A Tour of Machine Learning and its Subdomains

Rémi Emonet
Séminaire H.Curien Département Info/Image
2018-06-07

Overview

- What is Machine Learning (ML)?
- Some subdomains of ML

Overview

- What is Machine Learning (ML)?
- Some subdomains of ML

AI? Machine Learning? Deep Learning?

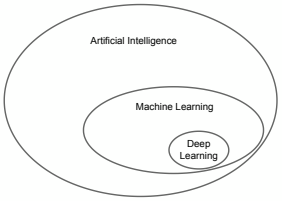

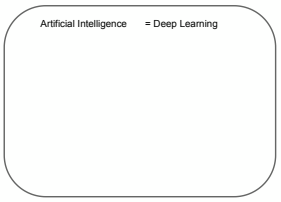

(in the news)

> *Machine learning is a subset of artificial intelligence in the field of computer science that often uses statistical techniques to give computers the ability to "learn" (i.e., progressively improve performance on a specific task) with data, without being explicitly programmed.* wikipedia

ML is your plain old curve fitting

(simplified) Machine Learning Principle

- Considering an input space and an output space
  - e.g. : the space of images
  - e.g. : an integer reprensenting the age of the person
  -
- Given a labeled training set
  -
  - set of inputs with their corresponding outputs
- Learn a "model"
  - that is able to predict the output from an input
  - that will be used on new unseen data
  - i.e., when we get a new input , predict
- NB: it is simplified (many settings don't fit here)

Task example: image classification

- : color image
- : an integer from 0 to 999
  - possible classes
  - incl. cat, dog, boat, ...

Task example: image (semantic) segmentation

- : an image
- : an segmentation
  - for every pixel, a class label
  - e.g., including car, person, tree, ...

Task example: tracking

- : videos
- : trajectory of object(s): sequence of position, shape, attributes, ...

Task example: credit card fraud detection

- x X: transaction with amount, credit history, ...
- y Y: whether the transaction is fraudulent

Task example: age estimation

- $x \in X$: an image of a person
- $y \in Y$: the age of the person

Task example: face verification

- $x \in X$: two face images
- $y \in Y$: whether they are of the same person

Task example: voice detection

- $x \in X$: an audio stream
- $y \in Y$: a segmentation voice/non-voice (or voice segments)

Task example: speech recognition

- $x \in X$: an audio stream
- $y \in Y$: the spoken text

Task example: many more...

- Goal: learn to generalize on unseen data
- Study of the soundness of approaches and their guarantees (learning theories)
- Propose new formalisms, algorithms, generalization bounds, approximations, ...

What is Machine Learning?

Overview

- What is Machine Learning (ML)?
- Some subdomains of ML

About the supervision signal

Supervised vs Unsupervised

Supervised
(known labels)

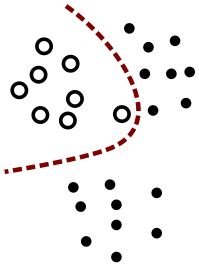

Unsupervised
(only inputs)

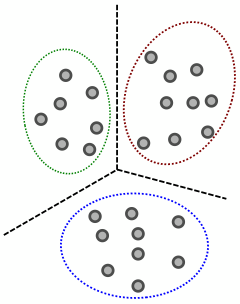

What kind of supervision do we get?

- Supervised, semi-supervised, weakly, webly, unsupervised
- Batch learning, incremental learning, sequential learning, active learning
- Reinforcement learning

About Transfer Learning

Transfer Learning: Multi-* Learning

Multi-Task Learning

- Covered a lot in a recent summer school talk
- (At least), different output for each task, e.g.,
  - different classification task: dog-vs-cat and domestic-vs-wild
  - different output kind: image segmentation and image classification
  - ...

Multi-View Learning

- Input have multiple views, e.g.
  - different viewpoints of an object
  - multi-modal perception (auditory and visual)
  - different medical tests on a patient
  - different sets of features extracted from images
  - ...
- There could be missing views for some input data

(we'll come back to this)

Multi-domain Learning?

Domain Adaptation: What and Why?

When do we need Domain Adaptation (DA)?

- The training distribution is different from the testing distribution

Example Domain Adaptation task?

- Given: labeled images (e.g., fruits images)
- Task: what fruit appears on this unlabeled images of trees

Blueberry Almond Blueberry Almond

- How can we learn, from one distribution, a low-error classifier on another distribution?

Attribution

By Sewaqu - Own work, Public Domain

CC by genevieveromier (Flickr)

CC by anmuell (Flickr)

CC by sgillies (Flickr)

CC by mustetahra (Flickr)