Introduction to COMP219

Dr. Xiaowei Huang

https://cgi.csc.liv.ac.uk/~xiaowei/


ALPHAGO
00:00:38
LEE SEDOL
00:01:00
AlphaGo
Google DeepMind

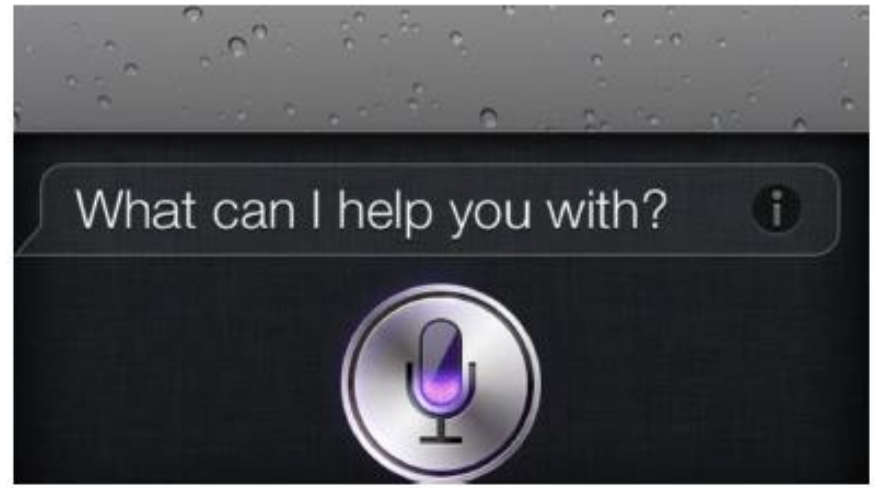
What can I help you with?


ID 8534569


ВЫХОД В ГОРОД
ACCESS TO CITY


405
SOUTH


PR2

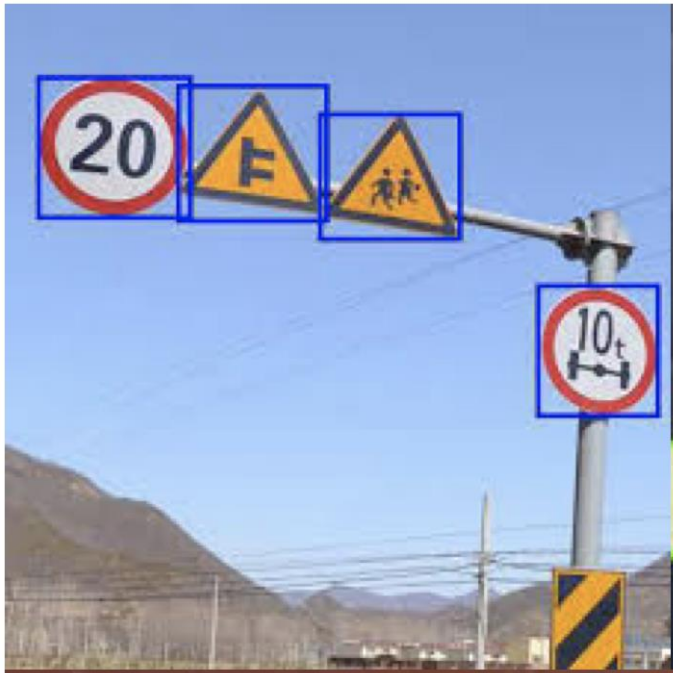
20
10t


189,413.27
(+12.64)
48,829.82
(+29.34)
52,394.62
(+7.30)
777.65
(+26.38)
1,002.66
15,127.81
16,696.99

