Machine Learning, AI, and people

Jim Smith

Professor of Interactive Artificial Intelligence

University of the West of England

What this talk is about

- How technology is changing AI
- How AI is changing as it comes out of the lab into the real world

And what it’s not …

Brief, very personal, histo

- Human “experts” provide “world model” in form of facts and rules
- Given a new situation (facts) system deduces new facts (diagnosis/actions/plans)
- Computationally hard to scale,
- Brittle
- Still useful in restricted scenarios:
- E.g. Simulated Air Traffic Control Environments
  - Fine to use ML for speech recognition
  - But ***absolutely not*** for speech & action generation
- Chatbots as front end to knowledge base
  - Consistent,
  - Separate *how* and *what* of questions
- Both driven by understanding of **human** requirements

Machine Learning:

- Human “experts” provide learning algorithm
- Given new examples the system learns a model
  - Learning often computationally expensive
- That can (hopefully) generalise to new cases
  - Deployed models may be computationally cheap
- For example Artificial Neural Networks
  - Hiatus between late 60s and early 1980s while people figured out how to have more than one layer of perceptrons
  - Hiatus between mid 1990s and mid 2000s while hardware caught up
  - “Deep learning / Deep Belief Networks now hugely successful
  - Variants such as LSTM (recurrence) and convolutional neural networks greatly widen scope
- GIGO still applies! So human intervention has been needed (MS racist chatbot, Facebook’s fake news, Google’s holocaust denial sites)


0:47 / 1:07

ML: allowing users to provide more forms of input

- EU project ”DynaVis” looked at human factors in ML for visual inspection of manufactured products
- This example is from partner Sony DADC
- Can’t always be sure from a picture
  - So we designed an interface that asked them how confident they were
- And then took account of that in training
- Error rates reduced by 50%!

(*Lughofer et al, IEEE Trans. Systems, Man & Cybernetics, 2009)*

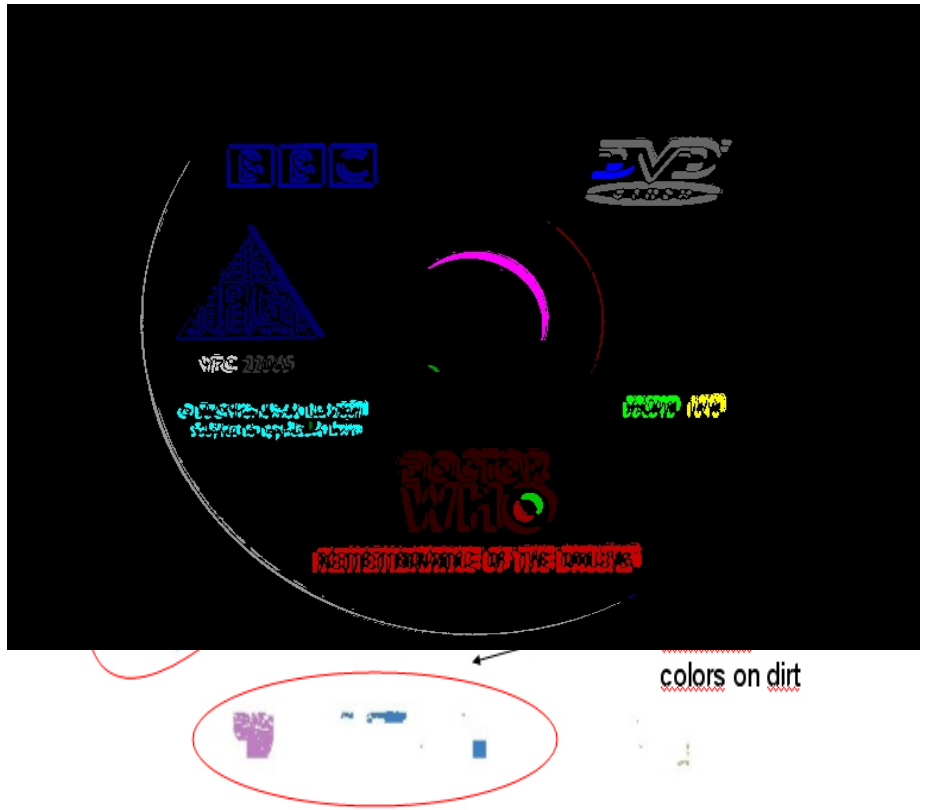

- **Affective computing** is the study and development of systems and devices that can recognize, interpret, process, and simulate human affects. It is an interdisciplinary field spanning computer science, psychology, and cognitive science. (wikipedia)

Lots of other applications being developed combining AI with everyday sensors:

- Providing input suitable to cognitive load (e.g. driver assistance),
- Healthcare (e.g. detecting huge range of diseases and conditions using data from smartphones)

Recent results suggests we can predict users' confidence as they perform visual tasks

1. Capture their gaze traces using eye-trackers
2. Find a suitable representation as an image
3. Use deep convolutional neural networks to predict confidence
    1. yes/no 89% or on a scale of 1-5 (MSE<1) (*Smith et al 2016,in review*)

