

DRAFT: Artificial Intelligence at Edinburgh

March 31, 2019

Note: The target audience is general: e.g. potential students, policymakers, etc.

Also, the final aim is for a webpage not a pdf, but this is convenient form to work on the text.

Artificial Intelligence (AI) is the field concerned with the automation of perceiving, reasoning and behaving in an intelligent (i.e. rational) manner, or like humans. Intelligence involves abilities like reasoning and problem solving, learning and memory, perception, taking actions and communication. Edinburgh has a long tradition of work in Artificial Intelligence, going back to the 1960s when it was one of the few centres in the world working on AI. The School of Informatics produced more world-leading and internationally excellent research ranked (4* and 3*) than any other university in the UK in the Research Excellence Framework (REF) 2014 assessment for computer science and informatics. Of the six research institutes in the School, four of these work on AI topics. Informatics is thus the leading centre for AI in the UK, with research covering a large fraction of the field. We are also a founder member of the Alan Turing Institute, the UK's national institute for data science and artificial intelligence.

One strand of AI research concerns *perception*, where we wish to understand signals in the world such as images, video and speech. For speech the aim is to turn the acoustic signal into words (1), while for vision we may wish to identify objects, agents and their actions (2). We can also combine modalities, e.g. with the joint analysis of speech and video data.

As well as perceiving the world, we may want to *reason* and answer questions about it. Logic provides a powerful framework for this task, as applied e.g. to mathematical theorem proving (3). Another important aspect is being able to reason properly under uncertainty when there is incomplete information, for example in a medical application where there may be multiple possible diagnoses, some more likely than others (4).

Artificial intelligence systems need to *adapt* and *learn* (4). In supervised machine learning labelled data is provided, along with an evaluation of how well the task is being achieved, and the system adapts to improve its performance. For example, a system can be trained to label the objects in an image, like bird or car. We also develop algorithms for unsupervised learning, to detect such structure in the data without the benefit of labels.

We build agents that can *act* in their environment to achieve useful tasks, be it the physical world for robots, or virtual worlds for softbots. Such agents need to learn, reason and plan, whether they are acting fully autonomously, or in collaboration with humans and computer-based entities. Controlling dextrous motion and physical interaction (5) provides the basis for intelligent action and control of robots (6). Generating such actions can also be applied to virtual agents in computer graphics and animation (7). Concepts of robust autonomy and decision making can be extended to other domains such as predictive modelling and decision making in energy and environmental systems (8).

Language is a critical route by which we communicate and convey information to each other. We work on algorithms that understand the structure of sentences (9), determine the

meaning of the linguistic input, and then generating appropriate responses (10). Natural language understanding and generation can be deployed in range of areas; particular strengths in Edinburgh include machine translation (11), document summarization (12), and information retrieval, including from social media (13).

The study of natural systems can provide useful information for the development of AI, and vice versa. This can be at the level of computational neuroscience (14) which studies how the brain processes information; inspiration for behavioural control from insects (15); computational psychiatry, which uses reinforcement learning models to understand the factors behind individual variability and clinical disorders (16); and behavioural experiments with human subjects which can be used to underpin computational cognitive models (17).

It is increasingly important that designers of AI systems can provide assurances regarding properties like safety, trustworthiness, ethics and fairness. Rigorous computer-based reasoning about the agents' models and algorithms is one route to help address these concerns (18). In order to build AI systems that people actually find useful, we need to go beyond purely technological solutions. Design Informatics (19) connects AI specialists with designers, cognitive scientists, and human-computer interaction experts to create fair, safe, user-centric AI applications.

Researchers in the School develop applications of AI-based systems in a wide variety of areas. In addition to those mentioned above, these include: modelling care pathways in healthcare (20); quantitative modelling and real-time data analysis techniques for finance and business (21); and using transfer learning to study cancer evolution (22).

TODO: ADD MORE EXAMPLE AREAS INCLUDING INDY, COMMERCE, SCIENCE? SAY MORE ABOUT COMPANY LINKS?

1. [Centre for Speech Technology Research](#)
2. [Prof Bob Fisher, VICO research group, Dr Laura Sevilla, Machine Intelligence Group](#)
3. [Mathematical Reasoning Group](#)
4. [Machine Learning Group](#)
5. [Dr Michael Mistry, Dr Zhibin Li](#)
6. [Statistical Machine Learning and Motor Control Group, Dr Steve Tonneau](#)
7. [Computer Graphics and Visualization Research Group](#)
8. [Robust Autonomy and Decisions Group](#)
9. [The Cohort research group](#)
10. [Natural Language Processing Group](#)
11. [Statistical Machine Translation Group](#)
12. [Extreme Summarization](#)
13. [Dr Walid Magdy](#)
14. [Computational Neuroscience and Neuroinformatics](#)
15. [Insect Robotics Group](#)

16. [Dr Peggy Series](#)
17. [Joint Eyetracking Lab](#)
18. [Dr Vaishak Belle, Mathematical Reasoning Group](#)
19. [Design Informatics](#)
20. [Workflow FM](#)
21. [Dr Tiejun Ma](#)
22. [Prof Guido Sanguinetti](#)