UNIVERSITY OF LIVERPOOL

COMP 532

Machine Learning and BioInspired Optimization

Dr. Shan Luo

Department of Computer Science

shan.luo@liverpool.ac.uk

Today

- Admin & module info
- Modus operandi

- Overview module
- Introduction to nature-inspired computing

Module delivery

Shan Luo (lecturer)

H201C, George Holt building Building (second floor)

Email: shan.luo@liverpool.ac.uk

Jacopo Castellini (demonstrator)

Email: J.Castellini@liverpool.ac.uk

Module info page: see http://vital.liv.ac.uk

Slides and assignments will be available on vital as the course progresses

Tutorials/papers/additional notes

Module delivery

- Lectures times and locations:
  - 3-4 pm Monday, LIFS-LT1
  - 9-10 am Tuesday, BROD-305b
  - 2-3 pm Friday, ELEC-204 (E4)
  - *4-5 pm Friday, BROD-305a (tutorial with Jacopo)*

- There are three one-hour lectures and one hour of tutorials every week

- The tutorials will be organised very flexibly: working on assignments/tutorials

Module Aims and Objectives

In this module we focus on **learning agents** that interact with an **initially unknown world**. Since the world is **dynamic** this module will put strong emphasis on learning to deal with sequential data unlike many other machine learning courses.

Module Aims and Objectives

- To introduce and give an overview to the **state of the art** bio-inspired self-adapting methods.
- To enable students to not only learn to build models with reactive input/output mappings but **also build computer programs** that sense and perceive their environment, plan, and make optimal decisions.
- To familiarise students with **multi-agent reinforcement learning**, **swarm intelligence**, **deep neural networks**, **evolutionary game theory**, **artificial immune systems,** and **DNA computing**.
- To **demonstrate** principles of bio-inspired methods, provide indicative examples, develop problem-solving abilities and provide students with experience to apply the learnt methods in real-world problems.

Learning Outcomes

- A systematic understanding of bio-inspired algorithms
- In depth insight in the mathematics of biologically inspired machine learning and optimisation methods
- A comprehensive understanding of the benefits and drawbacks of the various methods
- Demonstrate knowledge of using the methods in real-world applications
- Practical assignments will lead to hands on experience

Module Syllabus (approximate)

- Parallel Problem Solving from Nature (2)
- Single-Agent RL (6)
- Multi-Agent RL (6)
- Swarm Intelligence (4)
- Deep Learning (6)
- Artificial Immune Systems (3)
- DNA Computing (3)

Module Syllabus (of others)

Calendar Topics may change somewhat. Exam times are firm. Readings are from Bishop.

Date	Topic	Reading	Assign
9/8	Intro: Basics, max likelihood estimation	1.1–1.5	HW0 Out
9/13	Intro: Bias/Variance, Basic Bayes	2.1–2.4; 3.2	
9/15	Regression: max likelihood, Ridge	3.1	HW0 Due; HW1 Out
9/20	Regression: Bayesian	3.2	
9/22	Classification: Discrim analysis; perceptron	4.1–4.2	
9/27	Classification: Logistic reg	4.3–4.5	
9/29	Classification: Support vector machines	7.1	HW1 Due
10/4	Kernels	6.1–3	HW1 Rev Due; HW2 Out
10/6	Neural nets: Basics, backprop	5	
10/11	*Holiday*		
10/13	Neural nets: optimization, conv, LSTM		
10/18	Neural nets: guest lecture		
10/20	**Exam 1 7:30 - 9:30 PM**		
10/25	Non-parametric prob models: Gaussian process	6.4	HW2 due; HW3 out
10/27	Ensemble methods	14	
11/1	Graphical models: message passing	8	HW2 rev due
11/3	Graphical models: EM	9	
11/8	Graphical models: sequential (HMM)	13.1–13.2	
11/10	Graphical modesls: NP Bayes, topic models		HW3 Due
11/15	Sampling methods	11	HW3 rev due; Proposal Due
11/17	Continuous latent variables	12	
11/22	Recommender systems		
11/24	*Holiday*		
11/29	Special topic		
12/1	**Exam 2 7:30 - 9:30**		
12/6	Nearest Neighbors/Density Estimation	2.5	
12/8	Reinforcement Learning		
12/13	Scalability		Project due

