


# AI and the Law—looking beyond the hype of AI

Alistair Knott (Dept of Computer Science, Otago)  
John Zerilli, Joy Liddicoat (Faculty of Law, Otago)




# The AI and Law in NZ Project



- 3 year project, funded by the NZ Law Foundation
- Phase 1 topic: uses of AI in NZ government departments
- Phase 2 topic: AI and employment in NZ

# Our Phase 1 report is coming out next week!



# Today's presentation

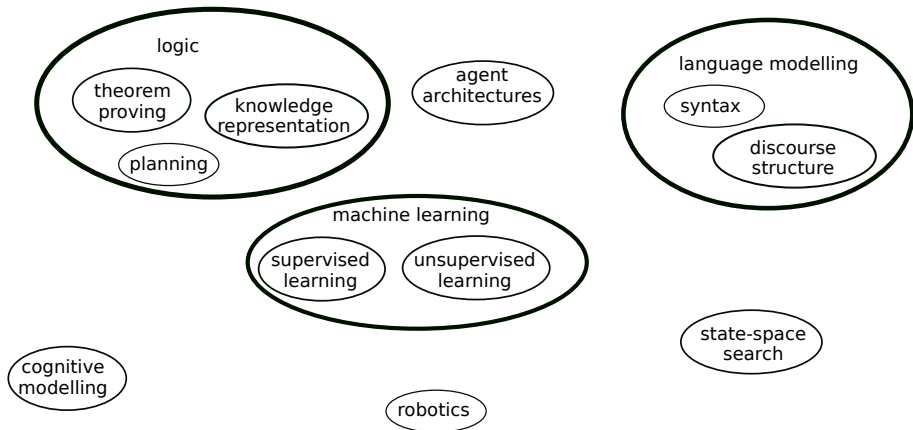
I will introduce AI systems. . .

John & Joy will discuss various legal issues that arise with AI systems.


# An introduction to AI systems

'In 15 minutes, for lawyers'

# The field of AI



# The field of AI



# The field of AI




supervised  
learning



# What's a supervised learning system?

# What's a supervised learning system?

A computer program that learns to predict some *output variable*, from a set of *input variables*.




# What's a supervised learning system?

A computer program that learns to predict some *output variable*, from a set of *input variables*.




# What's a supervised learning system?

A computer program that learns to predict some *output variable*, from a set of *input variables*.




# What's a supervised learning system?

A computer program that learns to predict some *output variable*, from a set of *input variables*.




# What's a supervised learning system?

A computer program that learns to predict some *output variable*, from a set of *input variables*.



# What's a supervised learning system?

A computer program that learns to predict some *output variable*, from a set of *input variables*.




# What's a supervised learning system?

A computer program that learns to predict some *output variable*, from a set of *input variables*.

input variables

legal  
clause



output variable


clause type  
- indemnity..  
- confidentiality..  
...



How does the learning happen?


## How does the learning happen?

In supervised learning, the program is given a set of *training examples*, where the input and the output are both known.



## How does the learning happen?

In supervised learning, the program is given a set of *training examples*, where the input and the output are both known.



## How does the learning happen?

In supervised learning, the program is given a set of *training examples*, where the input and the output are both known.



## How does the learning happen?

In supervised learning, the program is given a set of *training examples*, where the input and the output are both known.



## How does the learning happen?

In supervised learning, the program is given a set of *training examples*, where the input and the output are both known.

- Often, the output variable is the result of *human judgement*.



## How does the learning happen?

In supervised learning, the program is given a set of *training examples*, where the input and the output are both known.

- Often, the output variable is the result of *human judgement*.




After training, a good program will have learned *general principles* about how to map inputs onto outputs—which allow it to perform on *unseen* inputs.

## How does the learning happen?

In supervised learning, the program is given a set of *training examples*, where the input and the output are both known.

- Often, the output variable is the result of *human judgement*.




After training, a good program will have learned *general principles* about how to map inputs onto outputs—which allow it to perform on *unseen* inputs.




## A simple predictive system

Say our training examples map the input variable  $x$  onto the output variable  $y$ .



## A simple predictive system


Say our training examples map the input variable  $x$  onto the output variable  $y$ .



- A simple predictive model might find the straight line that most closely matches the training examples. (**Linear regression**)

## A simple predictive system

Say our training examples map the input variable  $x$  onto the output variable  $y$ .



- A simple predictive model might find the straight line that most closely matches the training examples. (**Linear regression**)
- The line lets us predict  $y$  values even for values of  $x$  that weren't in the training examples.

## Predictive systems aren't new!

Linear regression is a statistical technique that has been around since the 19th century.

- It's been widely used, both in industry and in government. (Especially since the 1950s.)

## Predictive systems aren't new!

Linear regression is a statistical technique that has been around since the 19th century.

- It's been widely used, both in industry and in government. (Especially since the 1950s.)

In many ways, the story of predictive models in industry and government has been one of *continuity*.

