Introduction to Machine Learning

Yu Cheng

EECS Northwestern

The Learning Problem - Outline

- Example of machine learning
- Components of learning
- Types of learning
- The road map of learning
- Conclusion

Example: credit approval

Deciding approve an application or not?

Application information:

item	value
age	23 years
gender	male
annual salary	$30,000
years in residence	1 year
years in job	1 year
current debt	$1000
…	…

Example: credit approval

Machine learning: find a function *f(data)* to predict the result (1,-1).

The essence of machine learning:

- Some data are observed
- A pattern exists
- We can not pin it down mathematically

Example: credit approval

Formalization:

Input: $\mathbf{x}$ (customer application)

Output: y (good/bad customer)

Target function: $f : X \rightarrow Y$ (ideal credit approval formula)

Data: $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_N, y_N)$ (historical records)

Hypothesis: $g : X \rightarrow Y$ (formula to be used)

Learning: find g that well approximates f

More applications of machine learning

- Spam (email spam…)
- Speech recognition (Siri…)
- Netflix
- Handwriting
- Climate modeling
- …

Components of learning

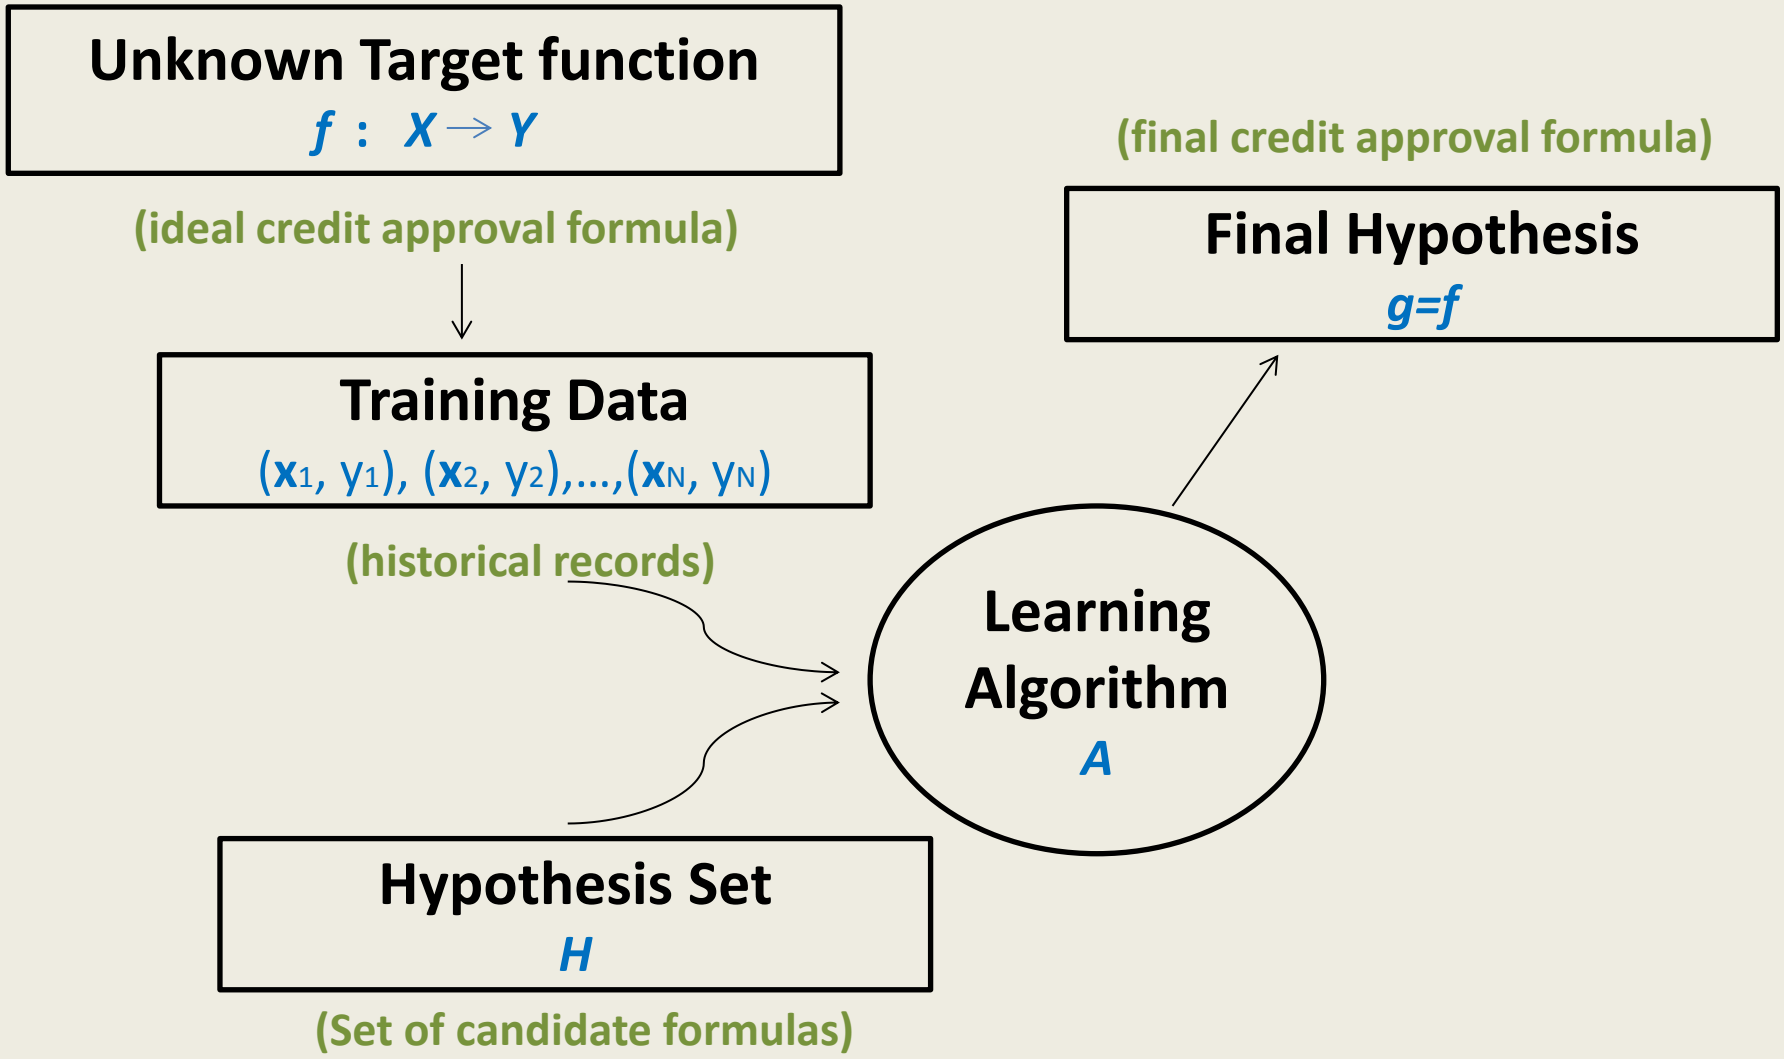

Solution components

Solutions components of the learning problem:

- The hypothesis set
  $H = \{h\}$ $\quad g \in H$
- The learning algorithm

Together, they are referred to as the *learning model*.

A simple case – "perceptron"

