人工智慧 (Artificial Intelligence)

機器人技術 (Robotics)

1092AI11
MBA, IM, NTPU (M5010) (Spring 2021)
Wed 2, 3, 4 (9:10-12:00) (B8F40)

Min-Yuh Day
戴敏育
Associate Professor
副教授
Institute of Information Management, National Taipei University
國立臺北大學 資訊管理研究所
https://web.ntpu.edu.tw/~myday

2021-06-02

課程大綱 (Syllabus)

國立臺北大學 National Taipei University

週次 (Week)	日期 (Date)	內容 (Subject/Topics)
1	2021/02/24	人工智慧概論 (Introduction to Artificial Intelligence)
2	2021/03/03	人工智慧和智慧代理人 (Artificial Intelligence and Intelligent Agents)
3	2021/03/10	問題解決 (Problem Solving)
4	2021/03/17	知識推理和知識表達 (Knowledge, Reasoning and Knowledge Representation)
5	2021/03/24	不確定知識和推理 (Uncertain Knowledge and Reasoning)
6	2021/03/31	人工智慧個案研究 I (Case Study on Artificial Intelligence I)

課程大綱 (Syllabus)

國立臺北大學
National Taipei University

週次 (Week) 日期 (Date) 內容 (Subject/Topics)

7 2021/04/07 放假一天 (Day off)

8 2021/04/14 機器學習與監督式學習
(Machine Learning and Supervised Learning)

9 2021/04/21 期中報告
(Midterm Project Report)

10 2021/04/28 學習理論與綜合學習
(The Theory of Learning and Ensemble Learning)

11 2021/05/05 深度學習
(Deep Learning)

12 2021/05/12 人工智慧個案研究 II
(Case Study on Artificial Intelligence II)

課程大綱 (Syllabus)

週次 (Week) 日期 (Date) 內容 (Subject/Topics)

13 2021/05/19 強化學習
(Reinforcement Learning)

14 2021/05/26 深度學習自然語言處理
(Deep Learning for Natural Language Processing)

15 2021/06/02 機器人技術
(Robotics)

16 2021/06/09 人工智慧哲學與倫理，人工智慧的未來
(Philosophy and Ethics of AI, The Future of AI)

17 2021/06/16 期末報告 I
(Final Project Report I)

18 2021/06/23 期末報告 II
(Final Project Report II)

Robotics

Outline

- Robots
- Robotic Perception
- Planning and Control
- Planning Uncertain Movements
- Reinforcement Learning in Robotics
- Humans and Robots

Stuart Russell and Peter Norvig (2020),

Artificial Intelligence: A Modern Approach,

4th Edition, Pearson

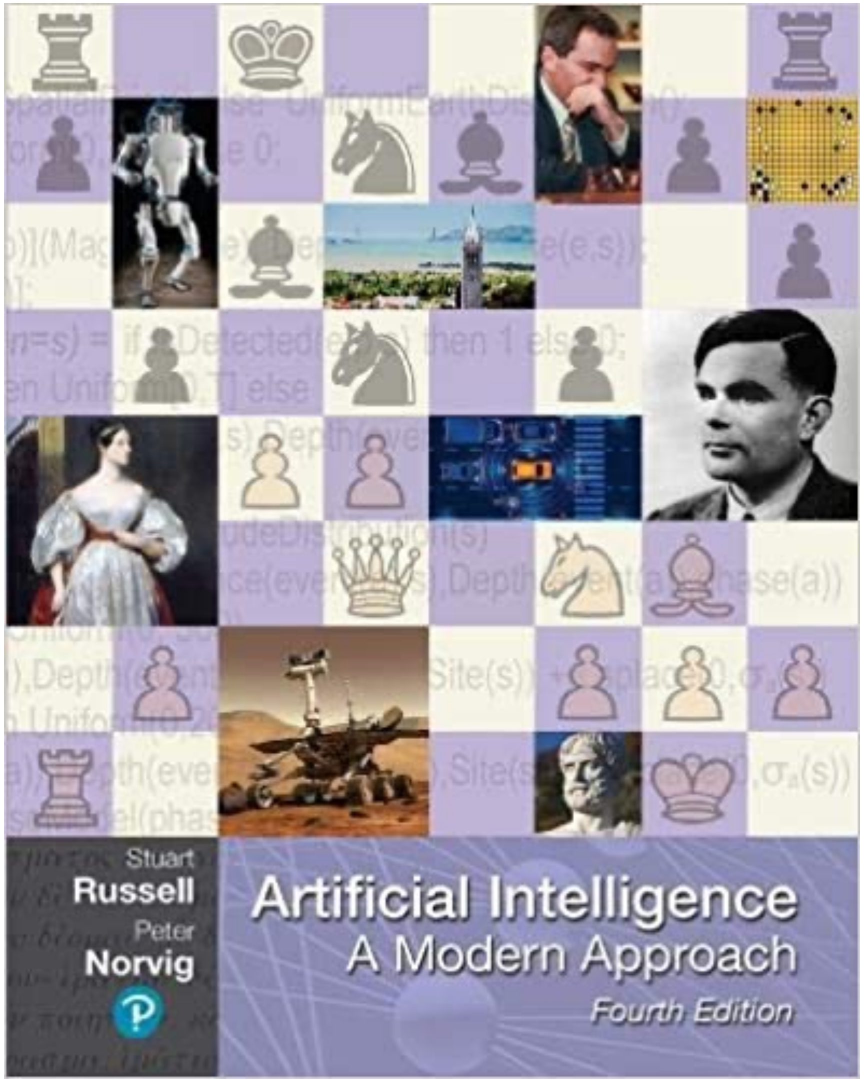

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

https://www.amazon.com/Artificial-Intelligence-A-Modern-Approach/dp/0134610997/

Artificial Intelligence: A Modern Approach

1. Artificial Intelligence
2. Problem Solving
3. Knowledge and Reasoning
4. Uncertain Knowledge and Reasoning
5. Machine Learning
6. Communicating, Perceiving, and Acting
7. Philosophy and Ethics of AI

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Artificial Intelligence: Communicating, perceiving, and acting

Artificial Intelligence:
6. Communicating, Perceiving, and Acting

- Natural Language Processing
- Deep Learning for Natural Language Processing
- Computer Vision
- Robotics

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Artificial Intelligence: Robotics

- Robots
- Robot Hardware
- What kind of problem is robotics solving?
- Robotic Perception
- Planning and Control
- Planning Uncertain Movements
- Reinforcement Learning in Robotics
- Humans and Robots
- Alternative Robotic Frameworks
- Application Domains

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Reinforcement Learning (DL)

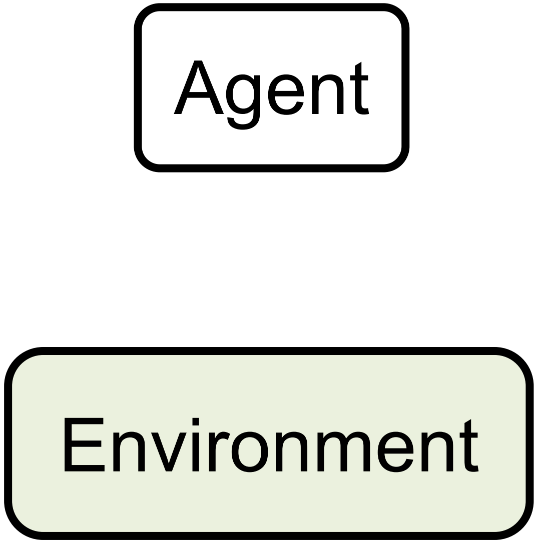

Source: Richard S. Sutton & Andrew G. Barto (2018), Reinforcement Learning: An Introduction, 2nd Edition, A Bradford Book.