MIT course 6.867: Machine Learning

Module Syllabus (approximate)

W1
* 29 Jan: lec 3-4pm, PPSN 1
* 30 Jan: lec 9-10am, PPSN 2
* 2 Feb: lec 2-3pm, RL 1
- 2 Feb: tutorial 4pm-5pm

no class (make pairs of two)

W2
* 5 Feb: lec 3-4pm, RL 2
* 6 Feb: lec 9-10am, RL 3
* 8 Feb: lec 2-3pm, RL 4
* 9 Feb: tutorial 4pm-5pm

W3
* 12 Feb: lec 3-4pm, RL 5
* 13 Feb: lec 9-10am, RL 6
* 16 Feb: lec 2-3pm, MARL 1
* 16 Feb: tutorial 4pm-5pm, task 1

W4
* 19 Feb: lec 3-4pm, MARL 2
* 20 Feb: lec 9-10am, MARL 3
* 23 Feb: lec 2-3pm, MARL 4
* 23 Feb: tutorial 4pm-5pm, task 1

Module Syllabus (approximate)

W5
* 26 Feb: lec 3-4pm, MARL 5
* 27 Feb: lec 9-10am, MARL 6
* 2 Mar: lec 2-3pm, SI 1
* 2 Mar: tutorial 4pm-5pm, task 1

W6 (Task 1 Due)
* 5 Mar: lec 3-4pm, SI 2
* 6 Mar: lec 9-10am, SI 3
* 9 Mar: lec 2-3pm, SI 4
* 9 Mar: tutorial 4pm-5pm, task 1

W7
* 12 Mar: lec 3-4pm, DEEP 1
* 13 Mar: lec 9-10pm, DEEP 2
* 16 Mar: lec 2-3pm, DEEP 3
* 16 Mar: tutorial 4pm-5pm, make groups of 3

W8 (task 2 paper selection)
* 9 Apr: lec 4-5pm, DEEP 4
* 10 Apr: lec 12-1pm, DEEP 5
* 13 Apr: lec 2-3pm, DEEP 6
* 13 Apr: tutorial 4pm-5pm, task 2

Module Syllabus (approximate)

W9

* 16 Apr: lec 3-4pm, AIM 1
* 17 Apr: lec 9-10am, AIM 2
* 20 Apr: lec 2-3pm, DNA 1
* 20 Apr: tutorial 4pm-5pm task 2

W10 (TASK 2 due)

* 23 Apr: lec 3-4pm, DNA 2
* 24 Apr: lec 9-10am, WRAP UP
* 27 Apr: lec 2-3pm, STUDENT LEC 1
* 27 Apr: tutorial 4pm-5pm task 2

W11

* 30 Apr: lec 3-4pm, STUDENT LEC 2
* 1 May: lec 9-10am, STUDENT LEC 3
* 4 May: lec 2-3pm, STUDENT LEC 4
* 4 May: tutorial 4pm-5pm, STUDENT LEC 5

W12

* 7 May: Early May Bank Holiday
* 8 May: lec 9-10am, STUDENT LEC 6
* 11 May: lec 2-3pm, QUESTIONS?
* 11 May: tutorial 4pm-5pm QUESTIONS?

