How teams of computers and humans can fight disasters

September 24 2015, by Oxford Science Blog

Credit: AI-generated image (disclaimer)

Over the past five years, researchers from Oxford University have been working on a collaborative project called ORCHID to develop new ways for humans and computers to work together.

This week, the team from Oxford joined their academic collaborators from the University of Southampton and University of Nottingham at the Royal Academy of Engineering yesterday to showcase their work. We spoke to Dr Steven Reece, a Senior Research Fellow at the University's Pattern Analysis and Machine Learning Research Group, to find out how the Oxford team has been using its research to help disaster response teams.

OxSciBlog: ORCHID attempts to integrate humans—and all of their foibles—with computers, so that they can work together as so-called human-agent collectives. Why is it important?

Steven Reece: Ninety percent of all recorded data that exists in the world has been generated in the past two years. This data is vast and mostly unstructured, made up of all kinds of text documents, photographs and videos. The problem is that humans and computers look at this data very differently. Humans are very good at understanding unstructured data—they can interpret the meaning of text and understand events depicted in a photograph better than any software, for example—but they can't work through that much of it. Computers, on the other hand, are better than humans at processing and spotting patterns in vast amounts of data very quickly. Human-agent collectives (HACs) take the best of both worlds, creating flexible teams of computers and humans to interpret large, unstructured data sets.

OSB: How do these HACs work?

SR: Traditionally, humans tell computers what to do; HACs turn that relationship on its head and allow computers to take control occasionally and request information from humans. Of course, humans and computers have their foibles: they can be unreliable, malicious, selfish and, in the case of humans, they can even get bored. But it was the goal of ORCHID to figure out how to mitigate these foibles: how to incentivise humans to contribute to the HAC, track performance, maintain the best teams and record the sources of information and decisions that are made.

OSB: Can you describe the kind real-world problems you've been applying that thinking to?

SR: As just one example, crisis responders need to know the extent of a natural disaster, what aid is required and where they need to get to as quickly as possible. This is what's known as 'situation awareness'. With the proliferation of mass media, a lot of data is now generated from the disaster zone via photographs, tweets, news reports and the like. With the addition of first responder reports and satellite images of the disaster area, there is a vast amount of relevant unstructured data available for situation awareness. A crisis response team will be overwhelmed by this data deluge—perhaps made even worse by reports written in languages they don't understand. But the data is also hard to interpret by computers alone, as it's difficult to find meaningful patterns in such a large amount of unstructured data, let alone understand the complex human problems that described within it.

OSB: How can you use HACs to help?

SR: Firstly, we can farm out satellite images and text to the 'crowd'. People want to help and they will happily use their skills to interpret a small number of text samples or satellite images. Computers can then build a model connecting features in the data to the interpretations supplied by the crowd. The computer can then use this model to trawl through the rest of the data and 'interpret' what it sees using these features. The computer decides what data to farm out to the crowd and who should be recruited from the crowd based on their reliability; The individuals in the crowd can decide if they want to take part and what tasks they are prepared to do. The computer aggregates the crowd responses intelligently and, in so doing, determines their individual reliabilities automatically. So we can use combinations of humans and computers to successfully aggregate and interpret vast amounts of unstructured data. This is just one example of where HACs can be used in disaster response—another is the coordination of a vast fleet of UAVs visually mapping aid requirements across the disaster area.

OSB: Have you been able to try any of these ideas in the real world?

SR: We've implemented the first approach I just explained, actually. It's in a system called 'CrowdScanner', and we used it for real immediately after the first major Nepal earthquake in April of this year. We used the crowd to locate settlements from satellite images, identified settlements that were not mapped on open sources such as OpenStreetMap, and our Search and Rescue partners deployed teams to reconnoiter these settlements.

OSB: Was it successful? Did you run into any pitfalls when marrying up computers and humans?

SR: First of all, the good news! It is not difficult to find a competent crowd to help out in a disaster situation. We are, however, finding it really difficult to build systems that are relevant to the disaster response community—mainly because we are trying to guess what their critical information requirements might be.

We were able to respond to the Nepal crisis because we just happened to have a working platform that we could adapt to the Nepal situation along with a satellite data source, and Rescue Global, who we were working with, were able to marry this data with their requirements of water filter placement and life detector placement. We were able to respond in a timely manner as a result.

In general, though, we're in the dark as to the generic problems faced by the crisis response community and specifically where we can help. Although we've attended field exercises with various disaster response organisations we need to sit down with them. That way, we'll be able to abstract information about where they work and their requirements to the level where we can start designing generic situation awareness algorithms to help them.

OSB: So what's next for the disaster response work?

SR: Imagine a service where people could post their resources, such as the availability of an aircraft or their plan to visit a location in the disaster area. This resource could be married with a crisis response team who want to use it to achieve their own goals, and machine learning can be used to link the responders' requirements to the people with the resources using crowd interpretations of the resource providers' offers. Another idea is that we could try to determine the probability of life in a collapsed building after an earthquake, to help responders prioritise their search. Crowdsourced interpretations of drone or satellite footage could be used to identify salient features and machine learning can then determine the probability of life.

There are many ways machine learning can be used in disaster response. The key now is to sit down with crisis responders and develop relevant data processing algorithms that will actually save lives.

Provided by Oxford University