Reinforcement Learning (DL)

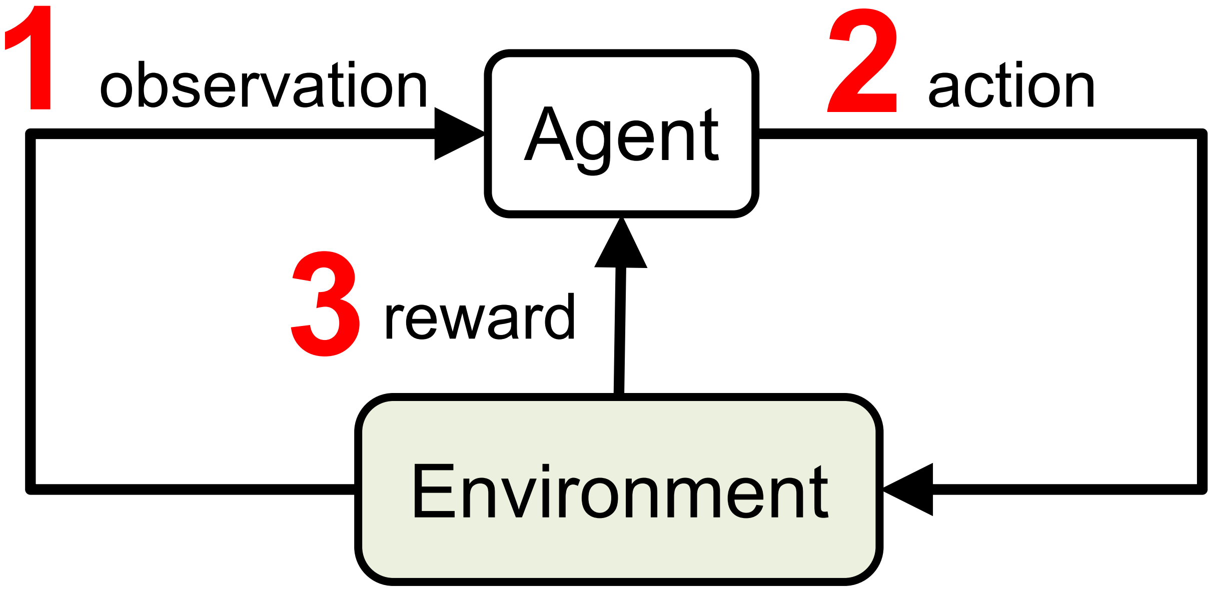

Source: Richard S. Sutton & Andrew G. Barto (2018), Reinforcement Learning: An Introduction, 2nd Edition, A Bradford Book.

Reinforcement Learning (DL)

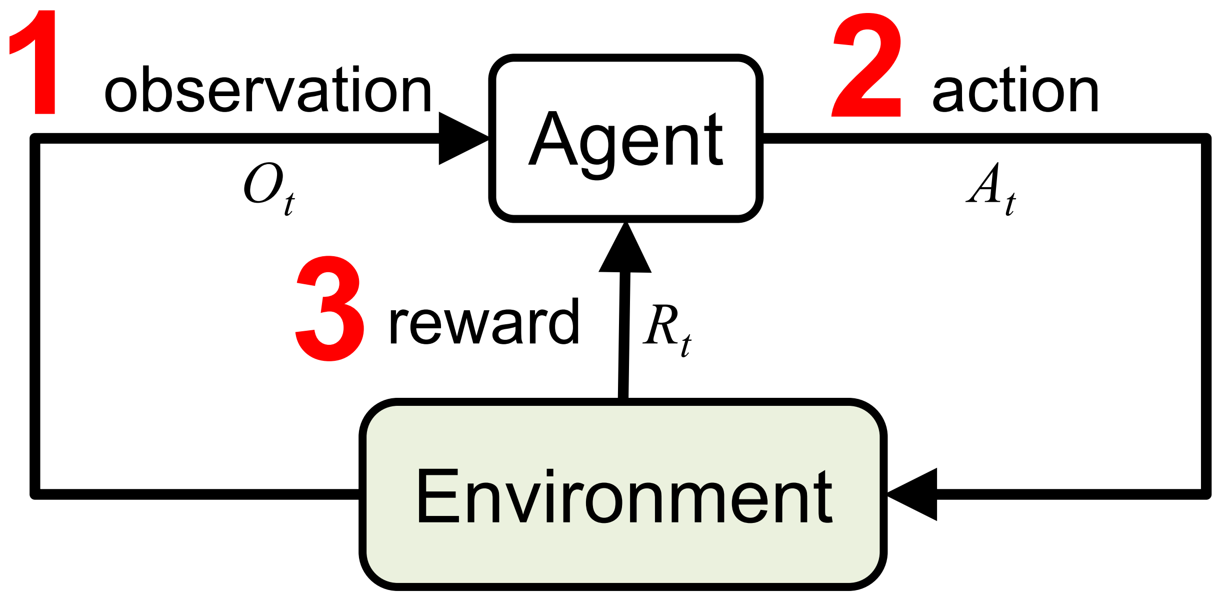

Source: Richard S. Sutton & Andrew G. Barto (2018), Reinforcement Learning: An Introduction, 2nd Edition, A Bradford Book.