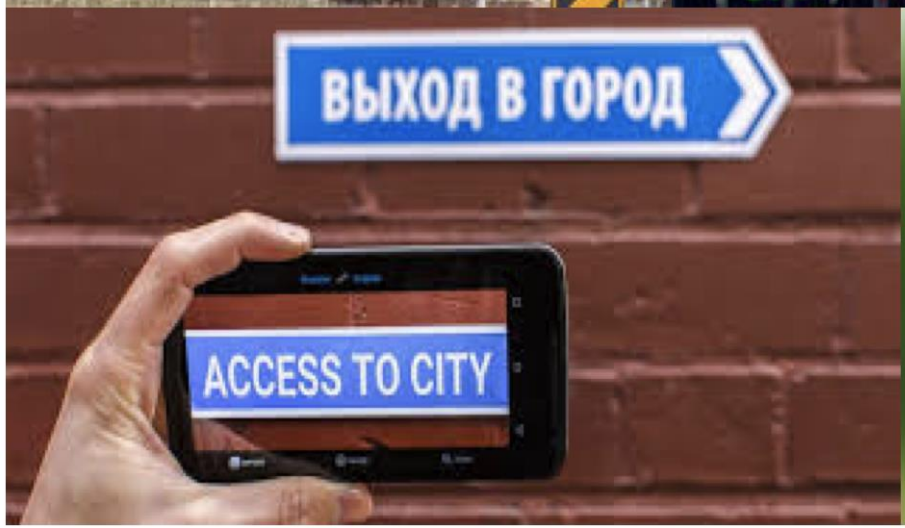
ВЫХОД В ГОРОД
ACCESS TO CITY

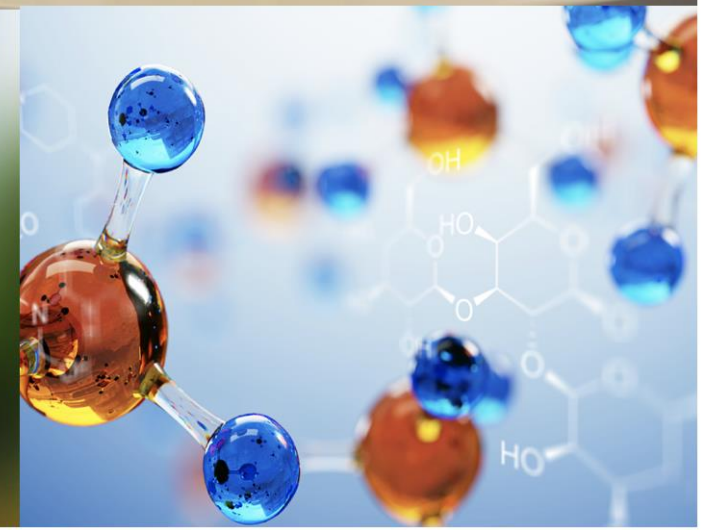
OH
HO
HO

Yes, they are AI-driven.

But to get there, it is nothing easy.

This module is not to do these fancy things, but to establish a foundation for you to be able to do them in the future.

Pre-requisite Knowledge - Probability

Number of Orders per Week x_i	Probability p_i
41	.03
42	.10
43	.15
44	.17
45	.25
46	.15
47	.10
48	.05

(1)

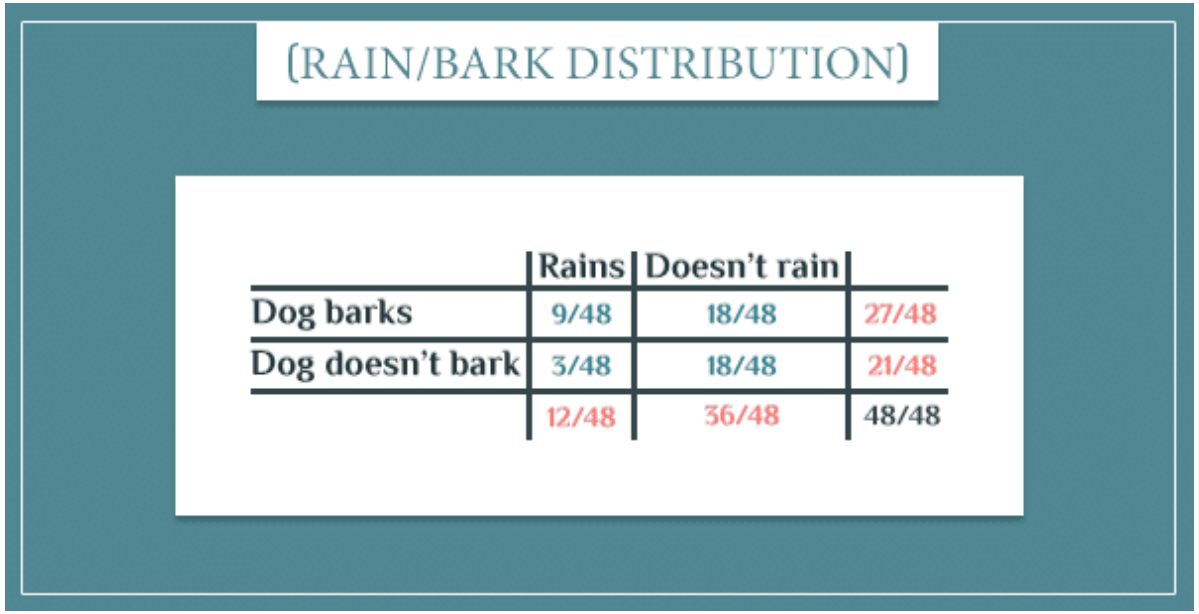

(2)

Pre-requisite knowledge – linear algebra

$$A = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \end{bmatrix}$$

(1)

$$x = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} * \begin{pmatrix} 9 \\ 8 \\ 7 \end{pmatrix}$$

$$1*9+2*8+3*7=46$$

$$4*9+5*8+6*7=118$$

$$x = \begin{pmatrix} 46 \\ 118 \end{pmatrix}$$

(2)

Pre-requisite knowledge – derivative and partial derivative

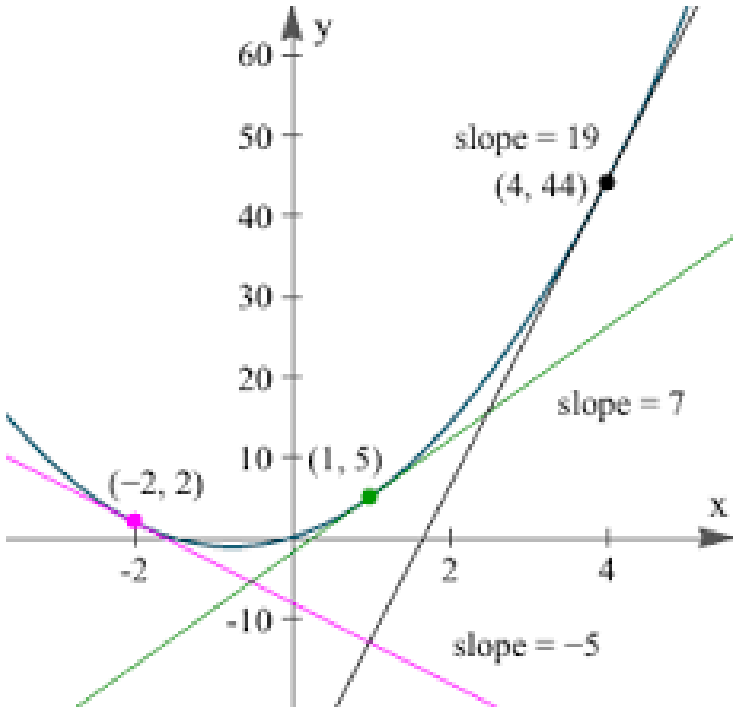

(1)