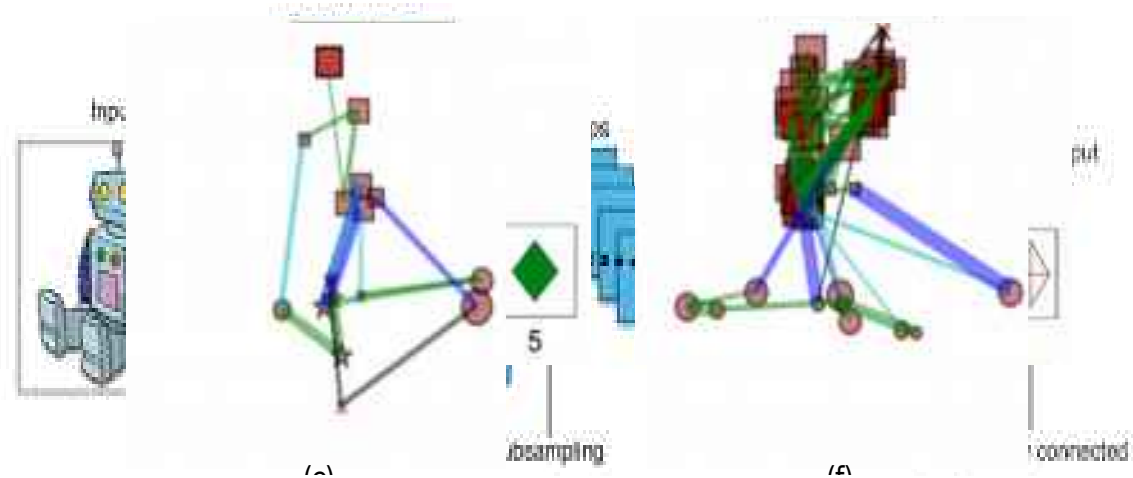

Similar history for optimisation

- AI approaches are now the state of the art for many problems
  - (Eiben and Smith, Nature 2016),
  - ACM-"Humies" sessions on "Human Competitive Results"
  - Patents have been awarded for "evolved" designs
- But these have traditionally relied on human-provided measures quality
- Various approaches challenge the need to state preferences up-front
  - e.g.: ask user for subjective score
- But why stop there? For example, iACO (*Simons et al 2014*)
  - allows multiple forms of user input, and
  - ML models the user's preferences to optimise interactions

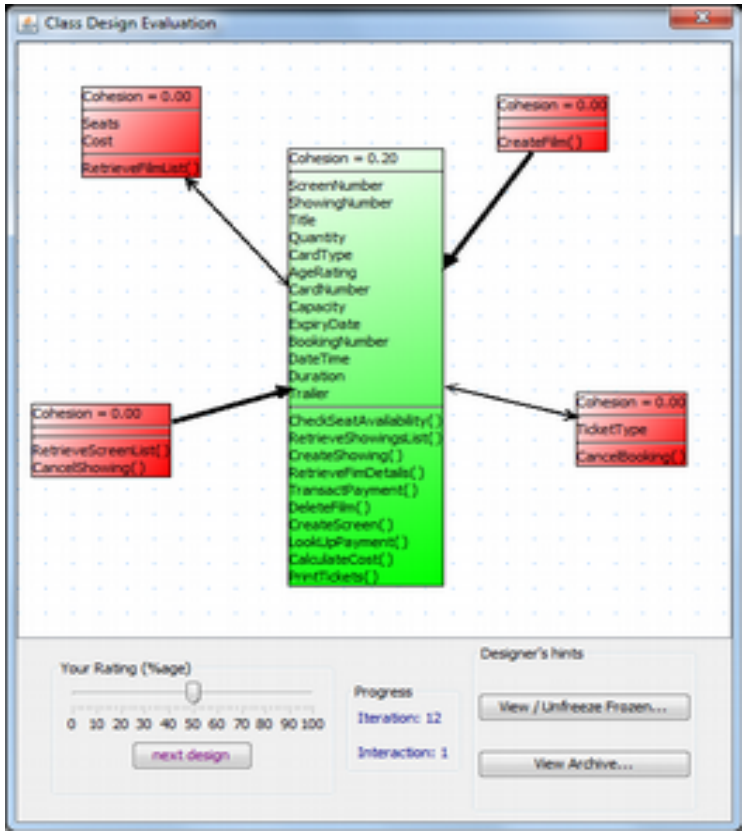

(One) vision for AI

- Advances in hardware, data handling, and algorithms let us:
  - learn models from massive disparate sources of information
    to enrich people's life experiences
  - Automatically produce tailored artefacts – whether physical, or virtual
    i.e. evolving a new bottle holder, or the firmware update for ***your*** fridge
  - Explore models through greater visualisation/sonification/haptic interfaces
  - Non-intrusively detect human reactions
- Increasingly different forms of AI are fused to create multi-modal environments for collaborative human-machine problem solving
  - Building on the strengths, and recognising the weaknesses, of each partner.
- (robo) ethicists are urgently needed to help a sensible informed debate about the implications for
  - Employment, safety, privacy, …