Multimodal data and machine learning by 2040?

Massimo Piccardi

EDTAS 2017, 10 April 2017

In 2040

- The only certainty about 2040 is that it is in 23 years from now
- Some predictions will be exceeded, others will be unmet
- A 2040 scenario is just an exercise

A 2040 scenario, then

- Driverless cars will be commonplace (like in aviation, accidents will not arrest development)
- Drones will be everywhere (ditto)

- Gadgets such as Google Home and Amazon Echo will completely replace remote controls
- Augmented reality will be a popular interface to a hybridised physical + digital world

- I feel social media as we know them will disappear, and turn into some kind of *digital social continuum*

Sensors

- Sensor data from:

Vehicles Street cameras Building sensors

Wearable sensors (GPS, accel, gyro, compass, light, temperature, altimeter, heart rate, glucose…)

Cameras & mics on smartphones and AR glasses

Social media Apps Online transactions

Humanoid robots

Sensors

An intimidating sensor?
(*pareidolia*?)

The same sensor - with legs

(courtesy of Aldebaran Robotics)

Intelligence

- Leaving aside important aspects about the data such as: privacy, ethical use, security, availability, integrity, reliability, communications, storage and more so, the question is: how to make sense of them?

- Two main trends in machine learning:

  – **Deep learning**

  – **Reinforcement learning**

Deep learning

- **Deep learning** is a renaissance of neural networks, just with more layers
- The intermediate layers compute various re-codings of the measurements and the output layers can tap into any of the existing conventional methods (logistic regression, CRFs, SVMs…)
- What does it require?
  - **Very large**, **annotated** training sets
  - Clearly specified loss functions

Reinforcement learning

- The requirements of deep learning are somehow restrictive:
  - annotating the “ground truth” of large datasets is expensive, and often subjective
  - specifying precise loss functions can be impractical
- **Reinforcement learning** offers a different paradigm, more suitable for online/evolutionary scenarios

Reinforcement learning

- The underlying model can be described as a partially-observable Markov decision process:

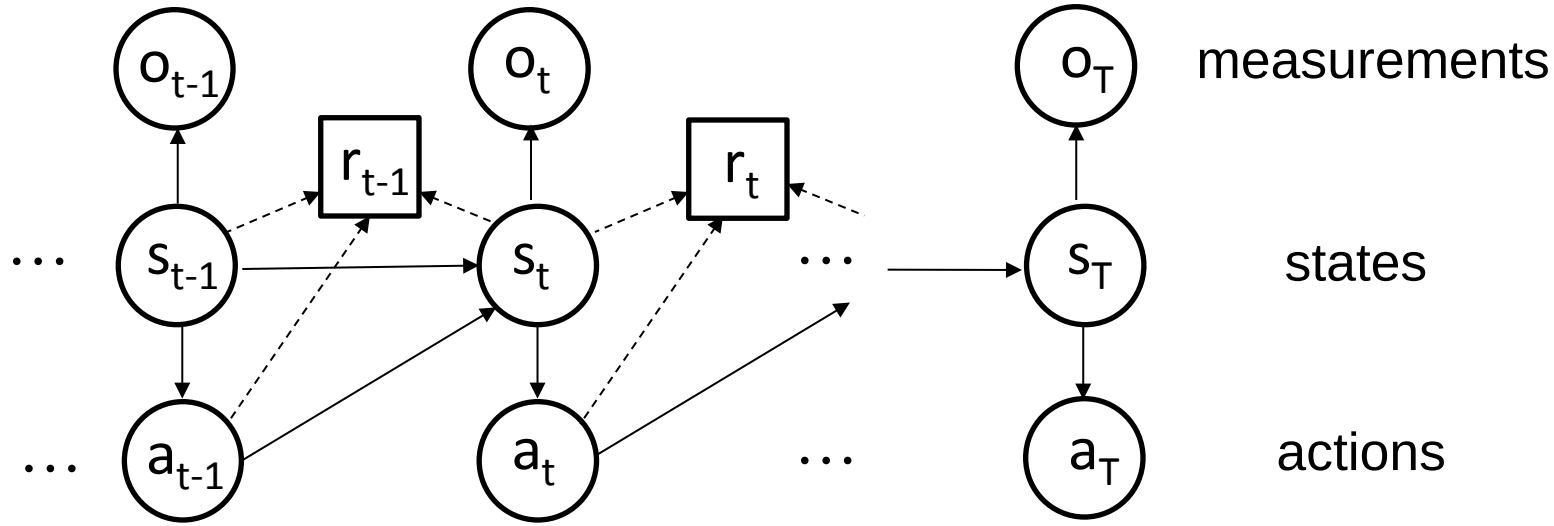

- Many variations on the theme: temporal differences, imitation learning, Q-learning, SARSA, “bandits”…

Reinforcement learning

- The notion of reinforcement learning tags spontaneously with robotics, but it is also used in computer vision, natural language processing, game playing, web design, business decisions and more
- This in reflection that an agent can be either physical or completely virtual
- **Deep reinforcement learning** mixes RL with DL

Q: will agents be *sentient* by 2040?