$$z = 3x^2 + 2xy - y^2$$

$$= 3(1)^2 + 2(1)(2) - (2)^2$$

$$= 3$$

$$\frac{\partial z}{\partial x} = 6x + 2y = 6(1) + 2(2) = 10$$

$$\frac{\partial z}{\partial x} = 2x + 2y = 2(1) + 2(2) = 6$$

(2)

Warning

- **A lot of maths in the first few weeks.**

- **If you choose this module, make sure that you are prepared (with knowledge, passion, persistency, etc)**

Today’s Content

- Module Information
- Contents of the module

Module Outline

- The module consists of
  - 25~30 lectures
  - ~6 lab sessions

- Please ensure sufficient time on self study

Module Outline

- Assessment
  - a two-hour exam (80%)
  - two practical assignments (10% each)

- Module information on Vital or course webpage (https://cgi.csc.liv.ac.uk/~xiaowei/ai.html)
  - 2018 course webpage for general information (https://cgi.csc.liv.ac.uk/~xiaowei/ai2018.html)
  - We will update this year

Module Delivery: Demonstrators

- 6 lab sessions
- 216 students registered
- Who is going to support this?
  - Mr Wei Huang and
  - Mr Gaojie Jin

Timetable: Lectures

- Tuesday 11am
- Wednesday 11am
- Thursday 10am

- Will be away on Tuesday, 15th October (4th week)

- Slides will be distributed the day before the lecture (for example, I may distribute the slides Monday evening for Tuesday lecture)

Lab Session

- We prepared 6-7 lab exercises
- Other than this, please follow the book “Python Machine Learning” to practice your ML skills
  - All codes are available at https://github.com/rasbt/python-machine-learning-book
- Demonstrators will try to help you

