Integrated Learning

Gal A. Kaminka

May 7, 2018

Learning and Intelligence

Do you need learning to be intelligent?

Learning and Intelligence

Do you need learning to be intelligent?

- No.

Learning and Intelligence

Do you need learning to be intelligent?

- No.
- But it helps.

Learning and Intelligence

Do you need learning to be intelligent?

- No.
- But it helps.

Why?

Why Learning is Important for Intelligence? (I)

Speedup

- Some (sub-) problems repeat
- If solved problem before, *not* learning is *not* intelligent
  - Why? (good question, e.g., for test)
- Law of Practice is relevant here

Why Learning is Important for Intelligence? (II)

Control problem (action execution) is tricky

- Effects, preconditions not known, or incorrect
  - Effects: Neuroscience *Forward Model*
- When have (virtual) body, it "has its own mind" sometimes :-)
- Uncertainty, failures

Why Learning is Important for Intelligence? (III)

Perception is hard

- What to pay attention to
- How to segment environment
- Recognize objects of import
- Recognize other agents (*IFF*)

Why Learning is Important for Intelligence? (IV)

Action selection is challenging

- CHOOSE(): What choices better? when?
- Separate choices in hierarchies:
  - Which decomposition to take (HTN), which child to choose (BIS)
  - Which next action to take (order)
- What part of state is relevant? How much of history?

Where does learning fit in the agent

Whenever there is a potential for knowledge failure

- As followup to CHOOSE()
  - Evaluation happens during execution
  - Can learn effects of CHOOSE()
- In response to TEST failing
  - If no actions selectable: something wrong with preconditions
- In followup to EXECUTE()
  - Effects of actions

Where does learning fit in the agent

Whenever there is a potential for knowledge failure

- As followup to CHOOSE()
  - Evaluation happens during execution
  - Can learn effects of CHOOSE()
- In response to TEST failing
  - If no actions selectable: something wrong with preconditions
- In followup to EXECUTE()
  - Effects of actions

Missing (for me):

- Not clear where perception learning occurs
- Learning by demonstration, learning by instruction?

General issues with integrated learning

- Utility problem: retrieving learned knowledge becomes slower
- Exploration vs Exploitation,
- Domain-model vs model-free

Examples of Learning in Agents

- Learning for action selection
  - Model-free reinforcement learning
  - Chunking, explanation-based learning
- Learning for perception
  - learning sensor models
  - recognize objects, etc. (e.g., using deep learning)
- Learning action models
  - Preconditions, effects (e.g., using regression, deep learning)
- Learning model of domain
  - Model-based reinforcement Learning
  - Grammar induction
  - learning affordances
  - Learning new actions, action sequences

Machine Learning (a sub-field of AI)

