

Applied Machine Learning

Applied machine learning (EDAN95)
Lecture 01 — Introduction to the course and the field
2018-11-05
Elin A. Topp

Goodfellow chapter 1
various sources

Statistics ... or how you ended up here

We had 192 initial applications for the course in LADOK by the deadline (60+15 places)

- 14 eligible exchange students (confirmed and agreed upon with the IO)
- 166 eligible students from programs D, C, BME, F, Pi
- 7 (eligible) students from programs not listing the course
- 5 students not fulfilling the prerequisites (from both types of programs)

We calculated (and recalculated every so often after discovering new issues) for all eligible except exchange students:

- number of credits taken within the study program one applied through
- including granted credits, credits taken in unfinished courses and including not yet granted but potentially valid credits from external courses and sources (if visible to us)

We subtracted this number from the required number of credits for the respective program. After some withdrawals, iterated adjustments, opening up another 9 places by maximising to the very last place in the lab sessions, being overruled in the last minute by LU administration, and finally (luckily) four more withdrawals, we had on Friday afternoon

84 students on **84** lab places - hopefully you are not afraid of getting close to other people!

What is Machine Learning?

ML is an area of (AI) research providing *powerful tools* that enable machines (computers) to find models describing data and the correlations between them, to

- *predict* future outcomes or developments given previous data (weather, stock market)
- *classify* unknown input given known, classified data (scene contains pedestrian or not)
- *identify structures* in unseen, unlabeled data (grouping people according to different attributes)
- *decide* upon next *steps or actions* to take to maximise reward (new measurement, robot action)

Some cool videos

Learning to flip a pancake

<http://kormushev.com/videos/robot-motor-skill-coordination-with-em-based-reinforcement-learning/>

Learning to play golf

https://www.youtube.com/watch?v=6o8was_tZZ4

Learning hand-eye coordination

<https://www.youtube.com/watch?v=V05SuCSRAtg>

Learning to sense a snapfit

<https://www.youtube.com/watch?v=TEIq5lQr4nk>

A success story

Silver et al, Nature, Oct 18, 2017:

“Mastering the game of Go without human knowledge”

A long-standing goal of artificial intelligence is an algorithm that learns, *tabula rasa*, superhuman proficiency in challenging domains. Recently, AlphaGo became the first program to defeat a world champion in the game of Go. The tree search in AlphaGo evaluated positions and selected moves using deep neural networks. These neural networks were trained by supervised learning from human expert moves, and by reinforcement learning from self-play. Here we introduce an algorithm based solely on reinforcement learning, without human data, guidance or domain knowledge beyond game rules. AlphaGo becomes its own teacher: a neural network is trained to predict AlphaGo’s own move selections and also the winner of AlphaGo’s games. This neural network improves the strength of the tree search, resulting in higher quality move selection and stronger self-play in the next iteration. Starting *tabula rasa*, our new program AlphaGo Zero achieved superhuman performance, winning 100–0 against the previously published, champion-defeating AlphaGo.

...

What is Machine Learning NOT?

ML is NOT

- the holy grail
- the sole supporting pillar of AI
- the hammer that helps with all nails (Go is complex, but finite and observable)
- magic - although it might appear that way

Word of caution #1

Robert A. Burton, New York Times, May 22, 2017:
“**Donald Trump, Our A.I. President**”

...

If conventional psychology isn't up to the task, perhaps we should step back and consider a tantalizing sci-fi alternative — that Trump doesn't operate within conventional human cognitive constraints, but rather is a new life form, a rudimentary artificial intelligence-based learning machine. When we strip away all moral, ethical and ideological considerations from his decisions and see them strictly in the light of machine learning, his behavior makes perfect sense.

Consider how deep learning occurs in neural networks such as Google's Deep Mind or IBM's Deep Blue and Watson. In the beginning, each network analyzes a number of previously recorded games, and then, through trial and error, the network tests out various strategies. Connections for winning moves are enhanced; losing connections are pruned away. The network has no idea what it is doing or why one play is better than another. It isn't saddled with any confounding principles such as what constitutes socially acceptable or unacceptable behavior or which decisions might result in negative downstream consequences.