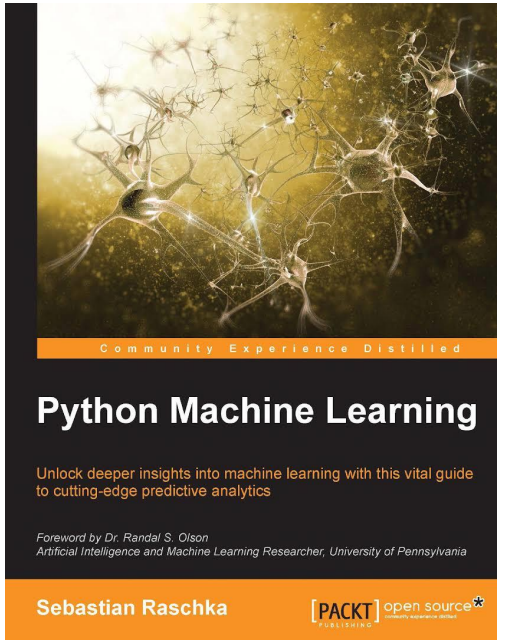

Reading

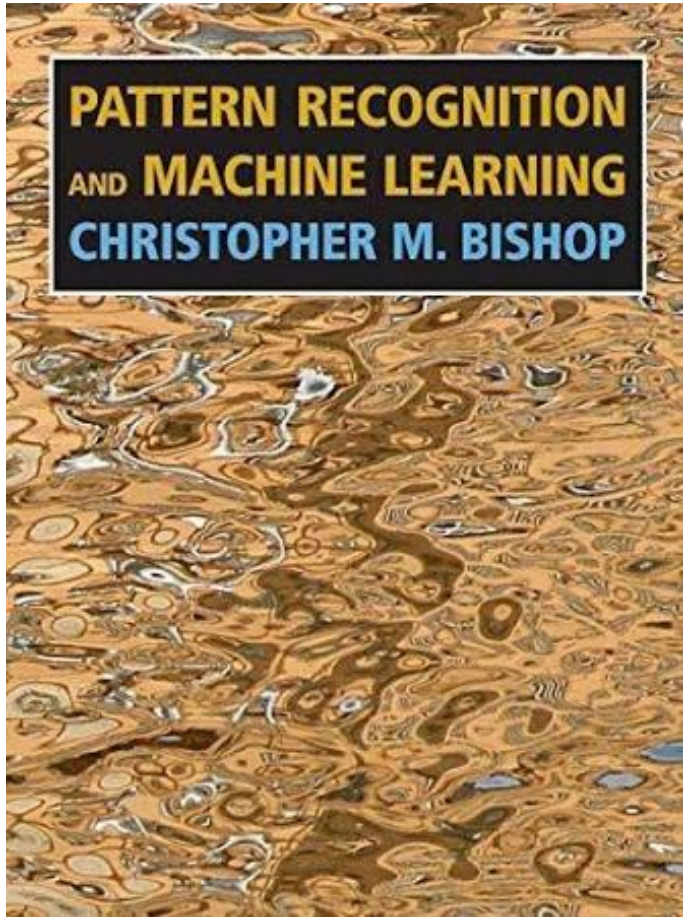
PATTERN RECOGNITION
AND MACHINE LEARNING
CHRISTOPHER M. BISHOP

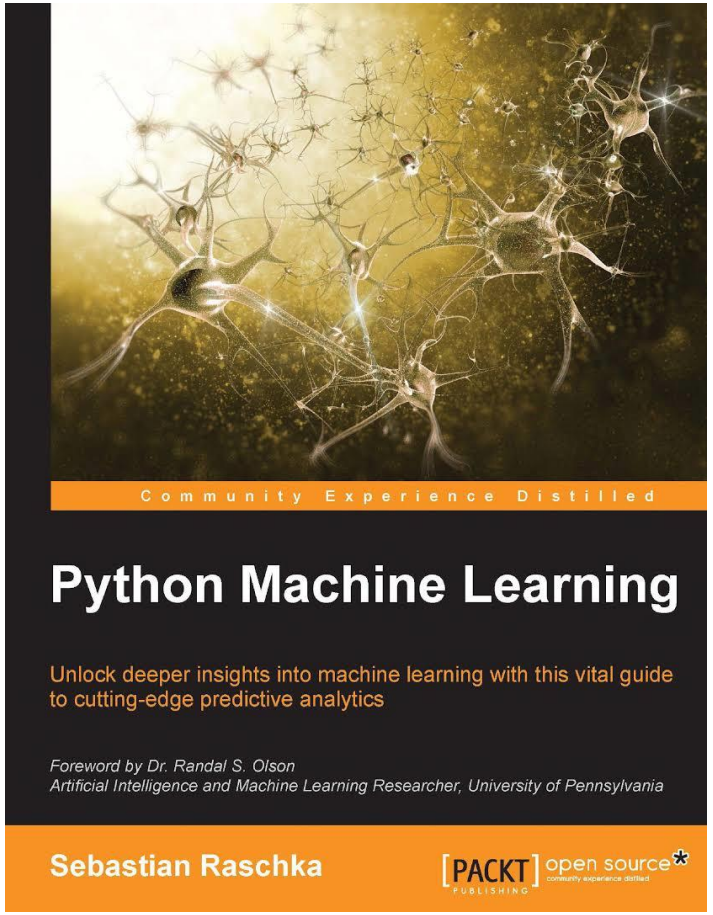
Community Experience Distilled
Python Machine Learning
Unlock deeper insights into machine learning with this vital guide to cutting-edge predictive analytics
Foreword by Dr. Randal S. Olson
Artificial Intelligence and Machine Learning Researcher, University of Pennsylvania
Sebastian Raschka
[PACKT] open source
PUBLISHING

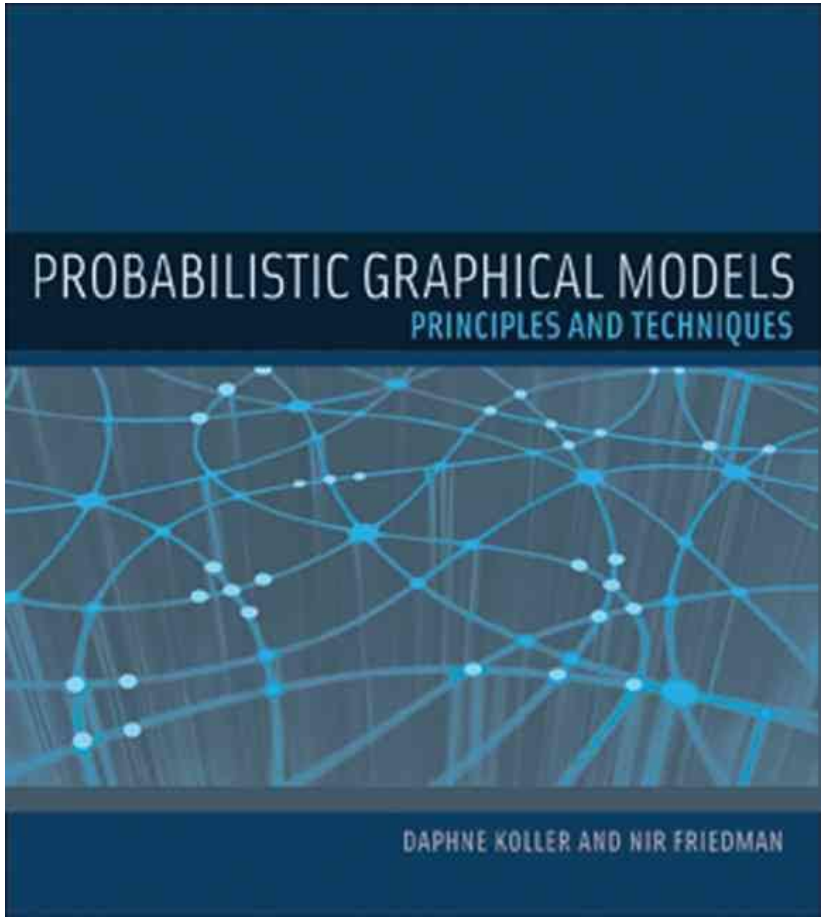
PROBABILISTIC GRAPHICAL MODELS
PRINCIPLES AND TECHNIQUES
DAPHNE KOLLER AND NIR FRIEDMAN