Agents interact with environments through sensors and actuators

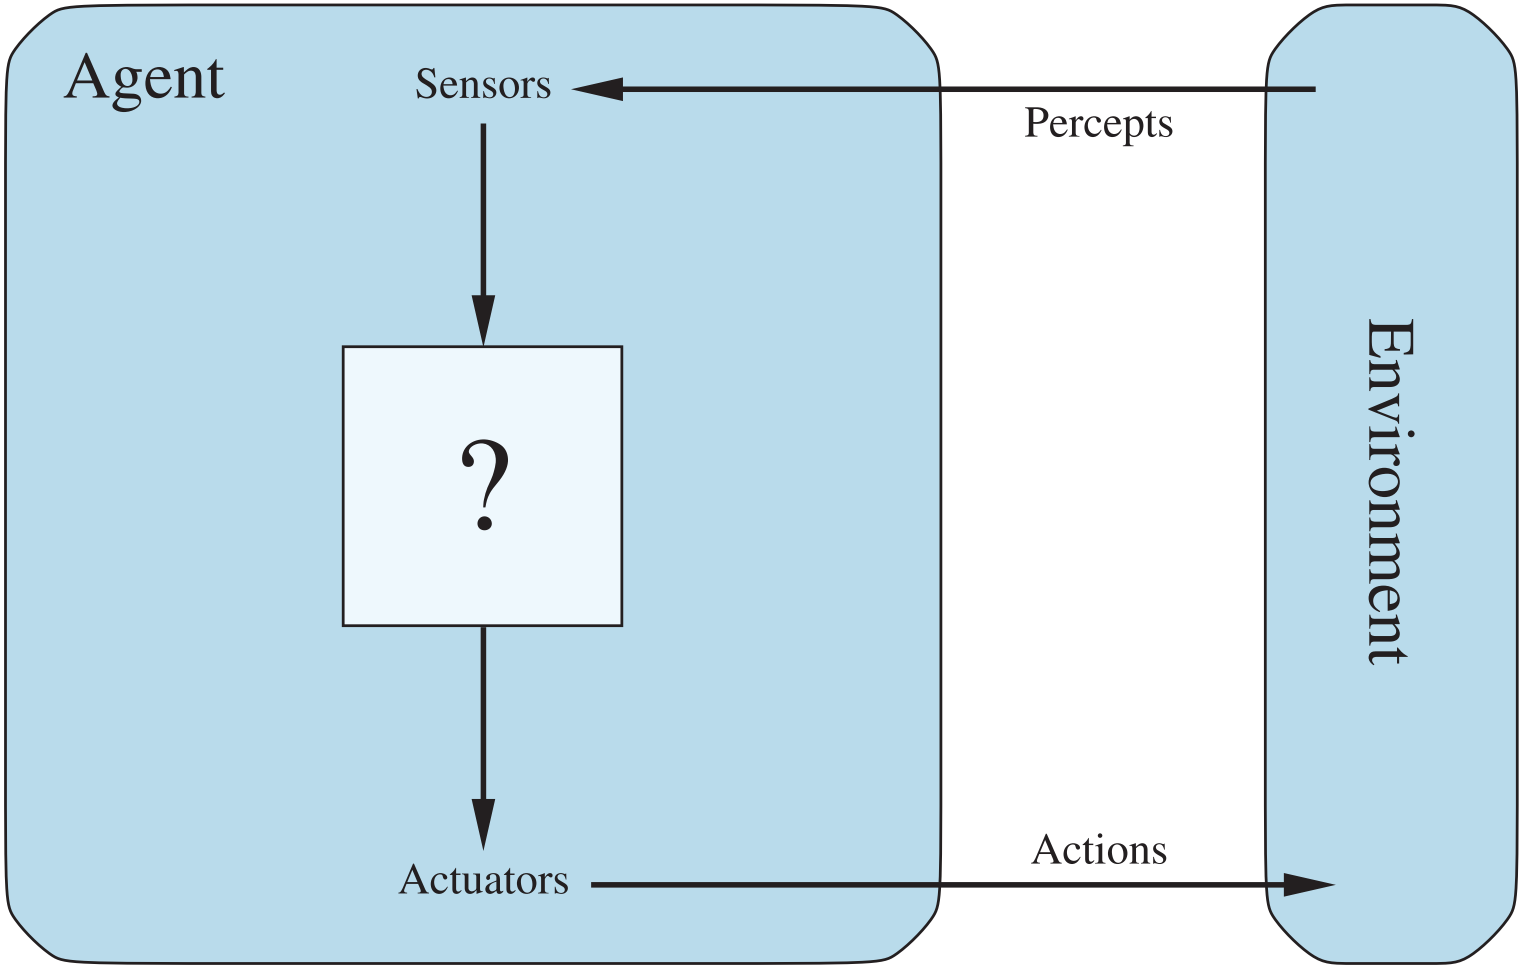

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Robotics

Artificial Intelligence: Robotics

- **Agents** are endowed with **sensors** and **physical effectors** with which to move about and make mischief in the real world.

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Boston Dynamics: Spot

Automate sensing and inspection, capture limitless data, and explore without boundaries.

Source: https://www.bostondynamics.com/spot

Boston Dynamics: Atlas
The world's most dynamic humanoid robot

Atlas is a research platform designed to push the limits of whole-body mobility

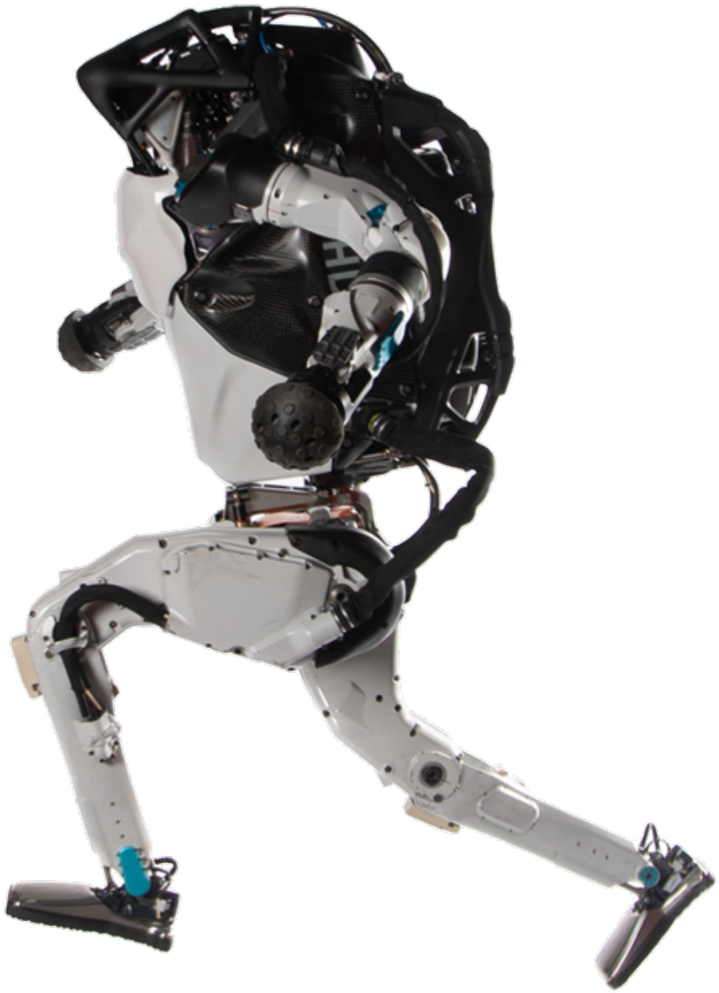

Source: https://www.bostondynamics.com/atlas

Boston Dynamics: Atlas

#13 ON TRENDING

What's new, Atlas?

https://www.youtube.com/watch?v=fRj34o4hN4I

Humanoid Robot: Sophia

https://www.youtube.com/watch?v=S5t6K9iwcdw

Can a robot pass a university entrance exam?

