

TEACHING COMPUTERS TO UNDERSTAND OUR LANGUAGE

ARTIFICIAL INTELLIGENCE HAS ADVANCED IN LEAPS AND BOUNDS IN RECENT YEARS AND HAS BECOME AN INCREASINGLY INTEGRATED PART OF OUR LIVES. HOWEVER, IT STILL LACKS THE CAPACITY TO TRULY UNDERSTAND HUMAN LANGUAGE; SUBTEXT, CONTEXT CLUES AND EMOTIONAL EMPHASIS REMAIN DIFFICULT TO INTEGRATE. PROFESSOR YULAN HE, AT THE UNIVERSITY OF WARWICK IN THE UK, IS WORKING ON NOVEL METHODS TO CROSS THIS BRIDGE

TALK LIKE A COMPUTER SCIENTIST

ARTIFICIAL INTELLIGENCE (AI) – computer systems able to perform tasks that typically require human intelligence

COGNITIVE PSYCHOLOGY – the study of mental processes, such as language use and memory

CONTEXT – the circumstances surrounding an event

EVENT TRIPLE – the three constituent parts of an event: subject, predicate and object

NATURAL LANGUAGE PROCESSING – a branch of AI that aims for computers to be able to fully understand the meaning of human language

SEMANTIC – related to the meaning of language

SITUATION MODEL – a representation of an event from a cognitive perspective

‘Are you up to anything this weekend?’ This sentence provides a simple example of how different contexts can alter a sentence’s intended meaning. If you are asked this by a colleague or teacher, they are likely just making small talk. However, if you are asked this by a friend, they might be seeking to find out whether you have any free time to socialise with them. As social animals, we naturally understand and interpret subtle differences like this, which occur during hundreds of our interactions every day. But, this concept of understanding the subtleties of language is very difficult to instil in computers.

Professor Yulan He is a computer scientist at the University of Warwick, working in the field of natural language processing. She is training artificial intelligence (AI) to overcome this challenge through sophisticated reading comprehension exercises. “Language understanding is extremely difficult for computers due to the variety, ambiguity, subtlety and expressiveness of human languages,” she says. She is drawing information from a multitude of sources to provide a knowledge base that computer models can ‘learn’ from, and contribute to, to build the models’ understanding of how we use language.


READING COMPREHENSION

“When we read a passage of text, we usually have an understanding of the meaning of words in the context of the events described, and can then refer back to the ‘theme’ of the passage as we read onwards,” says Yulan. “We can ‘read between the lines’ to establish a deep semantic understanding of text.” Our reading comprehension depends on us building an event structure that represents what is happening in the text, referred to as a situation model in cognitive psychology.

Consider the following pair of sentences: ‘I went to a coffee shop. I had a flat white.’ Even though it’s not explicitly stated, we safely assume the narrator bought the flat white from the coffee shop, given that fits best within the context. We build an event structure of the narrator entering the shop, buying a coffee and drinking it. We also use our background knowledge to know that ‘flat white’ is a type of coffee, despite both words being adjectives that can be used for many other purposes. Yulan is trying to find ways for computers to build similar event structures, using context clues and background knowledge to help them ‘infer’ exactly what is going on in a situation.

FLUENCY IN COMMON SENSE

“A common way of teaching computers to



PROFESSOR YULAN HE

Department of Computer Science,
University of Warwick, UK

.....

FIELD OF RESEARCH

Natural Language Processing

.....

RESEARCH PROJECT

Building artificial intelligence models that can understand written language, by incorporating background knowledge, context and associated emotions

.....