Other Reading:

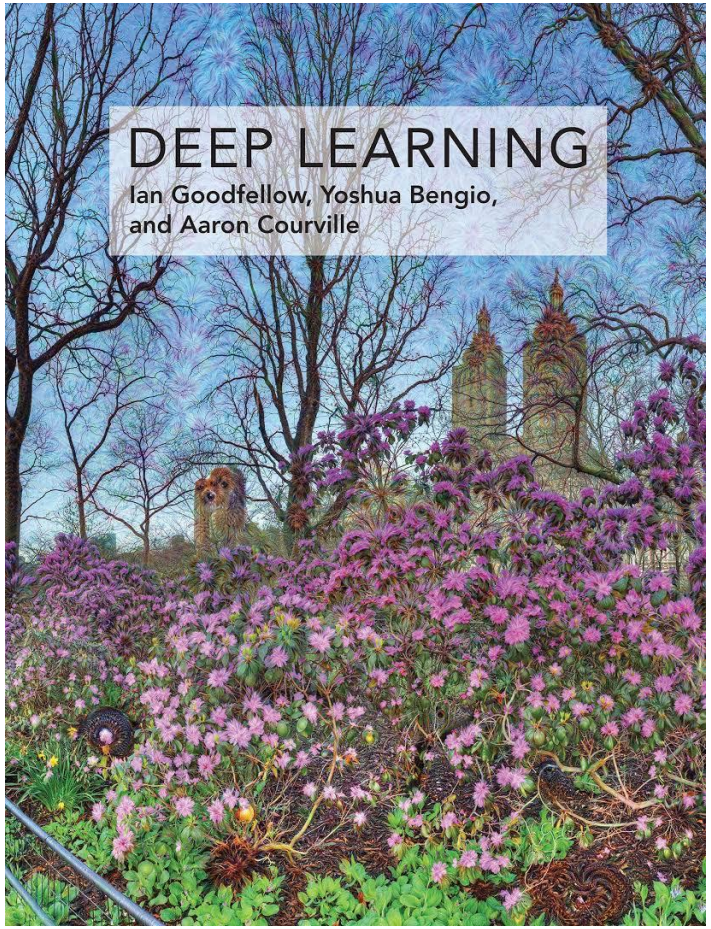

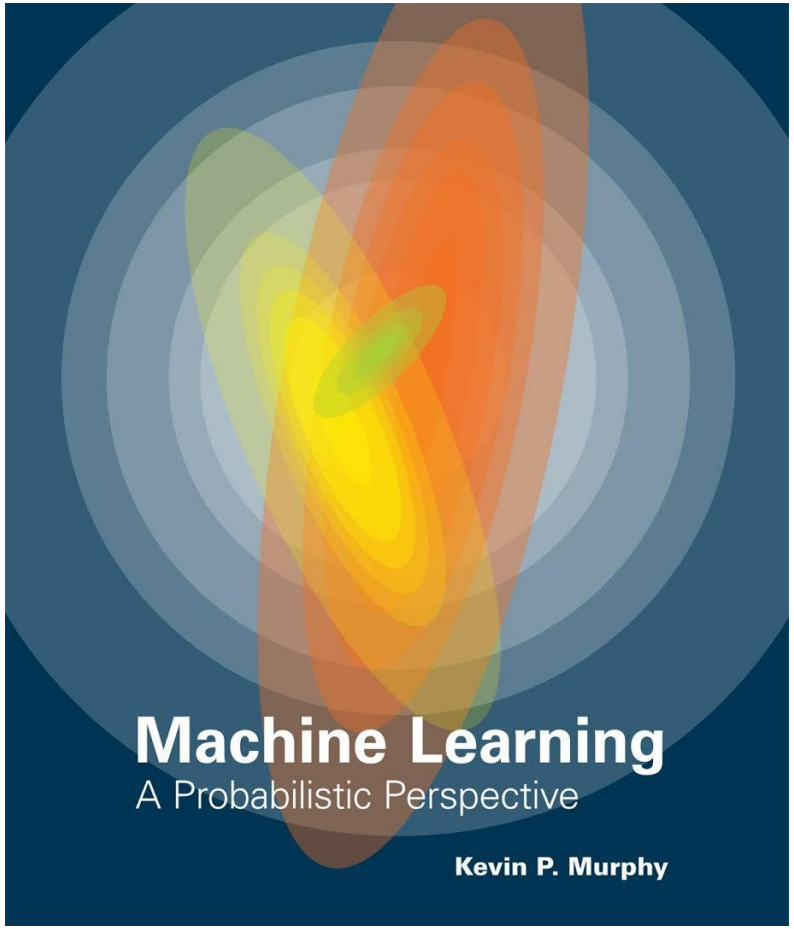

Other Reading:

- Tensorflow on-line documentations
- Github (plenty of resources, code, tutorials, etc)
- Various on-line courses
- Reddit, quite some good discussions. Experts are around there.
- Kaggle competitions, you can participate in to get more hand-on experience
- Wikipedia, for various concepts, key pointers, etc
- …. Many other on-line resources, please Google whatever you want

COMP111

- Brief history of AI including recent developments
- Intelligent Agents: A classification
- Search (applications: route planning, game playing)
- Knowledge Representation (applications: structured web search output)
- Reasoning under Uncertainty (application: almost everywhere)
- Learning (applications: face recognition, selfdriving cars)
- Philosophy and Ethics of AI (motivation: deducing sexual orientation from your picture ok? Visit `https:`