Assessment

- Written examination (75%)

- Two tasks (25%)
  - Task 1: Groups of 2
  - Task 2: Groups of 3

Assessment

- Task 1 on reinforcement learning, some programming and pen/paper exercises (10%)

- Task 2 student lecture/presentation on chosen topic (15%)

Useful resources

- Reinforcement learning
  - Courses:
    - David Silver, UCL[1]
    - Alexandre Proutiere, KTH[2]
  - Book
    - Reinforcement Learning: An Introduction, by Sutton & Barto[3]

[1] http://www0.cs.ucl.ac.uk/staff/d.silver/web/Teaching.html
[2] http://www.it.uu.se/research/systems_and_control/education/2017/relearn
[3] http://incompleteideas.net/book/the-book-2nd.html

Useful resources

- Deep learning
  - Courses:
    - Sebastian Trun, Udacity[1]
    - Fei-Fei Li *etc.*, Stanford (for visual recognition)[2]
  - Book
    - Deep Learning, by Goodfellow and Bengio[3]

[1] https://eu.udacity.com/course/deep-learning--ud730
[2] http://cs231n.stanford.edu/
[3] http://www.deeplearningbook.org/

Introduction
BioInspired Machine Learning

BioInspired Robotics

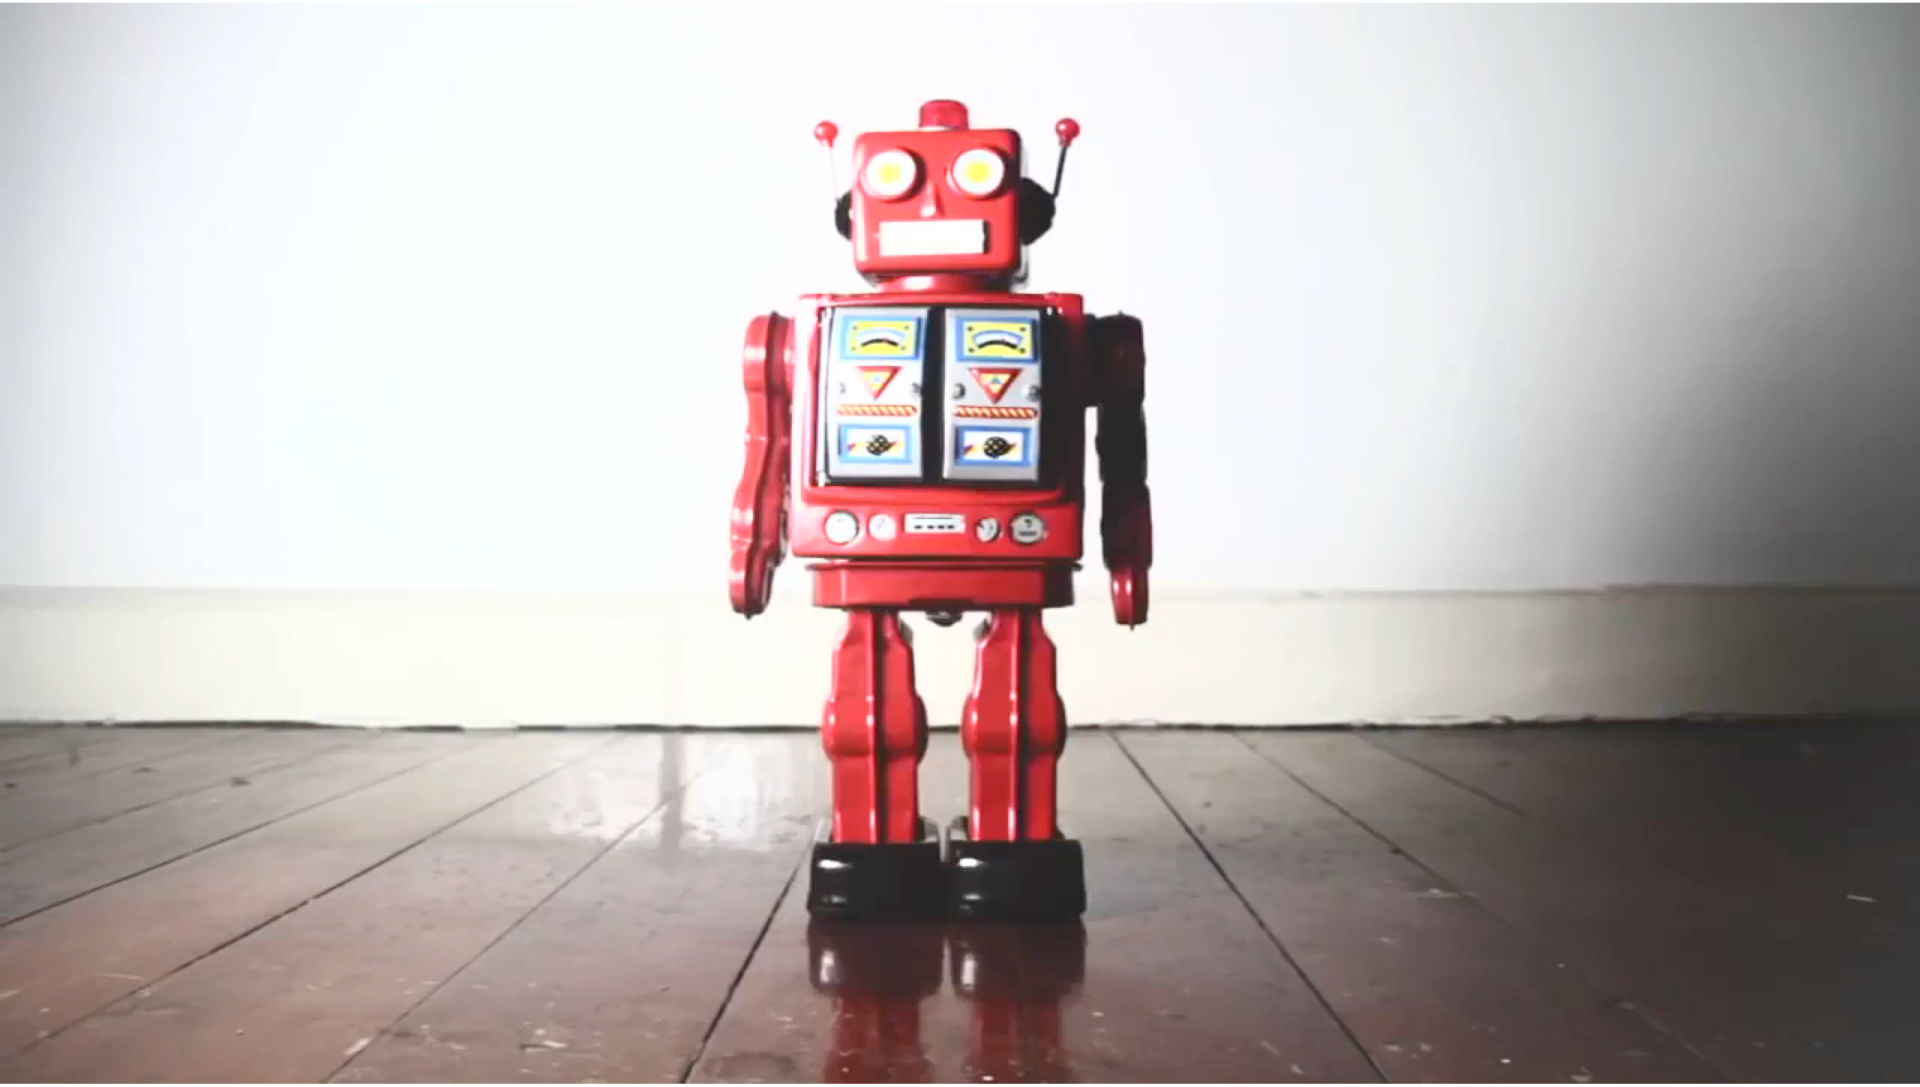

https://www.youtube.com/watch?v=-ojisOW4jzg

Bio or NatureInspired learning

- Nature Inspired Computation
- Examples:
  - Reinforcement Learning
  - Membrane Computing
  - DNA Computing
  - Evolutionary Computation
  - Swarm Intelligence
  - Evolutionary Game Theory
- BioInspired Machine Learning for Adaptive Agents/Systems
- Very young field – how young?
- What is Machine Learning?