...

Word of caution #2

Nguyen, Yosinski, Clune, CVPR 2015 / AAAI Video Competition 2016:

“Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images”

...

Here we show a related result: it is easy to produce images that are completely unrecognizable to humans, but that state-of-the-art DNNs believe to be recognizable objects with 99.99% confidence (e.g. labeling with certainty that white noise static is a lion). Specifically, we take convolutional neural networks trained to perform well on either the ImageNet or MNIST datasets and then find images with evolutionary algorithms or gradient ascent that DNNs label with high confidence as belonging to each dataset class. It is possible to produce images totally unrecognizable to human eyes that DNNs believe with near certainty are familiar objects, which we call “fooling images” (more generally, fooling examples). Our results shed light on interesting differences between human vision and current DNNs, and raise questions about the generality of DNN computer vision.

...

[<https://www.youtube.com/user/aaaivideocompetition>] (all videos)

[<https://www.youtube.com/watch?v=LO5xERK70ZM>] (this particular contribution)

Word of caution #3

Ribeiro, Singh, Guestrin, ArXiv, Aug 9, 2016:

“Why Should I Trust You?": Explaining the Predictions of Any Classifier

...

Often artifacts of data collection can induce undesirable correlations that the classifiers pick up during training. These issues can be very difficult to identify just by looking at the raw data and predictions. In an effort to reproduce such a setting, we take the task of distinguishing between photos of Wolves and Eskimo Dogs (huskies). We train a logistic regression classifier on a training set of 20 images, hand selected such that all pictures of wolves had snow in the background, while pictures of huskies did not. As the features for the images, we use the first max-pooling layer of Google's pre-trained Inception neural network [25]. On a collection of additional 60 images, the classifier predicts “Wolf” if there is snow (or light background at the bottom), and “Husky” otherwise, regardless of animal color, position, pose, etc. We trained this bad classifier intentionally, to evaluate whether subjects are able to detect it.

...

Why this course?

EDAN95 aims to give you an overview of and explain different ML tools to enable you to

- decide whether the problem or task you face can be tackled using ML
- apply the right ML tool(s) for the right problem(s) in the correct way
- handle data in a suitable way
- reflect upon the consequences of applying these tools
- control the amount of magic in your work

EDAN95 Machine Learning

- The boring part -

- Course website: <http://cs.lth.se/edan95>. Our assumptions are:

The website is sufficient as communication link from us to you, i.e., anything we publish on the website is considered as communicated properly to you.

You will have checked the website before asking anything that you wonder about.

You contact us if you find something missing or discover broken links, etc.

- The course has
 - 14 lectures, Mondays and Wednesdays, 1pm to 3pm.
 - 7 lab sessions with 3 instances each, Thursdays (2x) and Fridays (1x), 1pm-3pm
 - 7 lab assignments to be worked on in pairs and presented in the corresponding lab sessions
 - 1 written exam (with one scheduled retake instance) on January 8, gives final grade
 - 7.5 hp / ECTS credits. 4.5 are given for passing all lab assignments, 3 for passing the exam
 - 4 Teachers (Elin, Jacek, Pierre, and Volker) and 3-4 assistants (Erik, Marcus, (Dennis), (Matthias))
 - The maximum number of participants, and then some...

Compulsory items

- To keep your seat in the course you **NEED** to participate in lab session I actively.
 - Should you fail to sign up for a lab session by the deadline (**WEDNESDAY, Nov 7, 11:00am**), your seat will be given to a student on the waiting list immediately.
 - Should you sign up but fail to show without notification by the end of the session or a valid and confirmable cause (e.g., illness), your seat will be given to a student on the waiting list immediately after your session.
- Lab sessions 2-7 are weakly compulsory: you need to present the solution to the assignments within the lab session instances you signed up for, if not another agreement has been made explicitly. You should try to follow the schedule as closely as possible, but “retake presentations” in later sessions are acceptable within manageable limits - priority is given to students following the schedule.

