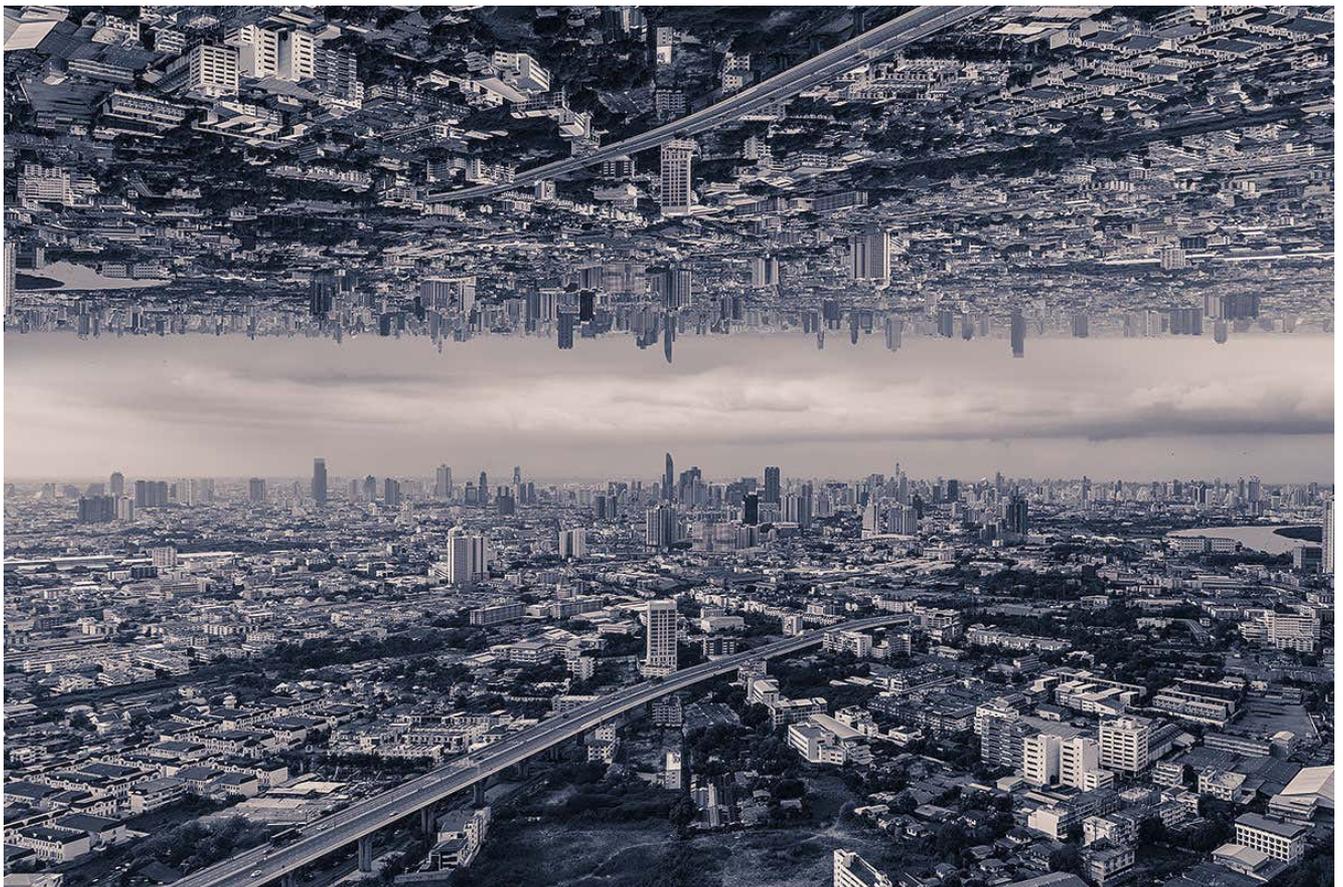


Predicting the future is now possible with powerful new AI simulations

If you thought Cambridge Analytica had scary tech, wait until you see this. A new form of AI modelling promises accurate simulation of the behaviour of entire cities, countries and one day perhaps, the world

[Technology](#) 2 October 2019



City simulations can imagine almost anything

Getty Images

OCTOBER 2020. The US presidential election campaign is in its final days. Donald Trump is behind in the polls and the pundits are predicting a win for his Democrat challenger, former vice president Joe Biden. But Trump is unruffled. He boasts that he will win again. [Bigly](#).