Aims

- To equip students with the knowledge about basic algorithms that have been used to enable the AI agents to conduct the perception, inference, and planning tasks;
- To equip students with the knowledge about machine learning algorithms;
- To provide experience in applying basic AI algorithms to solve problems;
- To provide experience in applying machine learning algorithms to practical problems;

Perception-Cognition-Action Loop

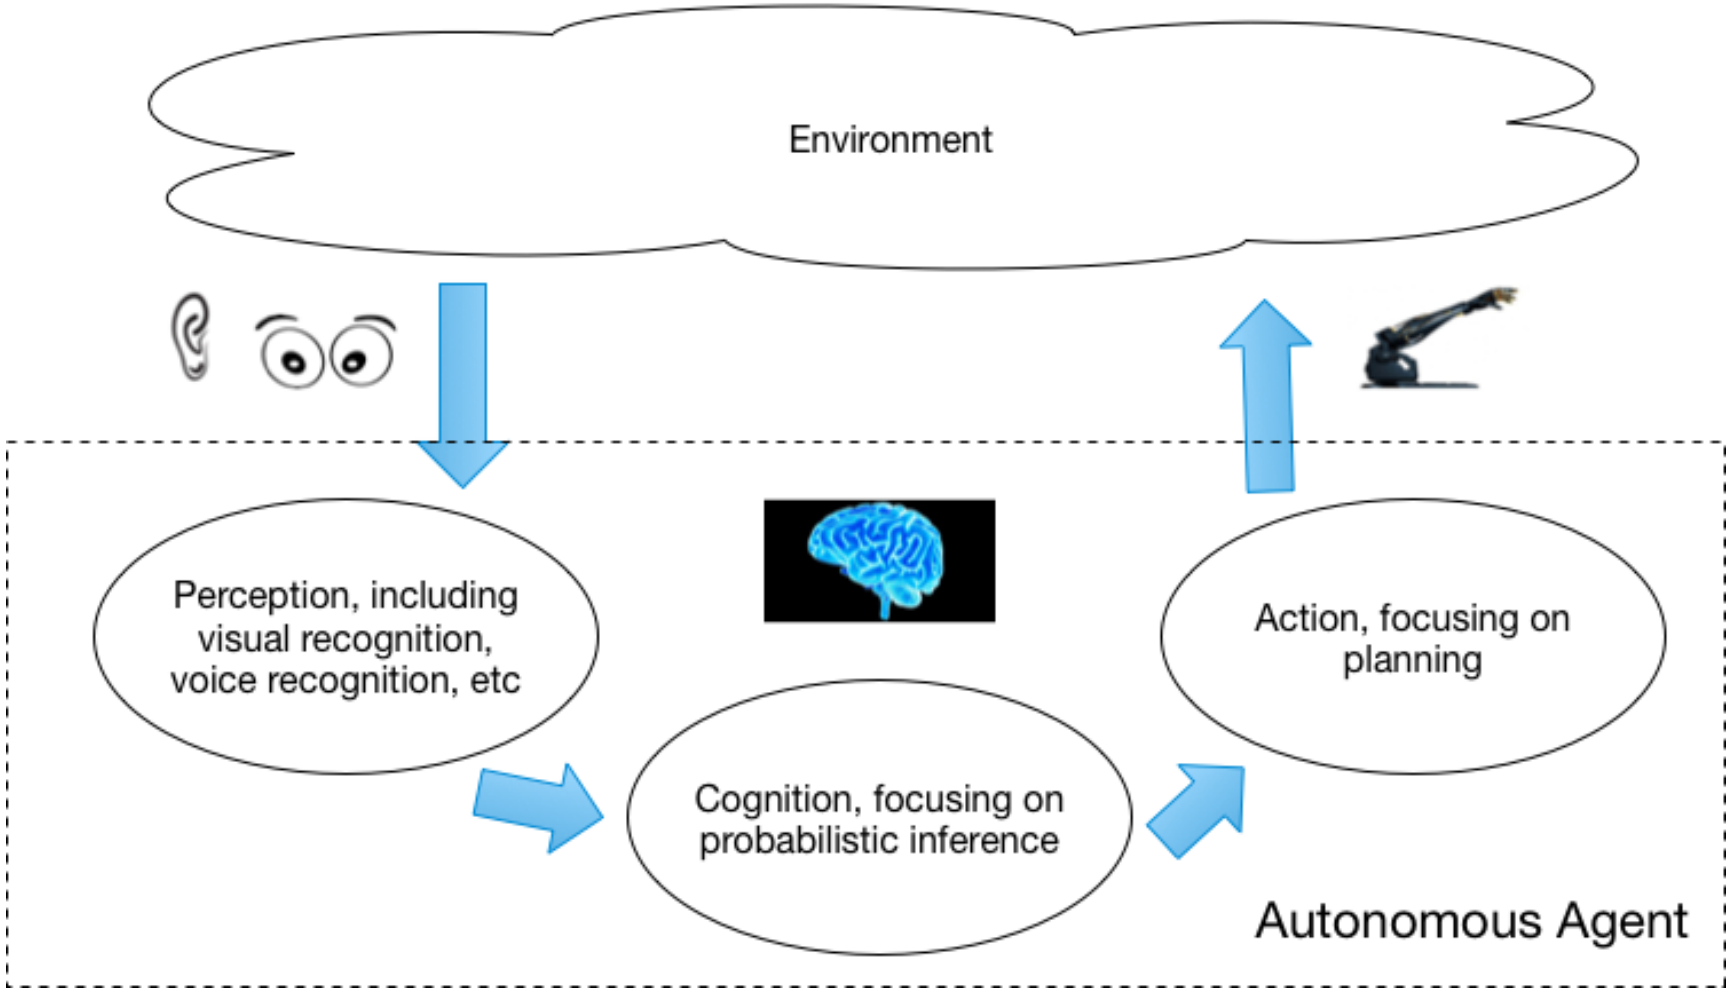

Teaching content:
traditional learning,
deep learning

Perception

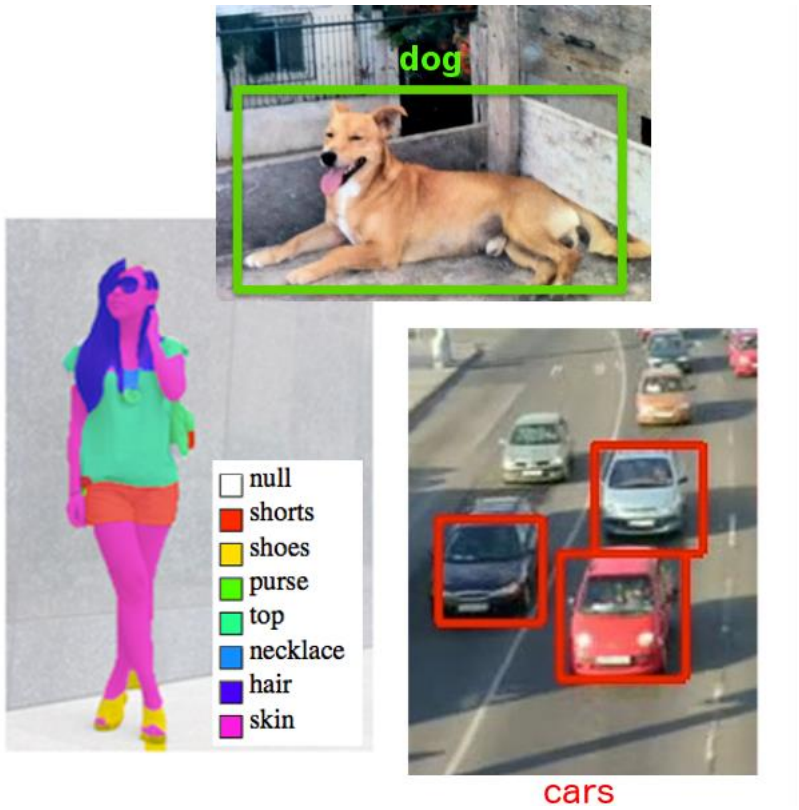

Visual Recognition