Learning in Machines

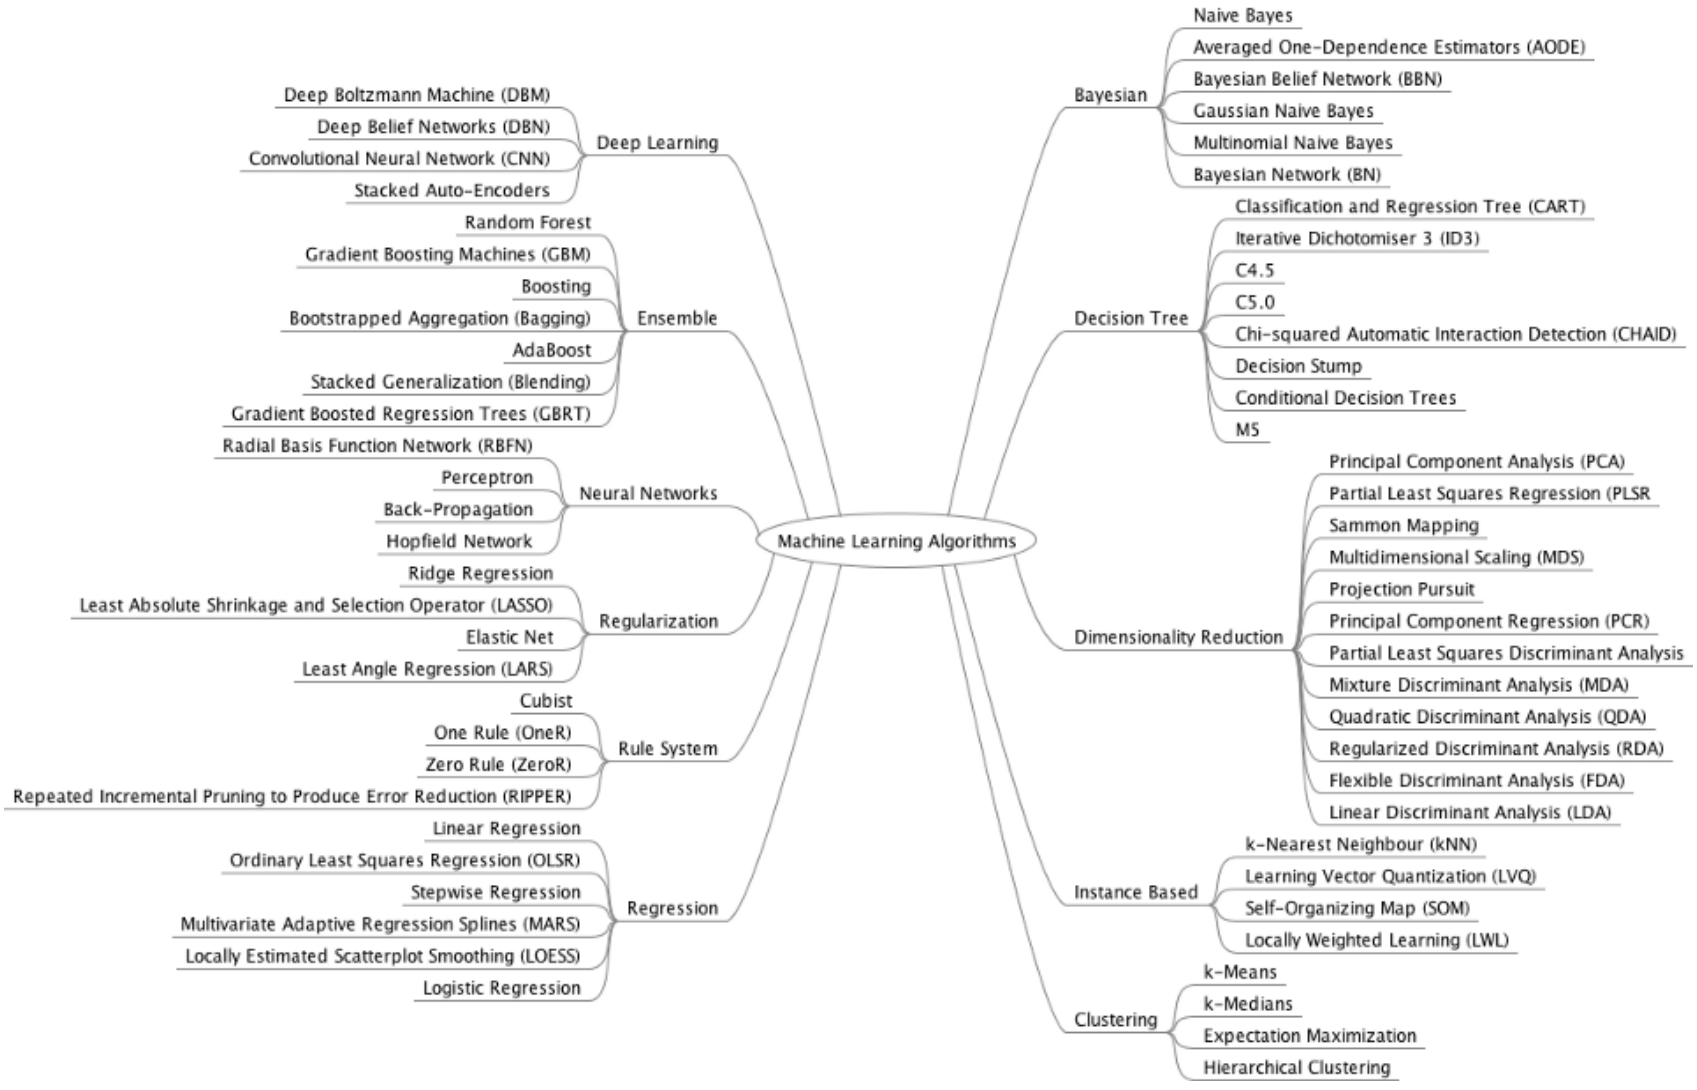

Figure 1: machinelearningmastery.com/a-tour-of-machine-learning-algorithms/

Most common

- Supervised:
  - Input: Examples with goal classes/values/structures
  - Output: a procedure to predict goal class given new input
- Unsupervised: Examples/data without any labels
  - Output: "interesting" patterns or substructures, anomalies
  - Output: abstractions, reduced dimensions
  - Output: model parameters (e.g., in HMMs)

Others: genetic programming, grammar induction, ...

Supervised Learning

Learning Algorithm

- Input: Examples of instances, and their *labels*
- Output: A **classifer**
  - a procedure for predicting label of *new instances*
- Distinct: *generative* or *descriptive* classifiers

Supervised Learning

Learning Algorithm

- Input: Examples of instances, and their *labels*
- Output: A **classifer**
  - a procedure for predicting label of *new instances*
- Distinct: *generative* or *descriptive* classifiers

Example Variants

- Regression (labels: scalar **values**)
- Classification (labels: two or more **categories**)
- Structured Prediction (labels: **structures**)
- Concept learning (single label, positive and negative examples)
- ...

Many of the neural networks/deep learning techniques in this class

Unsupervised Learning

Learning Algorithm

- Input: Examples of data
- Output: Invariants, Regularities, and Anomalies

Unsupervised Learning

Learning Algorithm

- Input: Examples of data
- Output: Invariants, Regularities, and Anomalies

Example Variants

- Clustering, hierarchical clustering (sub-groupings)
- Multivariate anomaly detection
- (Sequential) pattern mining
- Hidden Markov model, Bayes networks learning
- ...

Learning in Intelligent Agents

Types of Learning in Nature

- Learning by instruction
  - Procedures ("do like this")
  - Concepts ("This is a dog")
- Learning by analogy ("what do we learn from this story?")
- Learning by demonstration ("watch what I do")
- Mimicry, imitation ("I didn't know you watch what I do")
- Rote learning (e.g., the multiplication table)
- Learning by practicing, rehearsing
  - ("Practice makes perfect!")
- Reinforcement Learning ("Ouch/Yummy!")

Learning in Humans: The Law of Practice

- Practice improves: rapidly at first, then less so
- Response time decreases with number of practice trials
- Newell and Rosenbloom (1981): *Power Law of Practice*

$$Time = \alpha N^{-\beta} + c$$

- Heathcote, Brown, Mewhort (2000): *Exponential* Law

$$Time = \alpha e^{-\beta(N-1)} + c$$

But that's not what is interesting here

The **AMAZING** Law of Practice

- Regardless of task

The **AMAZING** Law of Practice

- Regardless of task
- Regardless of test subject

The **AMAZING** Law of Practice

- Regardless of task
- Regardless of test subject
- Regardless of intention to learn

It is a side-effect of how our mind is built to learn