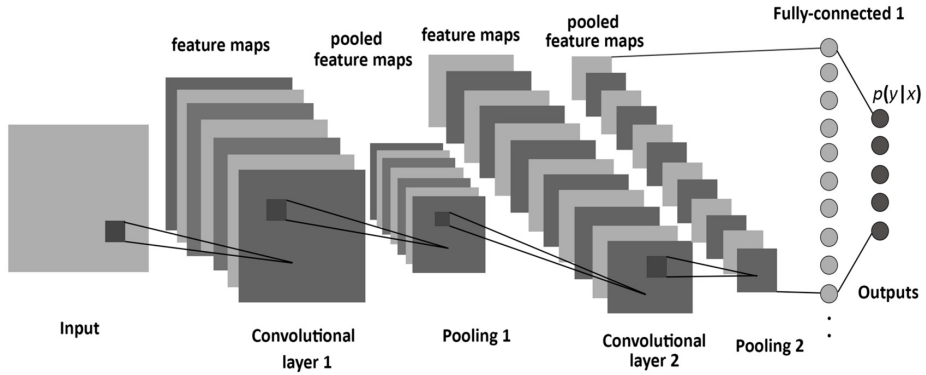

Networks for detection

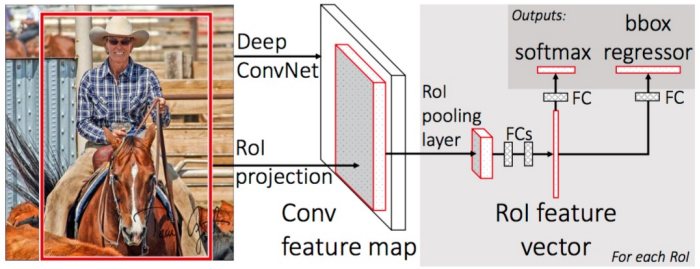

Networks for dense prediction

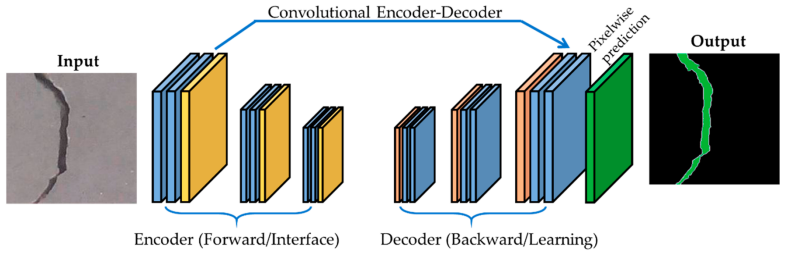

Generative models (GANs, VAEs)

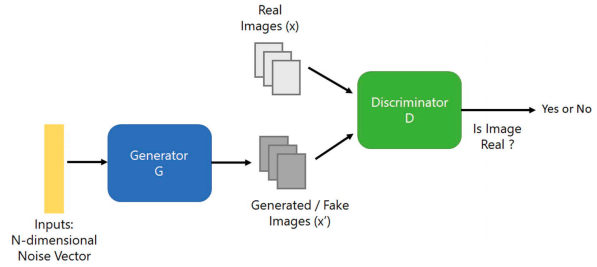

Recurrent models

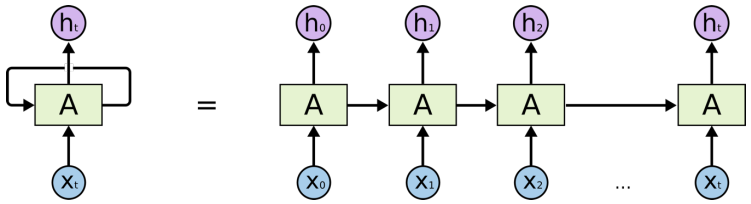

Transformers

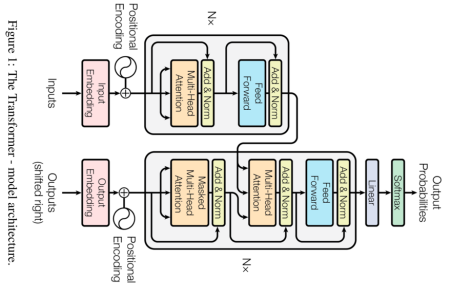

Deep reinforcement learning