Noriko Arai at TED2017

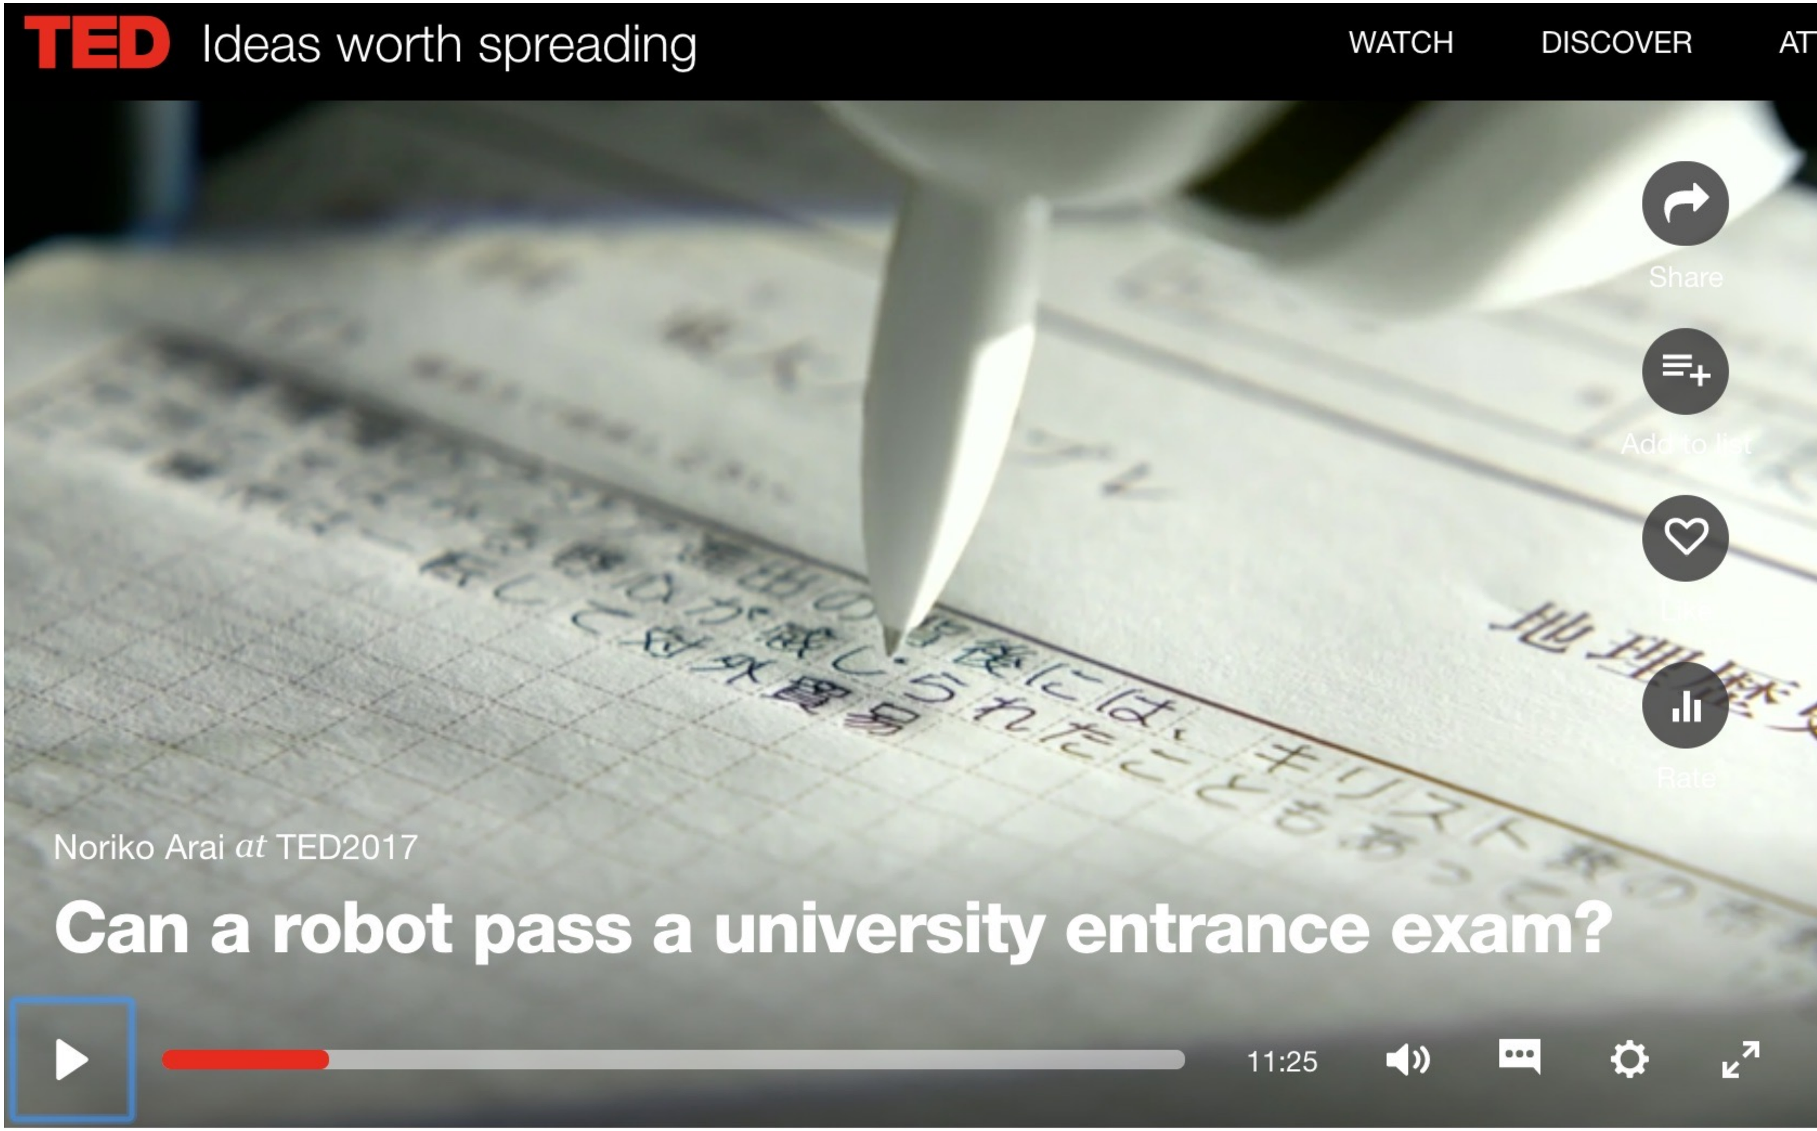

https://www.ted.com/talks/noriko_arai_can_a_robot_pass_a_university_entrance_exam

