

Building artificial intelligence for social good

Artificial intelligence is one of the most disruptive technologies nowadays and as such is considered both an opportunity and a threat to society. On the one hand, it could make us vulnerable as we trust machines to make more and more decisions that so far have been reserved for humans. On the other hand, it has the potential to revolutionise large-scale energy management, traffic control, the job market, and countless other areas of our economy. For this reason, **Dr Long Tran-Thanh**, a Lecturer at Southampton University, aims to harness AI to solve the most important societal challenges of our age.

As artificial intelligence becomes more and more ubiquitous, we need to learn how to live and interact with it. And of course, this goes both ways: AI systems also need to learn how to interact with humans, and take their preferences or choices into account. Dr Long Tran-Thanh calls this making AI more “human-aware” – something that so far has not been fully exploited. Often, AI researchers will design a beautifully efficient and rational digital system, only for it to fail when faced with the unpredictability of its human users. Dr Tran-Thanh wants to change this and develop AI platforms that can be used to tackle some of society's biggest problems. He wants AI to take into account that human users are not perfectly efficient nor perfectly rational, that we are creatures of habit and are often stubbornly stuck in our ways.

SMOOTHING OUT THE ENERGY PEAKS

One of the first attempts he made at making this vision a reality was a home energy management system. The main



issue with operating a nationwide electricity grid is that the demand for energy varies throughout the day: there are peaks and troughs all the time. For the supplier, this translates into keeping so-called “peak plants” up and running at all times just in case there is a sudden peak in energy consumption. In practice, such plants are rarely used but need to be manned and supplied with fuel nonetheless which leads to increased costs and greenhouse gas emissions.

One possible way of dealing with this is smoothing those peaks out by convincing people to change their energy consumption habits. To quantify the problem, a team of researchers at Massachusetts Institute of Technology fitted a few households with sensors monitoring the energy consumption of every home appliance. After that, numerous other teams tried to unleash AI

– and come up with the most efficient energy consumption plan within such derived constraints. However, another issue quickly came to light. People, bombarded every day with hundreds of incoming emails, phone calls and social media notifications didn't want to be bothered by yet another set of queries eating up their time and attention.

INNOVATIVE AI DESIGN

This is where Dr Tran-Thanh's innovation shone out. His AI was programmed to do as much as possible on its own, and query the user only when necessary. To achieve this, Dr Tran-Thanh's team had their algorithm take into account the “bother cost”, which reflected the increasing annoyance caused by asking questions. The result was a system designed to take preferences into account and maximise energy efficiency while minimising the “bother cost” at the same

Dealing with traffic management in a big city like Singapore or London is crucial – make one mistake and you cause instant mayhem

on the gathered data to find the most efficient ways of reducing energy costs in homes. But it quickly transpired that algorithms which had been super-efficient in simulations, proved nearly useless in the real world. The reason was simple – people rejected their recommendations as incompatible with their daily routines. If you have been doing your laundry on Saturday afternoon all your life, it will take more than a machine's recommendation to reschedule for 3am on Tuesday.

Dr Tran-Thanh and his team then thought about a system that would simply ask users about their preferences – for example, “Would you mind turning your kettle on an hour later on Wednesdays?”

time. And it worked. Overall, it has managed to reduce the participating households' energy consumption by 35%. In the not-so-distant future, Dr Tran-Thanh's AI will be implemented in Schoonschip, an experimental floating neighbourhood in Amsterdam.

There were a few other projects where Dr Tran-Thanh's team sought to apply their human-aware AI idea. One of them was built to cost-efficiently assign tasks to workers, each with his or her own hourly rate and skill level, hired through expert crowdsourcing platforms. Another dealt with applying AI to fight drug trafficking and other situations where networks needed to defend themselves



against adapting and malicious adversaries. This combined experience allowed Dr Tran-Thanh and his colleagues to take on one of the biggest challenges facing developed countries – traffic management.

INTELLIGENT TRAFFIC CONTROL

Thousands of sensors measuring traffic are fitted along streets and avenues to provide real-time data about how many cars drive past them and how quickly. This data is then transferred to data centres and processed by AI algorithms making decisions about managing the traffic lights all over the area. Solutions much like the one used in Dr Tran-Thanh's home energy management system will keep human operators in the loop without overloading them with information and systems resembling those Dr Tran-Thanh designed to tackle drug trafficking will protect the network from cyber attacks.