Bio or NatureInspired learning

- Tom Mitchell on Machine Learning:

"How can we build computer systems that automatically improve with experience, and what are the fundamental laws that govern all learning processes?"

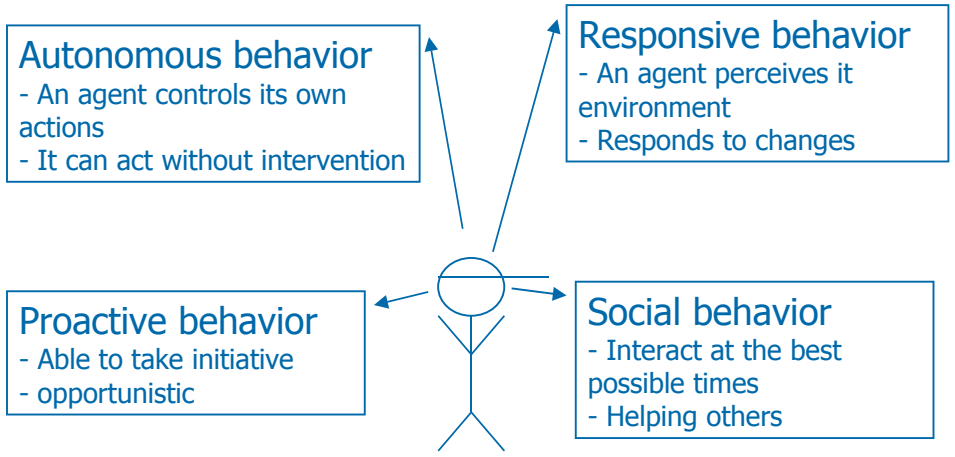

Bio or NatureInspired learning

BioInspired Machine Learning

- To optimise its products, nature exploits genetic operators: selection, mutation, crossover
- There are those mechanisms we can “easily” mimic
- Can we model emergent collective intelligence from local simple interactions at level of:
  - Cells
  - Neurons
  - Insects
- How do simple creatures like ants/bees organise themselves to show great intelligence at a group level?

Bio or NatureInspired learning

- This course is **NOT** about **Machine Learning** for **Datamining** or its **Applications**

  (Although it will prove useful for this)

- But it is about **agents** situated in a (potentially) unknown **environment**, **learning how to behave** and **optimize** its action in order to achieve its goals, taking inspirations from **nature**.

Bio or NatureInspired learning

- Some important notes:
  - We are not necessarily mimicking nature
  - Inspirational

Natural systems

Artificial systems

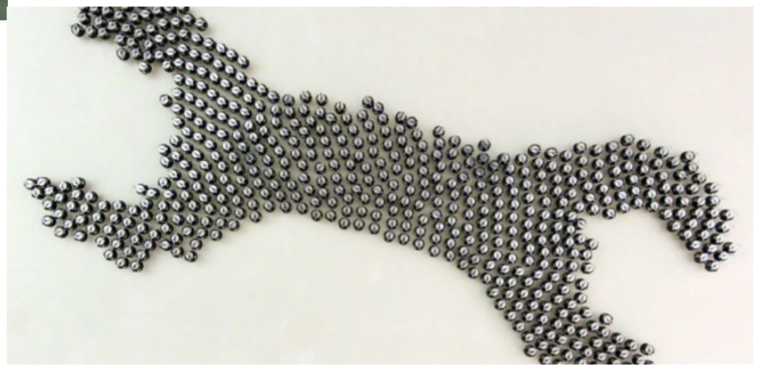

Bio or NatureInspired learning

Applications

- Making a robot walk or speak with humans or grasp an object

or… flip pancakes

[1] https://www.youtube.com/watch?v=W_gxLKSsSIE&list=PL5nBAYUyJTrM48dViibyi68urttMlUv7e

Bio or NatureInspired learning

Applications

- Portfolio optimisation
- Playing games better than humans
- Search engines
- …..

Summary

- Introduction to course setup
- Topic of the course
- Introductory examples