https://www.youtube.com/watch?v=XQZjkPyJ8KU

Artificial Intelligence (A.I.) Timeline

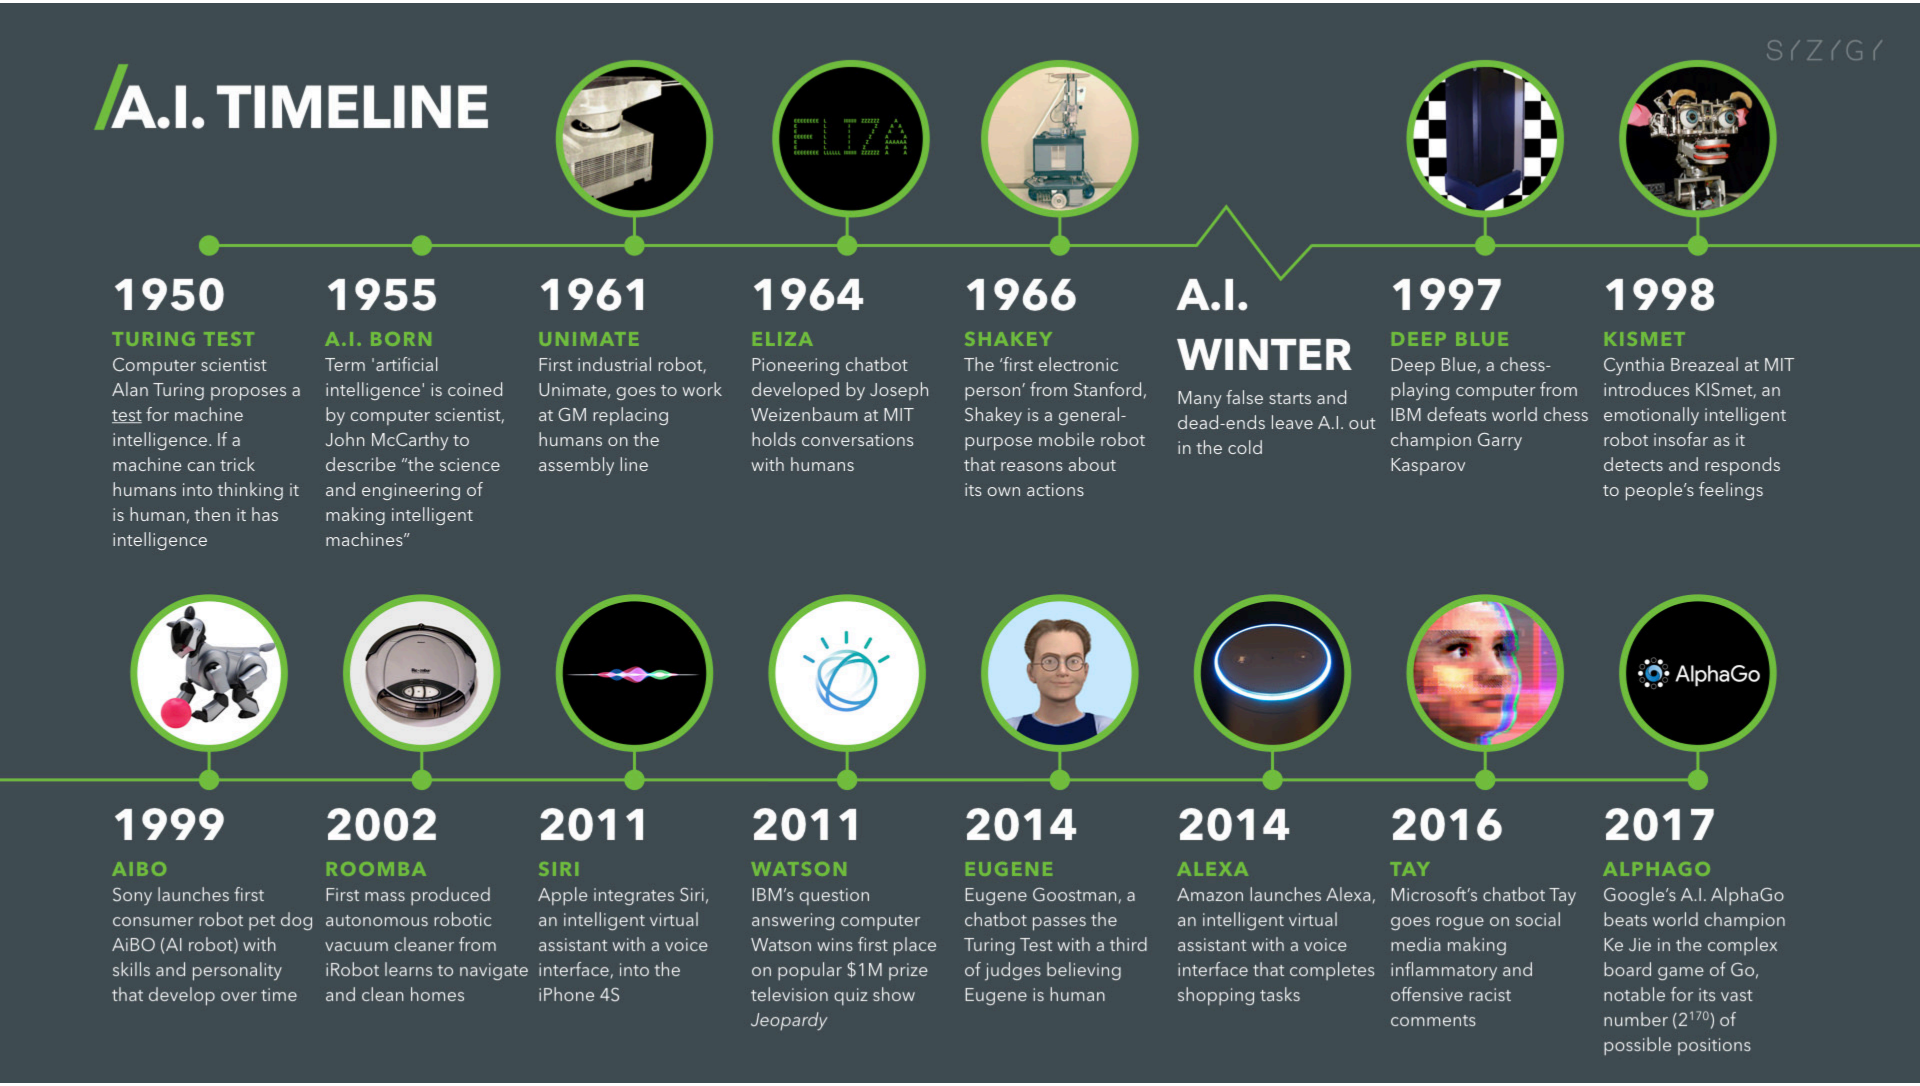

Source: https://digitalintelligencetoday.com/artificial-intelligence-timeline-infographic-from-eliza-to-tay-and-beyond/