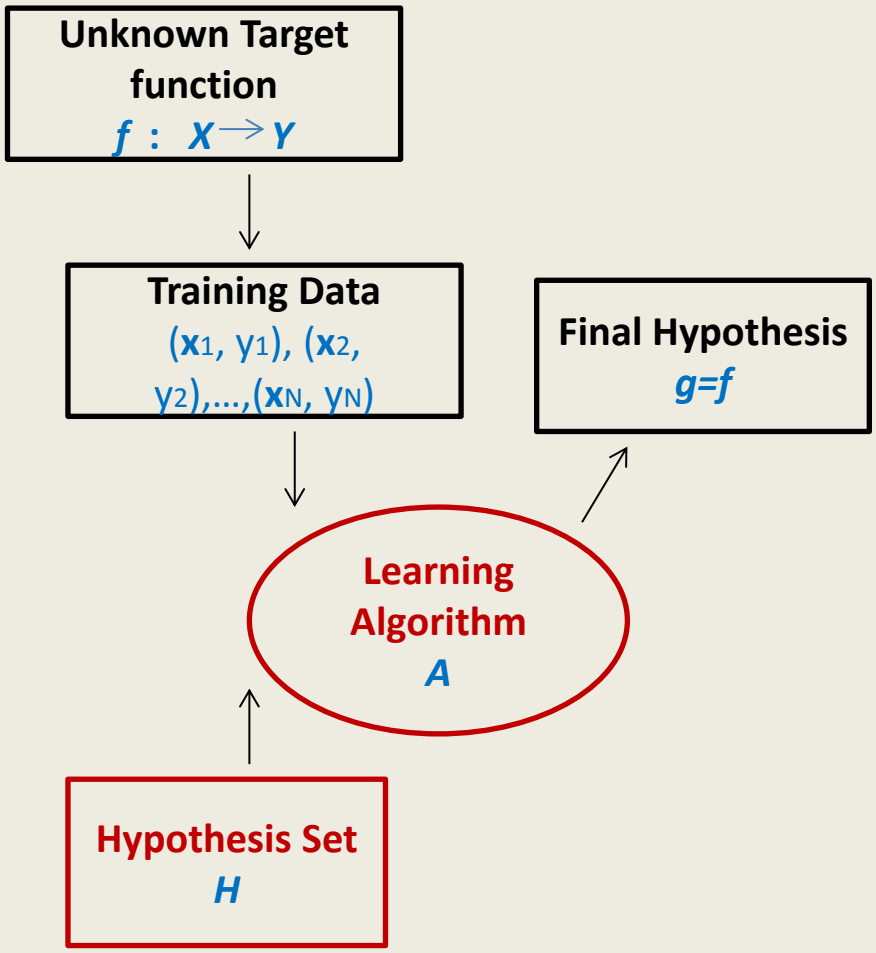

Basic premise of learning

What is machine learning?

"Using a set of observations to uncover an underlying process"

- Yaser Abu-Mostafaz

"The study of computer programs that improve automatically with experience"

- Tom M. Mitchell

"Algorithms for inferring unknowns from knowns"

broad premise $\longrightarrow$ many variations

Types of Learning

Supervised Learning

(input, correct output)

Two paradigms: classification, regression

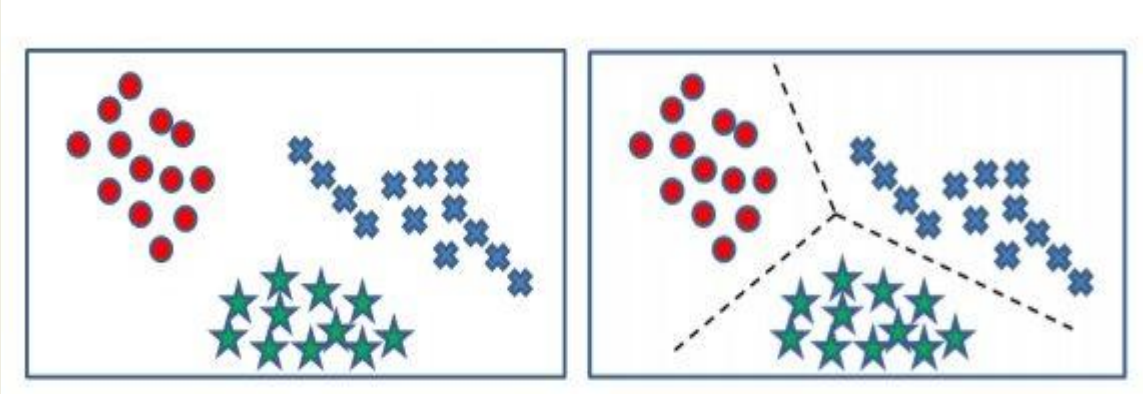

From Alex's Tutorial

Types of Learning

Unsupervised Learning

Instead of (input, correct output), we get (input, ?)