Voice Recognition

…

Cognition by Probabilistic Inference

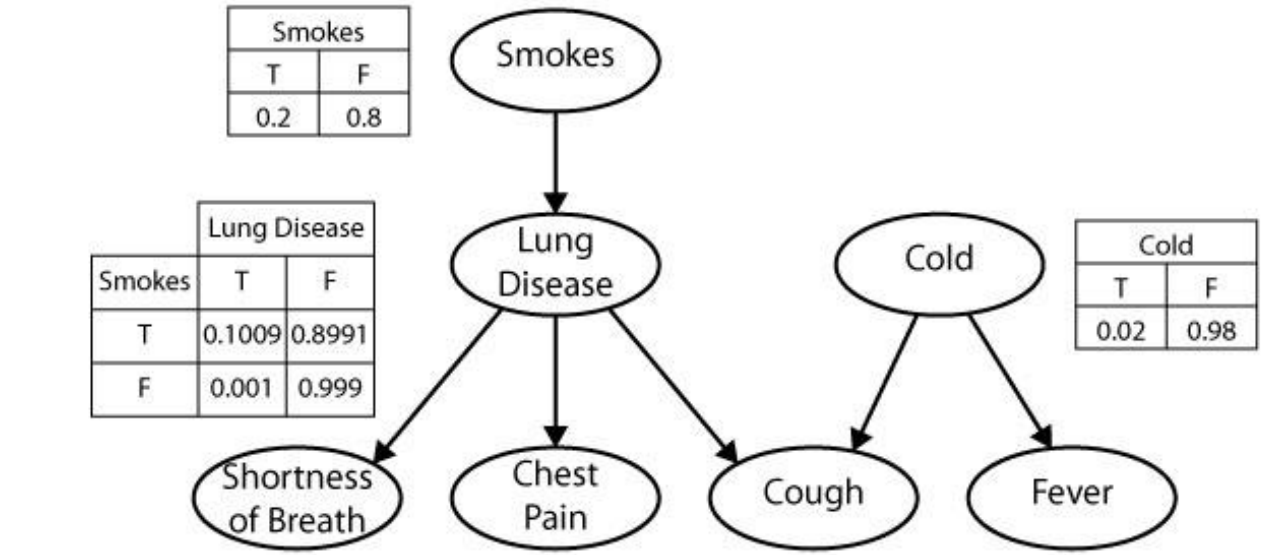

Smokes	
T	F
0.2	0.8

	Lung Disease	
Smokes	T	F
T	0.1009	0.8991
F	0.001	0.999

Cold	
T	F
0.02	0.98

	Shortness of Breath	
Lung Disease	T	F
T	0.208	0.792
F	0.01	0.99

	Chest Pain	
Lung Disease	T	F
T	0.208	0.792
F	0.01	0.99

		Cough	
Lung Disease	Cold	T	F
T	T	0.7525	0.2475
T	F	0.505	0.495
F	T	0.505	0.495
F	F	0.01	0.99

	Fever	
Cold	T	F
T	0.307	0.693
F	0.01	0.99

Q. how to automatically infer the disease (e.g., lung disease, cold, etc) from the symptoms (e.g., smokes, shortness of breath, chest pain, cough, fever, etc)?