It turned out algorithms that had been super-efficient in simulations proved nearly useless in real world applications

What makes this project stand out is the huge consequences of a potential failure. Bad design in a home energy management system, or even letting a drug trafficker escape, causes a relatively small-scale, local problem. But dealing with traffic management in a big city like Singapore or London is different – make one mistake and you cause instant mayhem. That's why there are numerous teams working on different aspects of the AI system being built to take on this challenge. Professor Nicholas R. Jennings, a vice-provost at Imperial College London leads a team responsible for the system's cyber security and Dr Bo An, Nanyang

Assistant Professor at NTU Singapore, is working on the game theory approach to eventual attacks. Dr Tran-Thanh's team is responsible for making algorithms which can recognise patterns in the vast amounts of traffic data generated by the sensors. Making an enterprise this huge a reality will take years. But with people like Professor Jennings, Dr Bo An, and Dr Tran-Thanh in charge, we can rest assured we will be well taken care of.



Behind the Bench

Dr Long Tran-Thanh

E: ltt08r@ecs.soton.ac.uk T: +1 442380593715 W: <http://users.ecs.soton.ac.uk/ltt08r/>

Research Objectives

Professor Tran-Thanh develops AI algorithms aimed at tackling societal problems and his work focuses on increasing security for AI.

Funding

EPSRC (EP/N02026X/1)

Collaborators

- Professor Nicholas R. Jennings (vice-provost of Imperial College London)
- Professor Alex Rogers (Oxford University)
- Professor Milind Tambe (University of Southern California Center for Artificial Intelligence in Society)
- Dr Sebastian Stein and Dr Sarvapali D. Ramchurn (University of Southampton)

- Dr Henry Ngoc Cuong Truong (SenseEye) and Dr Tim Baarslag (CWI Amsterdam)
- Dr Bo An (Nanyang Assistant Professor at NTU Singapore)
- Dr Vinh Thong Ta (University of Central Lancashire)
- Dr Matteo Venanzi, Dr Valentin Robu, Dr The Anh Han, Dr Avi Rosenfeld, Dr Trung Dong Huynh, Dr Michael Kaisers, Dr Marcin Waniek, Haifeng Xu

Bio

Long Tran-Thanh is a Lecturer (Assistant Professor equivalent) at the University of Southampton, UK. He obtained his PhD in Computer Science in 2012 at the same university, under the supervision of Nick Jennings and

Alex Rogers. He has been conducting active research in a number of key areas of Artificial Intelligence. His results have been published at top AI conferences and journals and he has received a number of prestigious awards.

Contact

Dr Long Tran-Thanh
University of Southampton
University Road
Southampton
SO17 1BJ
United Kingdom

Q&A

How much more efficient will traffic control get once the system you're working on is implemented?

Our system makes the smart traffic control more secure against malicious and adversarial behaviour, while keeping the traffic control efficiency at the same level (i.e., we protect the traffic system without compromising on its quality).

What are the possible consequences of a successful cyber attack carried out against the traffic management system?

A possible consequence is that terrorists can take over the traffic control system and create chaos, significantly increasing the casualty tally in case of a terrorist attack. Another example is a heist, in which the robbers, by being able to control the traffic system, will block the routes of the police while freeing up their escape

route (see, for example, the Ocean's 11 movie).

How is your AI for home energy management system going to work at Schoonschip?

We are planning to continue our collaboration with Kaiser's team at CWI to run a set of experiments at Schoonschip to collect data about different user behaviours. Once we have this data, we will be able to model these behaviours and fine-tune our energy management system against these models. Once this is done we aim to deploy the system in these real homes.

How do algorithms like yours help fight drug trafficking and what are the results so far?

This is still in a theoretical research stage. However, in the (near) future, we aim to introduce our results to the appropriate authorities to initiate a collaboration with them in order to deploy our algorithm in real life.

To what extent do you want to incorporate affective computing in your "human-aware" AIs and do you think they should take into account users' emotional states, such as asking questions when they are less stressed or angry?

Affective computing is definitely an important tool in order to achieve efficient human-aware AI. In the future, I am planning to incorporate it into my research. I am working closely with Dr Klaus-Peter Zautner (bio-inspired robotics) and Dr Hedwig Eisenbarth (psychology), both from Southampton, to combine emotion detection, robotics, and sequential machine learning, to make robots more aware of human annoyance/stress level, so they will know when it is appropriate to approach their human trainers for further guidance and instructions.