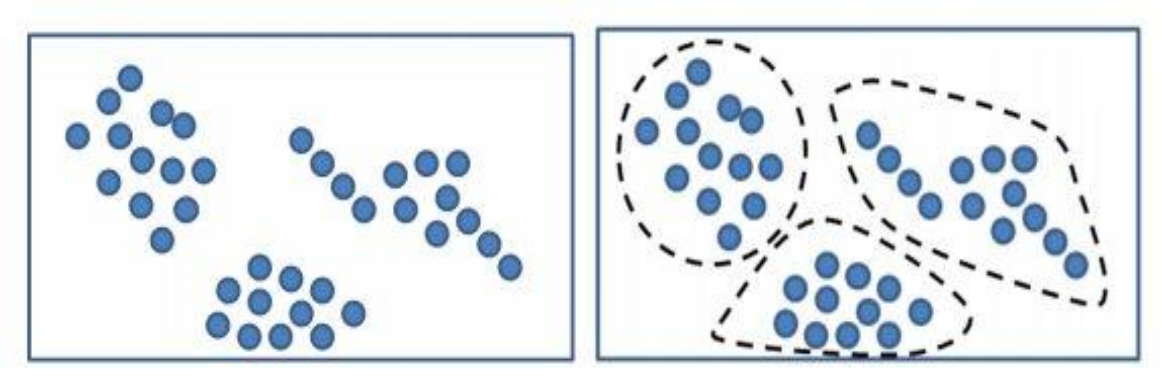

From Alex's slides

Types of Learning

- Passive learning VS Active learning
- Online learning VS Batch learning
- Generative learning VS Discriminative learning