Robots

- Robots are **physical agents** that perform tasks by manipulating the physical world.
  - To do so, they are equipped with **effectors** such as legs, wheels, joints, and grippers.
- **Effectors** are designed to assert physical forces on the environment.

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Robots and Effectors

- When they do this, a few things may happen:
  - the robot's state might change
  - the state of the environment might change
  - the state of the people around the robot might change

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Robots

- The most common types of robots are **manipulators** (robot arms) and **mobile robots**.
- They have **sensors** for perceiving the world and **actuators** that produce motion, which then affects the world via **effectors**.

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Robotics Problem

- The general robotics problem involves
  - **stochasticity**
    (which can be handled by MDPs)
  - **partial observability**
    (which can be handled by POMDPs)
  - acting with and around **other agents**
    (which can be handled with game theory)

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Robotic Perception

- We typically separate perception (estimation) from action (motion generation).
- Perception in robotics involves computer vision to recognize the surroundings through cameras, but also localization and mapping.

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Robotic Perception

- **Robotic perception** concerns itself with estimating decision-relevant quantities from sensor data.
  - To do so, we need an internal representation and a method for updating this internal representation over time.

Robot Perception

can be viewed as temporal inference from sequences of actions and measurements

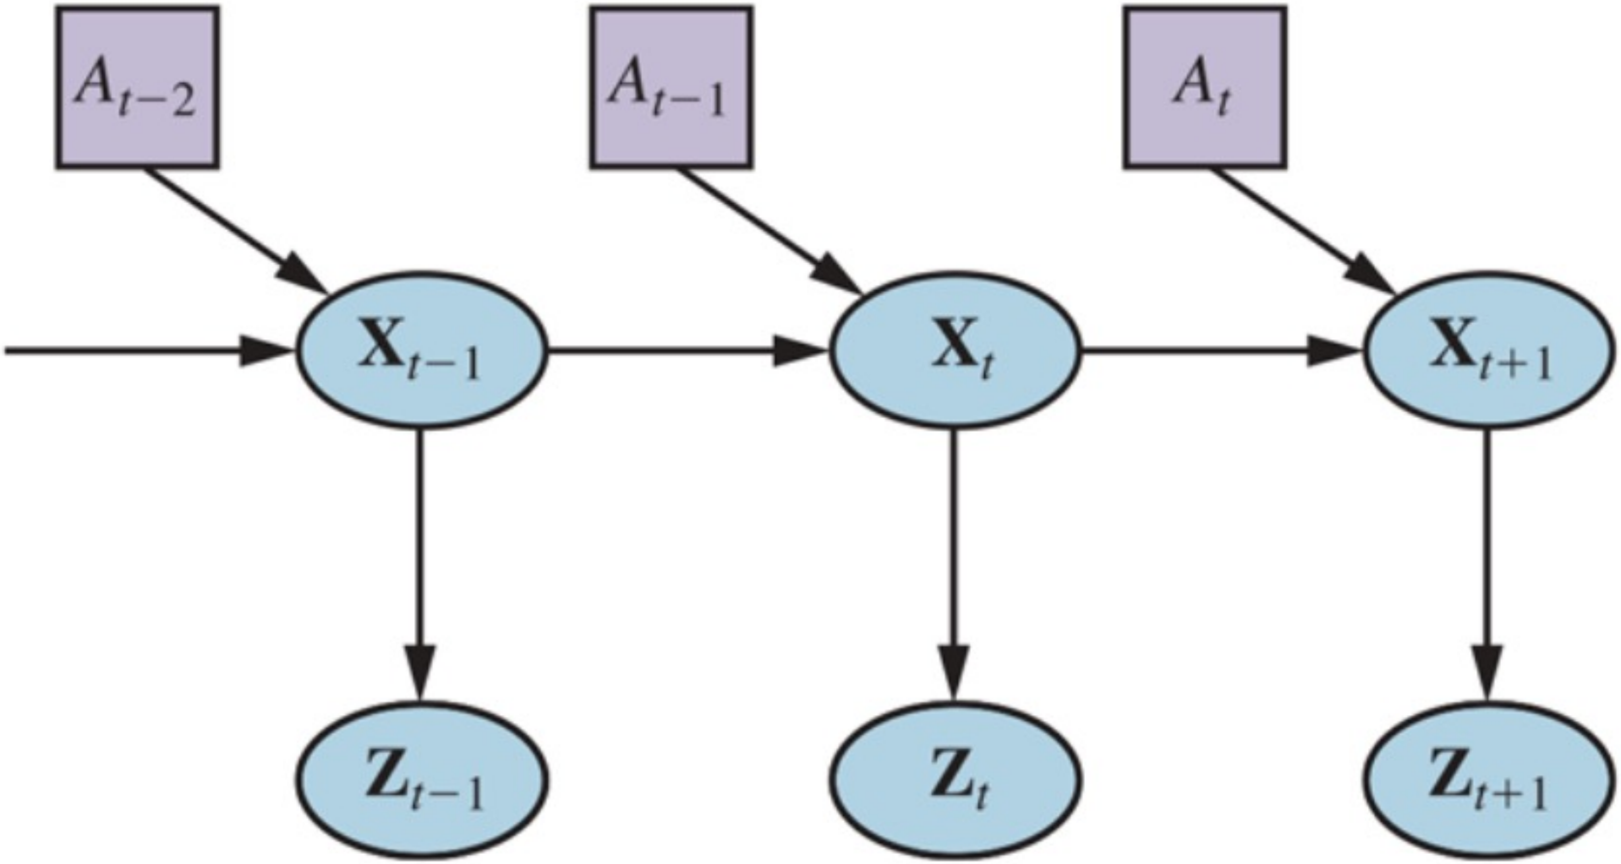

Dynamic Decision network

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Probabilistic Filtering Algorithms

- Probabilistic filtering algorithms such as particle filters and Kalman filters are useful for robot perception.
  - These techniques maintain the belief state, a posterior distribution over state variables.