Note: Symptoms obtained from perception.

Teaching content:
in other modules, e.g., COMP111, COMP222

Action by Planning

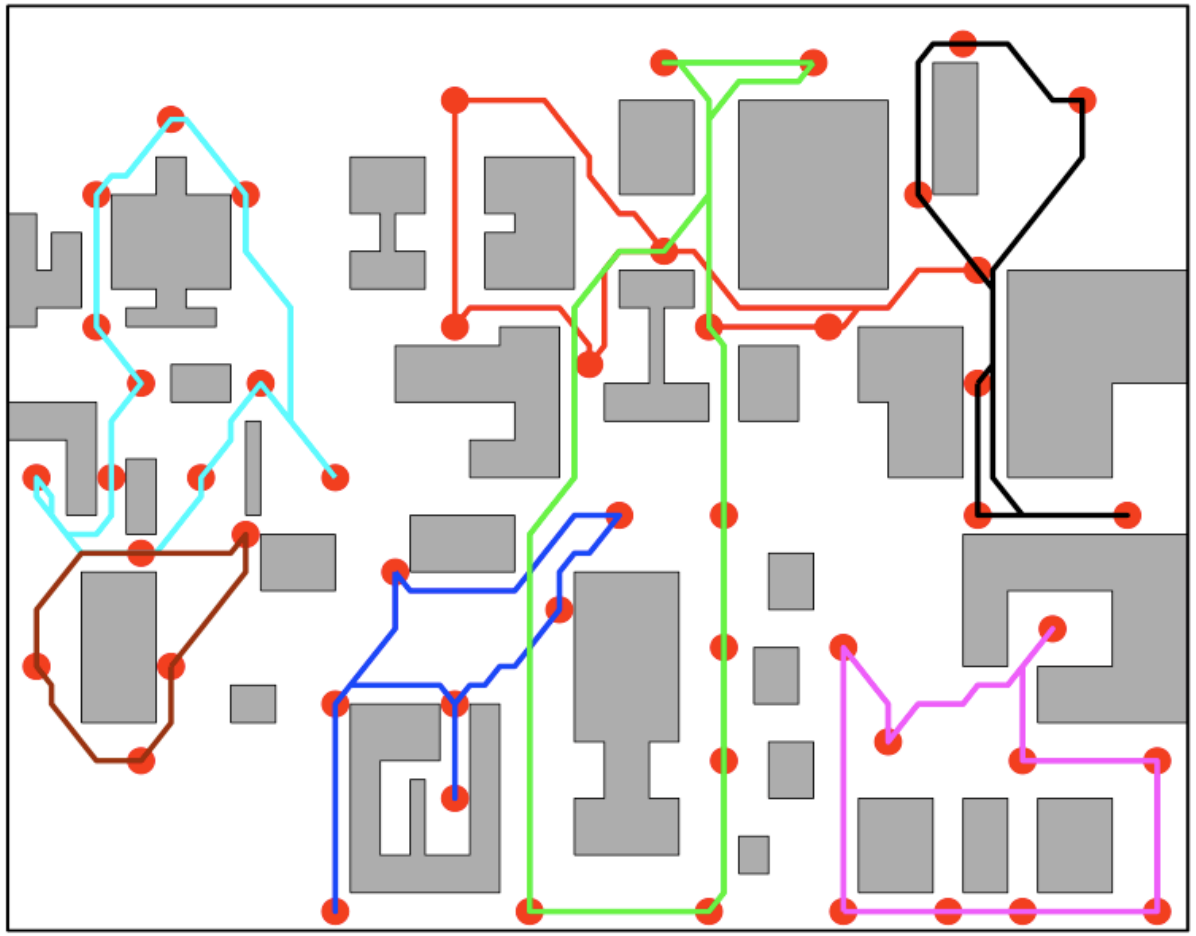

After cognition, we may use the obtained knowledge to react to the environment

Q: in the factory floor as shown in the left diagram, how many robots is needed to patrol the area? and **how to plan** their activities?

Learning Outcomes

- Ability to explain in detail how the techniques in the perceive-inference-action loop work
- Ability to choose, compare, and apply suitable basic learning algorithms to simple applications
- Ability to explain how deep neural networks are constructed and trained, and apply deep neural networks to work with large scale datasets
- Ability to conduct probabilistic inference.

Contents of this module

- Introduction
- preliminary knowledge (probabilistic foundation, linear algebra)
- Traditional machine learning (gradient descent, decision tree learning, K-nn, model evaluation, linear regression, naïve Bayes)
- Practical tutorial (python, tensorflow)
- Deep learning
- Probabilistic graphical models
- (optional) advanced topics

Credits

- I used many resources from the web