



# Artificial Intelligence Workshop



# Welcome and Introduction

---



- Who are we?
- Why are we here?
- Agenda



# What is AI?



**Strong AI**

**Weak AI**



# Examples



**Personal Assistants**



**Autonomous Vehicles**



**Translation Aids**



**Medical Imaging Evaluation**



**Facial Recognition**

# Machine Learning


## Definition:

“the use and development of computer systems that are able to learn and adapt without following explicit instructions, by using algorithms and statistical models to analyse and draw inferences from patterns in data.”

(Oxford Languages via Google)

# Supervised Learning

- Known goal, e.g. find the dogs.
- People label the input data.
- The algorithm is trained on that input data.
- After that training is complete, the algorithm is used on new data.








# Unsupervised Learning


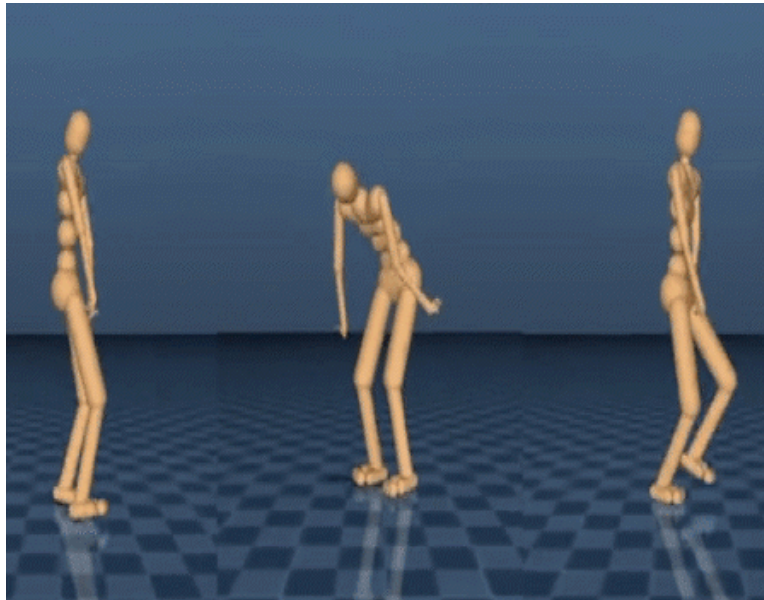

- Goal not supplied.
- The algorithm receives data that is not labelled.
- The algorithm finds the patterns.



# Reinforcement Learning

- Learning through trial and error.
- ‘Rewards’ for good behaviour.
- ‘Punishments’ for bad behaviour.





Images used with permission of DeepMind (<https://www.deepmind.com/blog/producing-flexible-behaviours-in-simulated-environments>)


# Example: DeepMind

- AI against human.
- The AI masters games without even being told the rules.
- Significant step towards general intelligence.







## Domains




## Knowledge




**AlphaGo** becomes the first program to master Go using neural networks and tree search (Jan 2016, Nature)



**AlphaGo Zero** learns to play completely on its own, without human knowledge (Oct 2017, Nature)



**AlphaZero** masters three perfect information games using a single algorithm for all games (Dec 2018, Science)



**MuZero** learns the rules of the game, allowing it to also master environments with unknown dynamics. (Dec 2020, Nature)

Images used with permission of DeepMind  
<https://www.deepmind.com/blog/muzero-mastering-go-chess-shogi-and-atari-without-rules>

# Some Problems with AI Systems

- The outcome is often not known or guessable up front.
- The reason for the outcome often can't be explained.
- A real-world AI system has:
  - **many parts** – hard to work out what went wrong.
  - **many participants** – hard to work out who is responsible.



# Case Study 1

---

## The Bank Loan





# Case Study 1: Predicting the Outcome

- We're designing a system for processing loan applications at a bank.

## Round #1:

- We don't use AI – we will use a 'deterministic algorithm' instead.
- What will it look like?



# Case Study 1: Predicting the Outcome

## Round #2:

- Supervised machine learning system.
- We feed it old loan applications and tell it whether borrower paid the loan back in time.
- The system builds a model internally.
- After we are happy it is getting it right enough of the time, the system goes live.



# Case Study 1: Predicting the Outcome

- Eliza applies for a loan.
- The AI says: 'No.'
- Eliza: 'That's ridiculous – Why won't you lend to me?'
- Bank doesn't know why.
- Why does that matter?





# Case Study 2

---

## The AI Lawyer



# Case Study 2: The AI Lawyer

- AI chatbot gives free legal advice to members of the public.
- Presents via a website: you tell it your problem, it provides the answer.
- Warns you that: it is an AI tool, not a real lawyer; and that it's experimental.
- The AI chatbot gives the wrong answer.







# Case Study 3

---

## The AI Journalist



# Case Study 3: The AI Journalist

- Online news website can't afford humans and is using an AI to generate articles.
- One of those articles accuses a well-known article of racism.
- The actor sues...





# Case Study 4


---

## The AI Driver



# Case Study 4: The AI Driver

- Autonomous vehicle is driving through town.
- It is fully self-driving but legally the human driver needs to be ready to intervene.
- It accelerates rapidly and hits a pedestrian before the human driver realises what is happening.
- Turns out: a hacked advertising board showed a false speed limit.





# What Next?

# Being a Tech Lawyer

- Thinking about new situations and how the law applies to them is a key part of being a tech lawyer.
- If you are tech lawyer, these sorts of problems may be your bread and butter.
- If being a lawyer sounds like it *might* be for you, consider:
  - Joining a debate club (even if that sounds scary at first);
  - Coming to Court to see advocates in action (free and public);
  - Listening to podcasts about the law (try [techlaw.chat](#) & others);
  - Finding work experience – look across the the legal sector (you never know what might interest you and it is great to see how it fits together).



# The Legal Profession wants Diversity

- There are stuffy old stereotypes of lawyers and judges.
- The legal profession is changing. No doubt there is still work to be done and the profession is committed to that.
- Already there are serious and effective drives to improve access.
- The legal profession needs to be representative. It wants to attract the very best people: they don't have a particular stereotype.

# Thank you

---



**Content:** Dr Matthew Lavy and Iain Munro, 4 Pump Court

**Videos and images:** [https://www.fiverr.com/next\\_animation](https://www.fiverr.com/next_animation)

**Voice Over:** Dr Jesse Mears (Transformation, Data Strategy and Service Design Consultant)