Will agents become morally relevant?

Do Artificial Reinforcement-Learning Agents Matter Morally?

by Brian Tomasik

Written: Mar.-Apr. 2014; last update: 29 Oct. 2014

Abstract

Artificial reinforcement learning (RL) is a widely used technique in artificial intelligence that provides a general method for training agents to perform a wide variety of behaviours. RL as used in computer science has striking parallels to reward and punishment learning in animal and human brains. I argue that present-day artificial RL agents have a very small but nonzero degree of ethical importance. This is particularly plausible for views according to which sentience comes in degrees based on the abilities and complexities of minds, but even binary views on consciousness should assign nonzero probability to RL programs having morally relevant experiences. While RL programs are not a top ethical priority today, they may become more significant in the coming decades as RL is increasingly applied to industry, robotics, video games, and other areas. I encourage scientists, philosophers, and citizens to begin a conversation about our ethical duties to reduce the harm that we inflict on powerless, voiceless RL agents.

A sample of our work

- Structured classification of text in non-English languages

Ezāfe	POS-tag	Inflexion	Token	NER-tag	Transliteration
O	N	N,SING,SURN	سید	B-PERS	Seyed
EZ	N	N,SING,PR,GEN	محمود	I-PERS	Mahmoud
O	N	N,SING,PR	محدث	I-PERS	Mohadess
EZ	N	N,SING,COM,GEN	مدیر	O	manager
EZ	N	N,SING,COM,GEN	اکتشاف	B-ORG	discovery
		N,SING,COM,GEN	شرکت	I-ORG	Company
		ADJ,SIM,GEN	ملی	I-ORG	National
		N,SING,COM,GEN	نفت	I-ORG	Oil

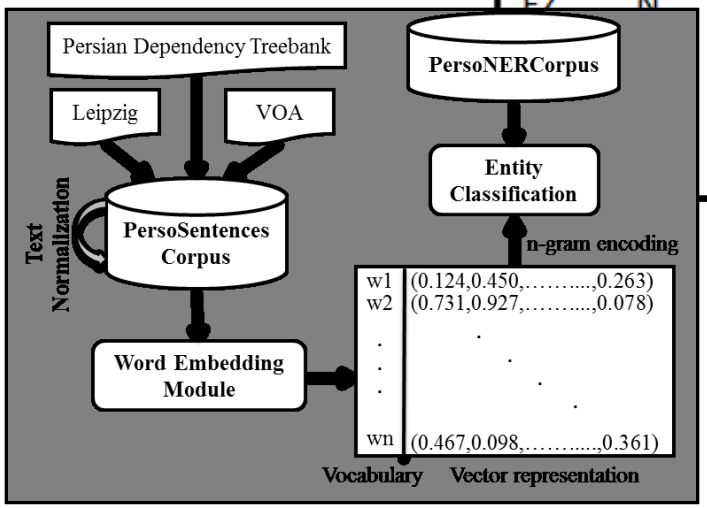

A sample of our work

- Structured co-prediction of actions and summaries in videos

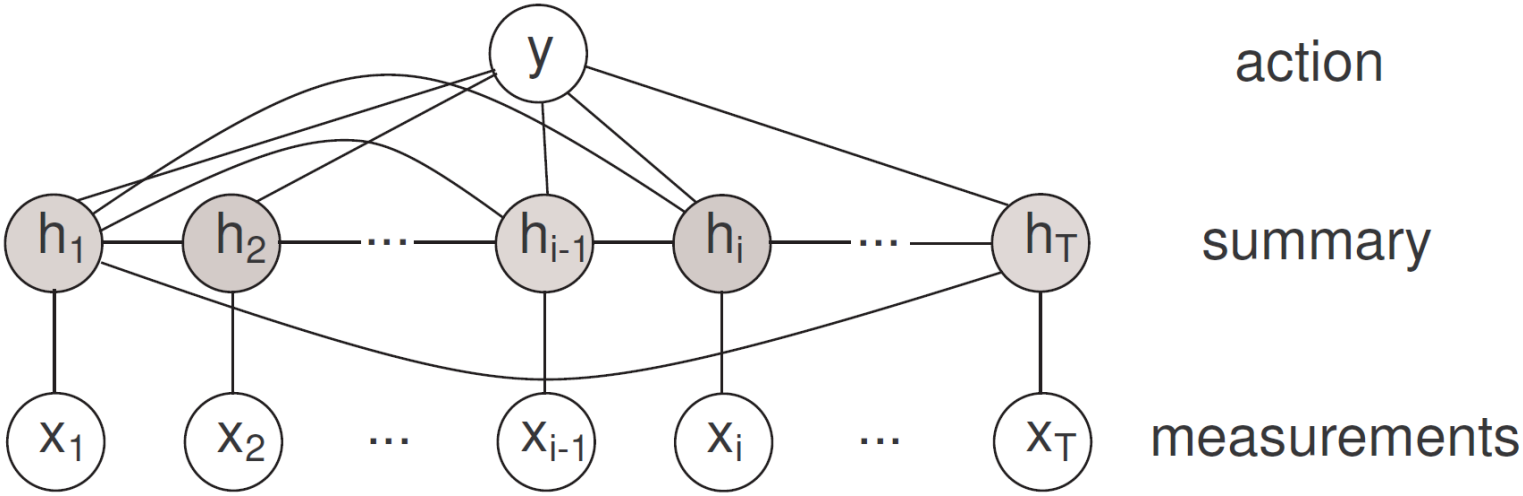

A sample of our work

recognition as *relationships*!