It’s a jungle out there

supervised learning

batch learning

VC

liner regression

Bayesian

SVM

online learning

neural network

complexity

Naive Bayesian

unsupervised learning

active learning

Gaussian process

cross-validation

Bayesian net

reinforcement learning

bias-variance

liner regression

decision tree

discriminative learning

The map of machine learning

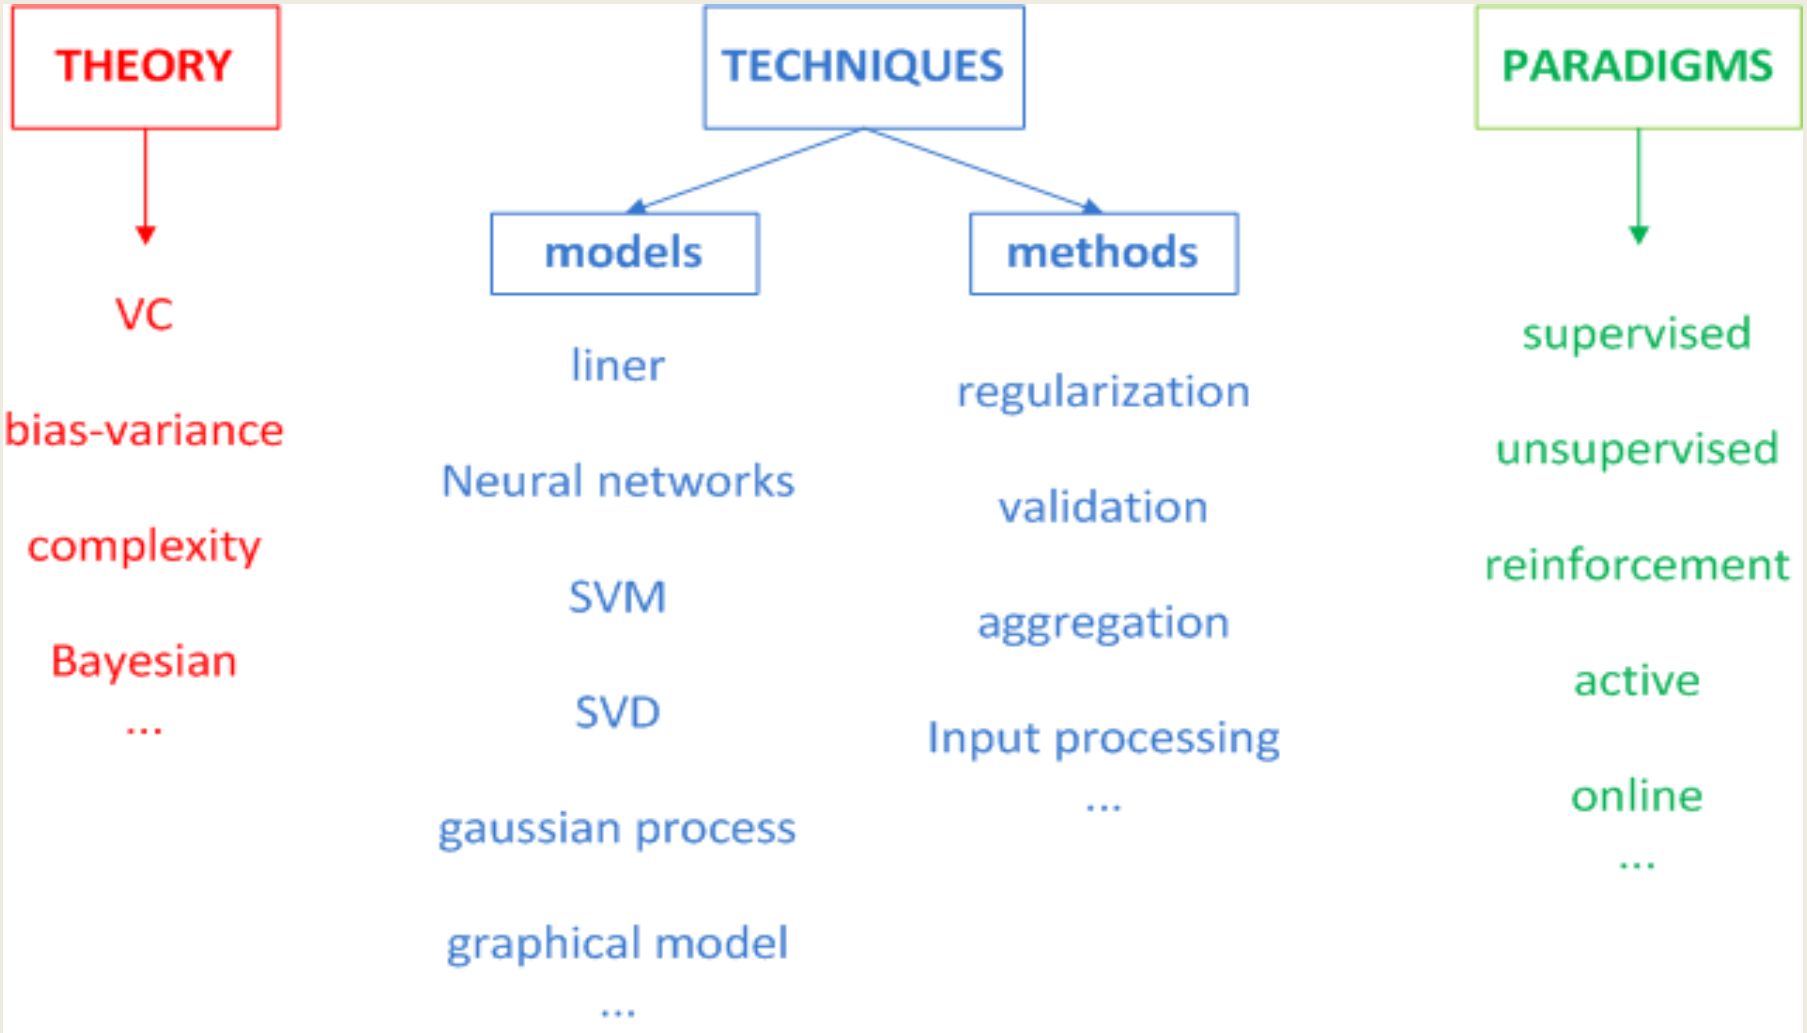

Conclusion – Some references

- Learning from data, by Yaser Abu-Mostafa, 2012.
- Pattern Recognition and Machine Learning, Christopher M. Bishop, 2007.
- Pattern Classification, Peter E. Hart et al, 2000.

- Machine Learning, Tom M. Mitchell, 1997
- The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Trevor Hastie et al, 2009

Thanks a lot!

Course Instructor: Dr. Downey

Slides from Doug Downey ,Alex Smola and Yaser Abu-Mostafa