## Predictive systems aren't new!

Linear regression is a statistical technique that has been around since the 19th century.

- It's been widely used, both in industry and in government. (Especially since the 1950s.)

In many ways, the story of predictive models in industry and government has been one of *continuity*.

But there are two newer developments.

## Predictive systems aren't new!

Linear regression is a statistical technique that has been around since the 19th century.

- It's been widely used, both in industry and in government. (Especially since the 1950s.)

In many ways, the story of predictive models in industry and government has been one of *continuity*.

But there are two newer developments.

- Current predictive systems are much more *complex*.

## Predictive systems aren't new!

Linear regression is a statistical technique that has been around since the 19th century.

- It's been widely used, both in industry and in government. (Especially since the 1950s.)

In many ways, the story of predictive models in industry and government has been one of *continuity*.

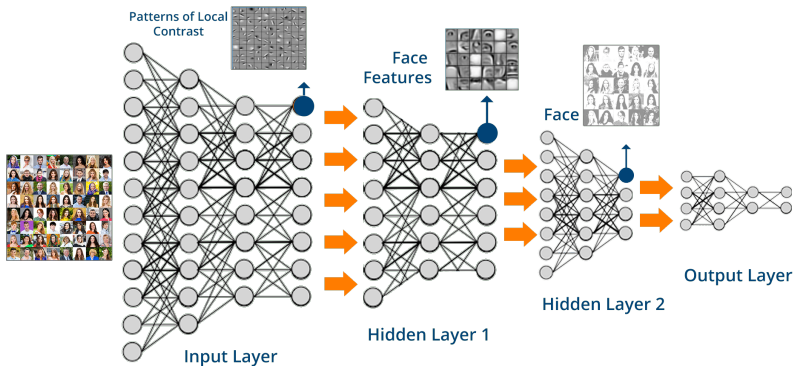
But there are two newer developments.

- Current predictive systems are much more *complex*.
- There's *more data* available for current systems to train on.



# A modern predictive system

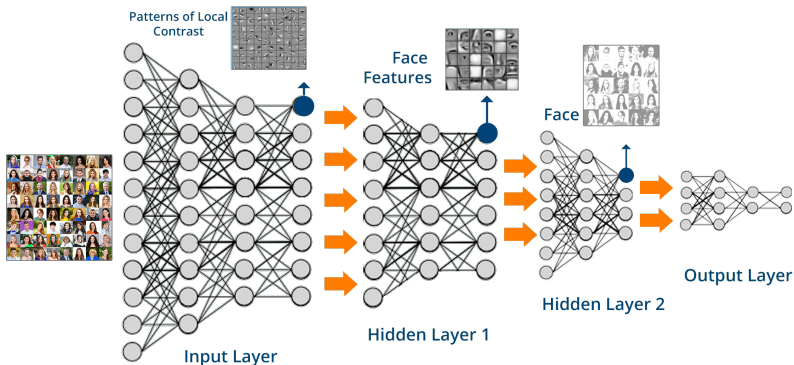
Here's a schematic **deep neural network** for image classification.



# A modern predictive system

Here's a schematic **deep neural network** for image classification.

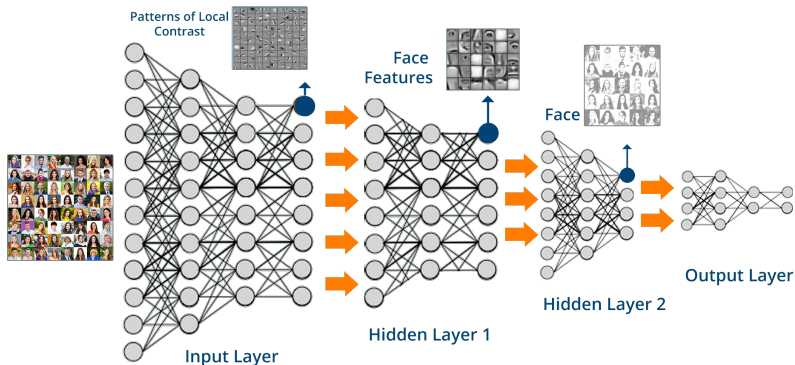
- The circles are neuron-like 'units', linked by 'connections' whose 'weight' can be adjusted.



# A modern predictive system

Here's a schematic **deep neural network** for image classification.

- The circles are neuron-like 'units', linked by 'connections' whose 'weight' can be adjusted.
- Learning involves adjusting all the weights, to reduce the network's 'error' on the training examples.



## Working with modern predictive systems

Deep networks (and similar methods) are the best learning devices we have.

# Working with modern predictive systems

Deep networks (and similar methods) are the best learning devices we have.

- But their operation is impossible to understand!  
Millions of units, billions of connections. . .

# Working with modern predictive systems

Deep networks (and similar methods) are the best learning devices we have.

- But their operation is impossible to understand!  
Millions of units, billions of connections. . .

How should we use them??  
oversee them??  
regulate them??