latent objects: “sky”, “road”, “desktop”, “coffee mug”…

superpixels: small, homogeneous regions in the image

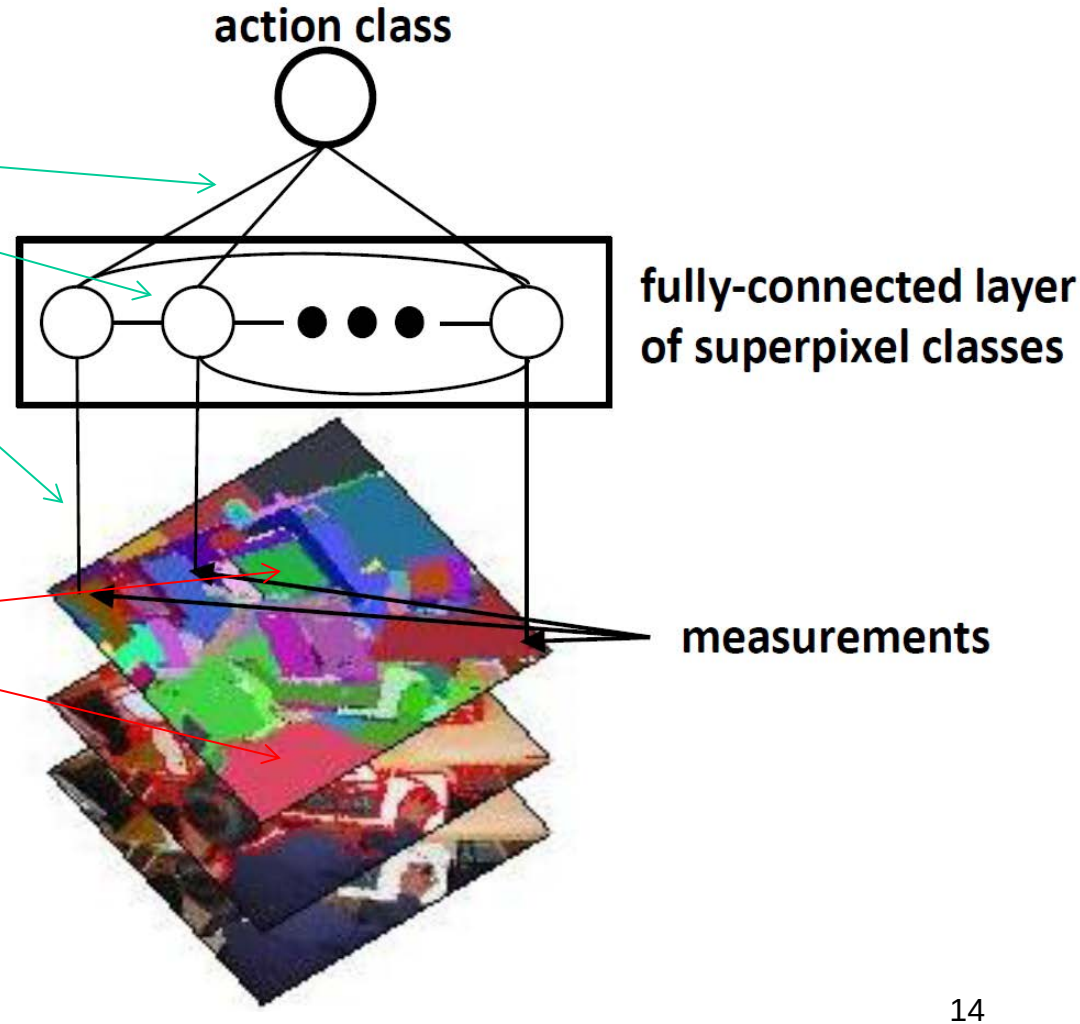

Federated approaches

- Research is becoming increasingly multi-disciplinary
- This is why in our centre (UTS:GBDTC), we try to cover all facets of big data:
  - Communications
  - Networks
  - IoT
  - Security
  - Machine learning (NLP, computer vision, multimedia)
  - Cloud computing
- More is needed

Thank you & let's stay in touch

- Prof. Massimo Piccardi

  Massimo.Piccardi@uts.com.au

  Global Big Data Technology Centre

  http://www.bdt.uts.edu.au/

- UTS' research centres – covering AI, digital innovation, behavioural economics and much more: https://www.uts.edu.au/research-and-teaching/our-research/research-centres/research-centres

  David Beins, Executive Manager (Centre Development)