FUNDER

Engineering and Physical Sciences
Research Council (EPSRC)

Yulan He is supported by a Turing AI Fellowship funded by the EPSRC (grant no. EP/V020579/1).

understand text is to first train a language model on large bodies of text, to capture word correlation patterns,” says Yulan. “These models can then be fine-tuned depending on the task they are needed for, such as answering questions.” This is the way that search engines or voice-activated virtual assistants work, for instance, but there are obvious limitations. “Even a small change in an input, such as paraphrasing a question, can decrease the model’s performance,” says Yulan. “They also struggle to automatically acquire common-sense or background knowledge to help their understanding.” For example, think about the times you have used voice-activated assistants to check a fact or play some music. Chances are, there have been times when the assistant has not understood your command or has not picked up on the context.

“We have made some attempts on building in common-sense knowledge for emotion analysis in text,” says Yulan. She gives the example of ‘Boxing Day’, a phrase which, when alone, generally has positive connotations given it is a public holiday. However, if someone texted you saying, ‘I need to work on Boxing Day’, you will assume the phrase was said with little enthusiasm and it is now associated with negative connotations. “We have developed a novel approach which adds an additional layer to existing language

models to capture topic transition patterns in human conversations,” says Yulan. “Both topics and common-sense knowledge are stored in a knowledge base that the model uses, or adds to, when attempting to detect emotion in written conversations.” This layer helps the model learn associations between particular phrases, the surrounding context and commonly related emotions.

PICTURES AND WORDS

Yulan is building a framework that extracts events from text, which follow a simple structure called an event triple: a subject (the thing performing the action), the predicate (the action that is taking place), and the object (the thing the action happens to). “Event representational learning aims to map these triples within a multi-dimensional space, where triples with similar semantic meanings are found in nearby locations,” says Yulan. “However, we argue that textual descriptions alone can only go so far.” The phrase ‘a picture is worth a thousand words’ has some truth to it, which is why Yulan’s team are using images as a less abstract way of teaching these models.

Think about the coffee shop example: ‘I went to a coffee shop. I had a flat white.’ An image of this event could tell you a lot more: how busy the

coffee shop is, the weather outside, the size of the coffee, the mood of the narrator, to name just a few. “We’ve developed a framework to learn event representations based on both text and images simultaneously,” says Yulan. “Experimental results show that our proposed framework outperforms many simpler frameworks.”

BENEFITS TO SOCIETY

“Since spoken and written communication play a central part in our lives, frameworks like ours could have profound impacts for society,” says Yulan. “It has applications for intelligent virtual assistants, drug discovery, or answering complex questions about financial or legal matters.”

Yulan’s team is already working on practical applications, including collaborating with pharmaceutical giant AstraZeneca to detect adverse drug events from biomedical literature. This use of AI will help extract important information about undesirable side-effects that otherwise might never be detected.

With AI becoming increasingly integrated in our lives and societies, we need computer scientists like Yulan to ensure these systems understand us.

ABOUT ARTIFICIAL INTELLIGENCE

Yulan tells us more about the joys and challenges of developing AI:

WHAT DO YOU FIND MOST REWARDING ABOUT YOUR WORK WITH AI?

AI has the potential to transform a wide range of sectors, from health and social care to financial services, from citizen services to manufacturing. The technologies and algorithms developed for AI can be used to make people's lives better and being part of this research process is very rewarding for me.

WHAT PRACTICAL CHALLENGES ARE ASSOCIATED WITH DEVELOPING AI?

Existing AI technologies are far beyond human capabilities in terms of computational intelligence. They have almost reached similar performance levels to us in sense intelligence, such as speech or image recognition. Yet,

they are still far away in terms of cognitive intelligence: understanding language, and knowledge abstraction and inference.

While we can learn from past experiences to perform abstract reasoning, it is difficult to train an AI model to emulate this. This is due to the challenges in collecting training data, and the near-infinite combinations of contexts that define real-world problems.

WHAT ETHICAL CHALLENGES ARE ASSOCIATED WITH DEVELOPING AI?

Different social groups behave differently, and if the systems we develop are based on data that exclude certain groups, this will cause issues. There is also the possibility of AI developers unconsciously introducing our own biases into the algorithms, so we make sure to take steps to mitigate this. Additionally, these large-scale models require significant resources to train, which are only accessible to a handful of giant tech

companies. This could lead to a monopoly for applications of AI, which could make the benefits it brings inaccessible for many.