Configuration Spaces

- For generating motion, we use **configuration spaces**, where a point specifies everything we need to know to locate every **body point** on the robot.
  - For instance, for a robot arm with two joints, a configuration consists of the two joint angles.

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Motion Generation

- We typically decouple the motion generation problem into
  - **motion planning**, concerned with producing a plan, and
  - **trajectory tracking control**, concerned with producing a policy for control inputs (actuator commands) that results in executing the plan.

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Motion Planning

- Motion planning can be solved via graph search
  - using **cell decomposition**
  - using **randomized motion planning** algorithms, which sample milestones in the continuous configuration space
  - using **trajectory optimization**, which can iteratively push a straight-line path out of collision by leveraging a signed distance field.

Planning and Control

- **Optimal control** unites motion planning and trajectory tracking by computing an optimal trajectory directly over control inputs.

Planning Uncertain Movements

- **Planning under uncertainty** unites perception and action by
  - **online replanning** (such as model predictive control) and
  - **information gathering** actions that aid perception.

Reinforcement learning in robotics

- **Reinforcement learning** is applied in robotics, with techniques striving to reduce the required number of interactions with the real world.
- Such techniques tend to **exploit models**, be it estimating models and using them to plan, or training policies that are robust with respect to different possible model parameters.

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Humans and Robots

- Interaction with humans requires the ability to **coordinate** the robot's actions with theirs, which can be formulated as a game.
- We usually decompose the solution into **prediction**, in which we use the person's ongoing actions to estimate what they will do in the future, and action, in which we use the predictions to compute the optimal motion for the robot.

Humans and Robots

- Helping humans also requires the ability to learn or infer what they want.
- Robots can approach this by learning the desired cost function they should optimize from human input, such as demonstrations, corrections, or instruction in natural language.
- Alternatively, robots can imitate human behavior, and use reinforcement learning to help tackle the challenge of generalization to new states.

Source: Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson

Papers with Code
State-of-the-Art (SOTA)