Compulsory items (2)

To sign up for the lab session instance (one sign-up for all 7 sessions!) do the following by **WEDNESDAY, Nov 7, 11:00am**:

- decide whom to work with (registering as single will pair you up with another single student)
- decide whether you **NEED** either of the session instances, i.e., Thursdays or Fridays, or whether you are **FLEXIBLE** (we need to fill all sessions to the absolute maximum!)
- go to <https://sam.cs.lth.se/LabsSelectSession?occasionId=566> (linked from course website)
- sign up for “Flexible” if you are, or for a fixed session otherwise
- (if flexible, find yourself assigned to a session by Wednesday, Nov 7, 3pm, on the sign-up site)

EDAN95 Applied Machine Learning

- The hopefully less boring part -

Types (supervised, unsupervised, (semi)-supervised) of ML approaches

Problem classes (knowledge discovery, prediction, classification, action learning)

Flavours (connectionistic / neural, statistical, probabilistic)

Some statistical methods

**Clustering
(unsupervised)**

**Logistic regression
(supervised)**

**Decision trees
(supervised)**

Some connectionistic methods

**Perceptron -> MLP
(supervised)**

**Convolutional NN
("semi"-supervised)**

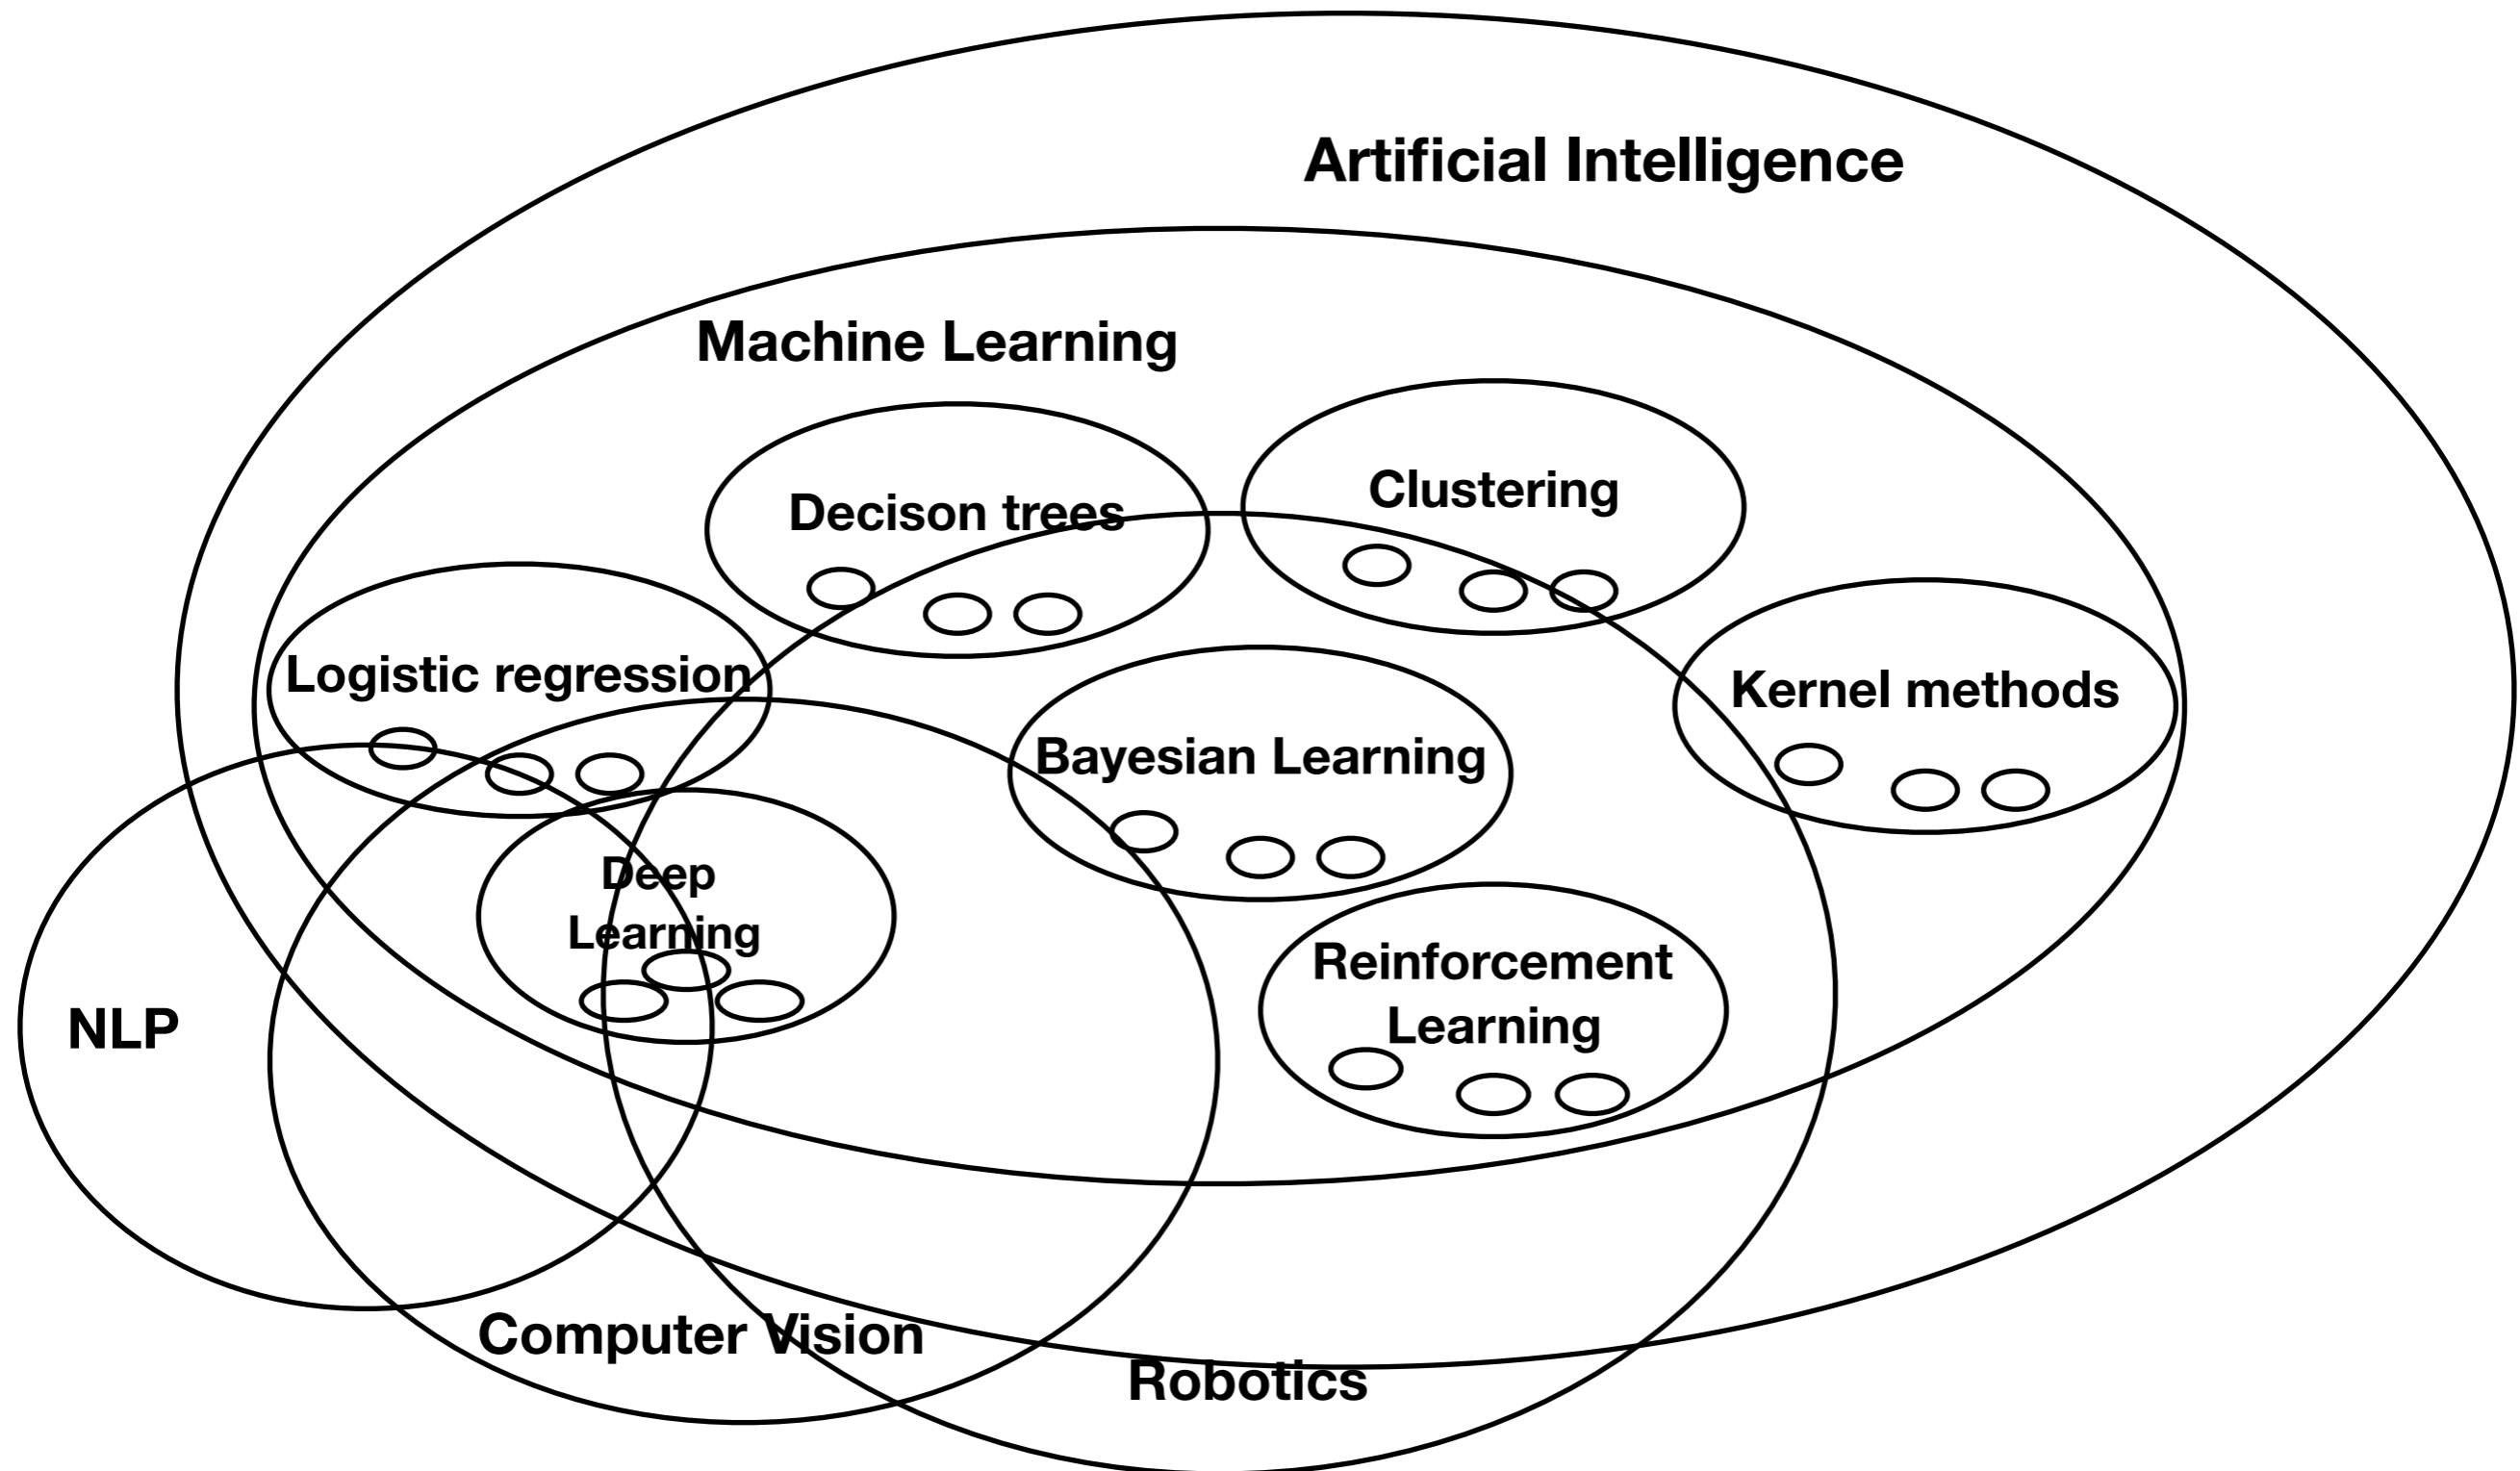
**Recurrent NN
(supervised)**

Some probabilistic methods


**Bayesian (probabilistic) methods
(super- / unsupervised)**

**Reinforcement learning
(unsupervised or “semi”-supervised)**

Drawing the (a?) EDAN95 ML Landscape (Methods and Application areas)



Or is it more like this?



Course syllabus

- Introduction (lecture 1), week 1
- Basics, weeks 1-3