WHAT ATTRIBUTES DO SUCCESSFUL AI COMPUTER SCIENTISTS HAVE?

Curiosity, passion, deep thinking, and an ability to set big goals are the most important. A strong mathematical background and good programming skills are also essential for developing models and algorithms in AI.

WHERE IS AI FOUND IN OUR DAILY LIVES?

AI can be found almost everywhere in our life. Google's search engine, and virtual assistants such as Apple's Siri and Amazon's Alexa, are all examples of AI. Face recognition software uses AI, many cars and homes now have AI-integrated sensors, and AI provides personalised recommender services in Netflix and YouTube.

EXPLORE A CAREER IN ARTIFICIAL INTELLIGENCE

- Yulan names Google (www.ai.google), Open AI (www.openai.com) and DeepMind (deepmind.com/careers) as major AI researchers. All have plenty of information about the latest AI developments and careers in AI on their websites.
- Yulan recommends learning a coding language as soon as possible. She says that HuggingFace (www.huggingface.co) provides an open-source library for many natural language processing technologies, as well as programming tutorials.

YULAN'S TOP TIPS

- 01** Build a strong mathematical foundation. This is important to help you understand the theories behind AI algorithms in your future studies.
- 02** Learn at least one programming language, such as Python. Then, practise using it to develop simple AI systems to solve problems that interest you.

PATHWAY FROM SCHOOL TO ARTIFICIAL INTELLIGENCE

At school, a strong mathematical background is essential, and computer science would also be advantageous if available. Studying psychology will lend useful insights into how humans process language.

Most universities offer undergraduate degrees in computer science or mathematics, which would provide the clearest path to a career in AI. Yulan recommends taking as many mathematics modules as possible at university. She particularly suggests linear algebra, calculus, statistics and probability, graphical models and convex optimisations. She says that subjects such as cognitive science, statistical modelling, machine learning, data science, natural language processing, computer vision, robotics, multi-agent systems, computational intelligence and human-computer interaction are also relevant.

HOW DID YULAN BECOME A COMPUTER SCIENTIST?

WHAT WERE YOUR INTERESTS WHEN YOU WERE YOUNGER?

I loved literature and won quite a few writing competitions while at secondary school. I was always fascinated by how language can influence our perception of the world and our thought processes. I was equally interested in mathematics and coding. I wrote my first computer program using BASIC at the age of nine.

HOW HAVE YOUR PREVIOUS STUDIES HELPED PREPARE YOU FOR YOUR CURRENT PROJECT?

My master's in computer engineering focused on developing algorithms for the retrieval

of scientific publications over the internet. I then pursued a PhD in spoken language understanding, under the supervision of Professor Steve Young, a world-renowned expert in spoken dialogue systems. Both these qualifications helped build my knowledge of statistical model learning, information retrieval and semantic understanding of languages, providing a solid foundation for my current project.

WHAT MOTIVATES YOU IN YOUR WORK?

The ultimate goal of AI is to create a thinking machine skilled in all aspects of intelligence. Achieving this ambitious goal requires a deep

understanding of how intelligence emerges. Language is crucial to general intelligence and natural language processing is a key area in AI.

HOW DO YOU SPEND YOUR FREE TIME?

I still love reading. I also like being active in the outdoors, through jogging and cycling.

(a): Food and Restaurant

A: Could I have some fish?

B: Certainly. And what vegetables would you like? 😊

A: Oh , spinach , I think.

A: I like drinking tea at teahouses. 😊

B: Oh, so do I.

A: Why don't we go for one now?

😊 B: Great. We can chat while enjoying a cup there.

A: Let's go! 😊

Utterances on particular topics carry certain emotions in dialogues. Yulan, and the Natural Language Processing Group at the University of Warwick, is training computers to recognise and understand these emotions conveyed within text.