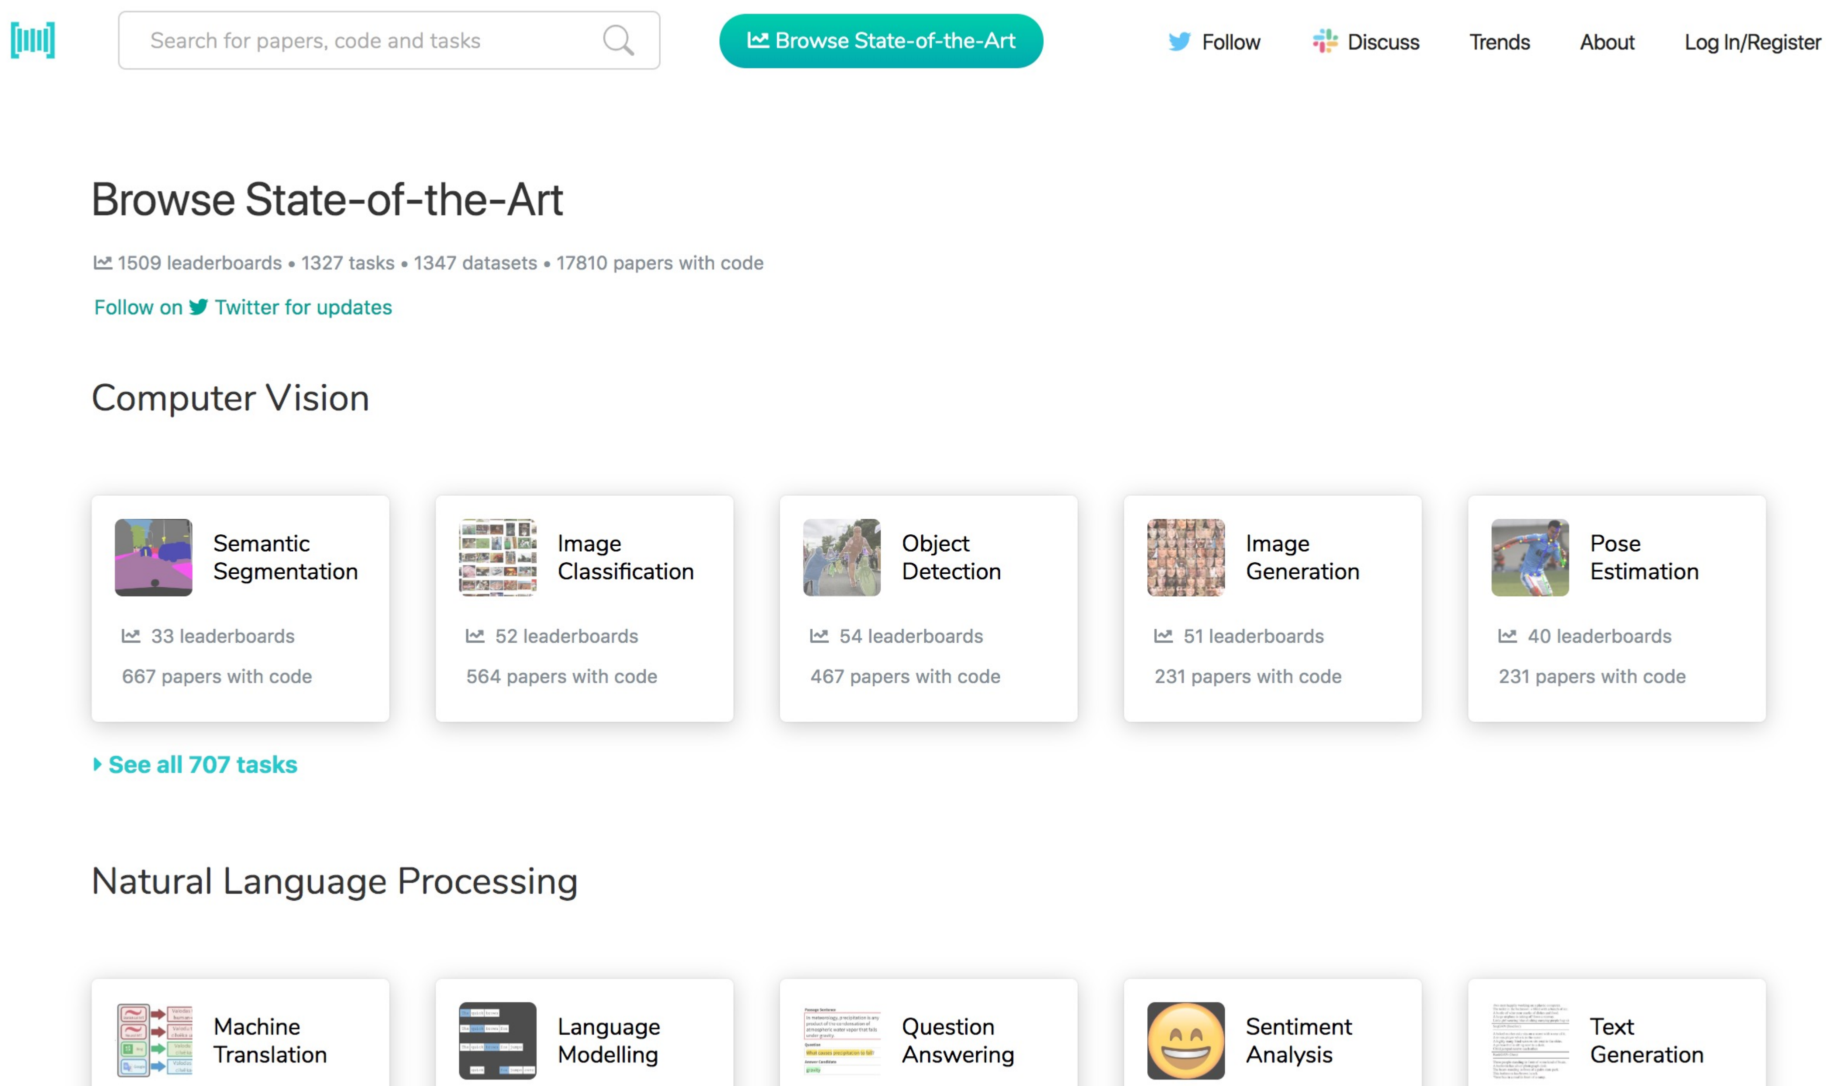

https://paperswithcode.com/sota

Aurélien Géron (2019),
Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems, 2nd Edition
O'Reilly Media, 2019

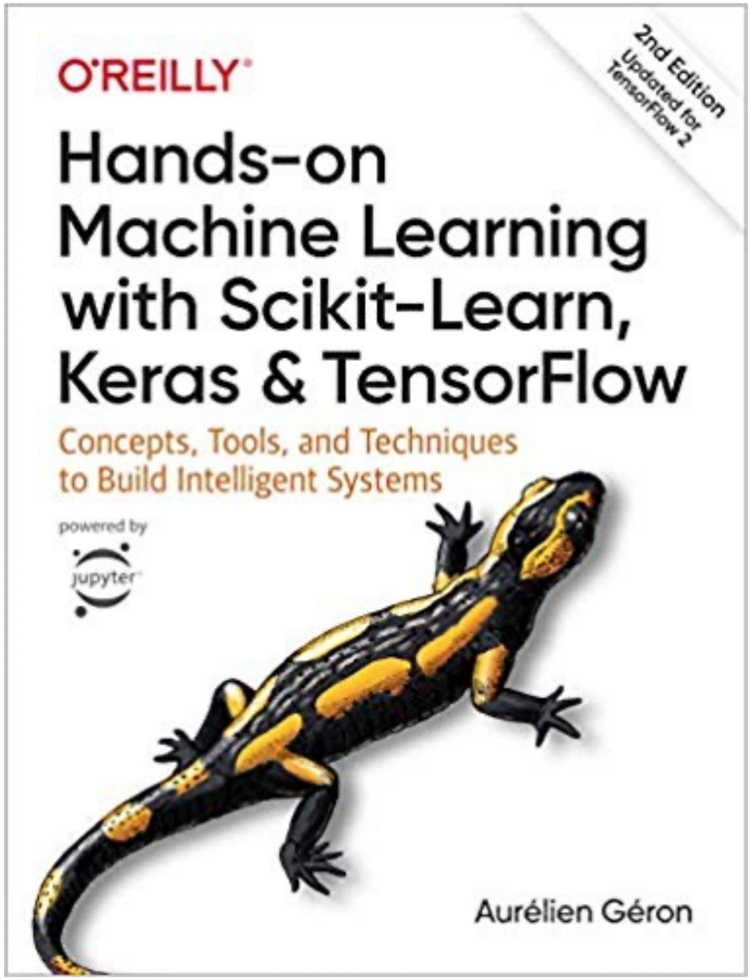

https://github.com/ageron/handson-ml2

Source: https://www.amazon.com/Hands-Machine-Learning-Scikit-Learn-TensorFlow/dp/1492032646/

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow

Notebooks

1. The Machine Learning landscape
2. End-to-end Machine Learning project
3. Classification
4. Training Models
5. Support Vector Machines
6. Decision Trees
7. Ensemble Learning and Random Forests
8. Dimensionality Reduction
9. Unsupervised Learning Techniques
10. Artificial Neural Nets with Keras
11. Training Deep Neural Networks
12. Custom Models and Training with TensorFlow
13. Loading and Preprocessing Data
14. Deep Computer Vision Using Convolutional Neural Networks
15. Processing Sequences Using RNNs and CNNs
16. Natural Language Processing with RNNs and Attention
17. Representation Learning Using Autoencoders
18. Reinforcement Learning
19. Training and Deploying TensorFlow Models at Scale

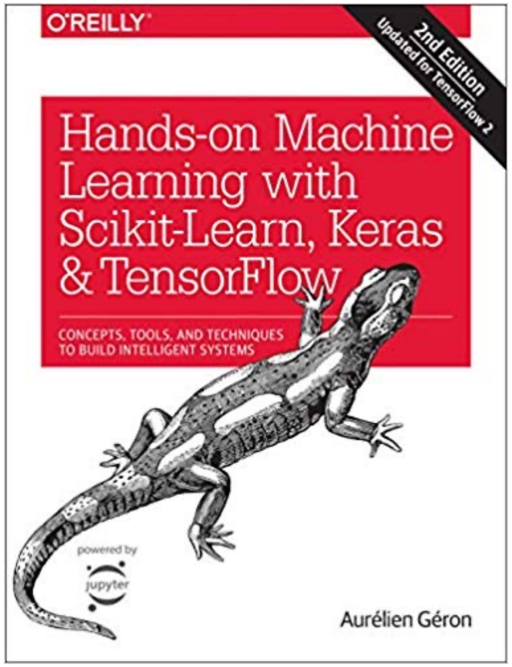

https://github.com/ageron/handson-ml2

Summary

- Robots
- Robotic Perception
- Planning and Control
- Planning Uncertain Movements
- Reinforcement Learning in Robotics
- Humans and Robots

References

- Stuart Russell and Peter Norvig (2020), Artificial Intelligence: A Modern Approach, 4th Edition, Pearson.
- Aurélien Géron (2019), Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems, 2nd Edition, O'Reilly Media.
- Min-Yuh Day (2021), Python 101, https://tinyurl.com/aintpupython101