 - Math and Python (lecture 2, lab 1)
 - Fundamental ML techniques (lectures 3-5, labs 2+3)
Concept learning, clustering, KNN, SOMs, Decision Trees, Feedforward networks, loss, regularization, ...
- Specific approaches, weeks 3-7
 - Deep learning techniques
Convolutional NNs, Recurrent NNs, LSTMs and GRUs, Autoencoders (lectures 6-9, lab 4)
 - Distributed techniques,
Spark / GPU (lecture 10, lab 5)
 - Probabilistic methods, Bayesian learning
MAP-learning, NBC, GMMs (lectures 11+12, lab 6)
 - Reinforcement Learning
Introduction, TD learning, Q-learning, Actor critic, applications in robotics (lectures 13+14, lab 7)

Course literature

- Listed books from course plan:
 - Kevin P. Murphy: Machine Learning, A Probabilistic Perspective. MIT Press, 2012, ISBN: 9780262018029.
 - Ian Goodfellow, Yoshua Bengio, Aaron Courville: Deep Learning. MIT Press, 2016, ISBN: 9780262035613.
 - Aurélien Géron: Hands-On Machine Learning with Scikit-Learn and TensorFlow, Concepts, Tools, and Techniques to Build Intelligent Systems. O'Reilly Media, 2017, ISBN: 9781491962299.
- Books not listed in course plan, but of some importance (and easily accessible):
 - François Chollet: Deep Learning with Python, Manning, 2018, ISBN: 9781617294433.
 - Tom Mitchell: Machine Learning. McGraw Hill, 1997, ISBN: 0070428077.
 - David L. Poole, Alan K. Mackworth: Artificial Intelligence - Foundations of Computational Agents (2e). Cambridge University Press, 2017, ISBN: 9781107195394.
- Other useful book(s):
 - Stuart Russel, Peter Norvig: Artificial Intelligence - A Modern Approach (3e). Pearson, 2010, ISBN: 10: 0132071487.