With two weeks to go, his campaign unleashes an offensive in the crucial swing states: adverts, Facebook posts, WhatsApp groups and tweets. They warn of violent crime and civil unrest driven by immigrants and gangs, playing up Trump's endorsement by evangelicals and smearing Biden as a closet atheist. The initiative works and Trump snatches another unlikely victory.

You probably think you have heard it all before. It is a replay of 2016, [when consulting firm Cambridge Analytica](#) used targeted messaging to apparently influence the outcome of the US election, right? Wrong. In this scenario, there is a new, even more persuasive technology: multi-agent artificial intelligence (MAAI). This tech allows predictions to be made with extraordinary accuracy by testing them in highly detailed simulations that amount to entire artificial societies. If, for example, a campaign team wants to decide how and to whom to pitch their messages – how to fight an election – it can do so, multiple times, inside a computer simulation.

The idea that the Trump campaign is planning to use MAAI is pure speculation. But in terms of technology, there is nothing to stop it. MAAIs are already being used to build digital societies that simulate real ones with uncanny accuracy. That

allows people to perform radical social experiments. Want to know what will happen if 20,000 Syrian refugees arrive in a city in western Europe? Build an artificial society and watch. Want to know how to make the integration of those immigrants peaceful? Build an artificial society, try things out and see what works. Want to stoke anti-immigrant hostility or design a disinformation campaign to win an election...?

In simple terms, an artificial society is just a computer model similar to those that have been used for decades to understand complex dynamic systems, such as the weather. The first were built by physicists and chemists in the 1960s, but as the models increased in complexity, they were embraced by biologists and, in the past decade, social scientists.

One of the most useful techniques is [agent-based modelling](#), which uses strings of computer code to represent agents, such as drivers navigating a route or companies competing in an economy. The agents are programmed to interact with one another and their virtual environment and change their behaviour accordingly. These models are useful for understanding how complex systems work, predicting how they will evolve and testing what happens if you intervene.

In 2014, for example, an [Ebola epidemic broke out in West Africa](#). As cases mounted, the US Defense Threat Reduction Agency asked computer modellers to predict how the epidemic would progress. Over seven months, they built an agent-based model that used real-world data on case numbers, infection rates, healthcare systems, population distribution,

demographics, economic and social interactions, travel patterns and even cultural behaviour such as funeral rites. It predicted that, left unchecked, the virus would infect 1.4 million people.

It was also used to test interventions to halt the spread. Medical teams were sent where the model said they would be most effective and people in affected areas were advised to adopt quarantine measures and safe burial practices. In the end, infections were restricted to 28,000 people. We can't know for sure that the model worked, that the interventions led to a lower number than it predicted, but this case is frequently cited as a successful use of agent-based modelling.

The model human

Even here, the agents are quite basic. Models are computationally expensive and modellers have to use their resources sparingly. Agents are thus endowed with the bare minimum of simple attributes – being more or less open to health messages, for example – and a small repertoire of behavioural responses, such as fleeing or staying put. Such models can produce surprisingly complex behaviour, but you would hesitate to call them an artificial society.

In the past couple of years, however, the game has changed, driven by a dramatic increase in the availability of four key raw materials: computing power, data, scientific understanding of human behaviour and, most crucially, artificial intelligence (AI).

“It has always been one of the ambitions of agent-based modelling to have intelligent agents,” says Nigel Gilbert, head of the Centre for Research in Social Simulation at the University of Surrey, UK. With the arrival of MAAI, that ambition has been fulfilled.

With AI, the models suddenly become more realistic. “One of the things that has changed is an acceptance that you really can model humans,” says [F. LeRon Shults](#), director of the Center for Modeling Social Systems at the University of Agder in Norway. “Our agents are cognitively complex. They are simulated people with genders, ages and personalities. They can get married, have children, get divorced. They can get a job or get fired, they can join groups, they can die. They can have religious beliefs. They’re social in the way humans are. They interact with each other in social networks. They learn from each other, react to each other and to the environment as a whole.”

The increase in computing power also means that the number of agents in a model can be vastly increased, from a few thousand to tens of millions. “We can model a city the size of London,” says Saikou Diallo at the [Virginia Modeling, Analysis and Simulation Center](#) at Old Dominion University in Virginia.

Shults says the next milestone is 320 million, the size of the US population, and from there 1.4 billion to model China. Ultimately, the goal is the whole world.

The way to harmony

The result is a revolution in agent-based modelling. “We can replicate how real societies work to explore real-world questions,” says Diallo. If this sounds something like *The Sims*, that’s because it is. But whereas *The Sims* is a game, virtual societies powered by millions of artificial intelligence-driven agents are deadly serious.

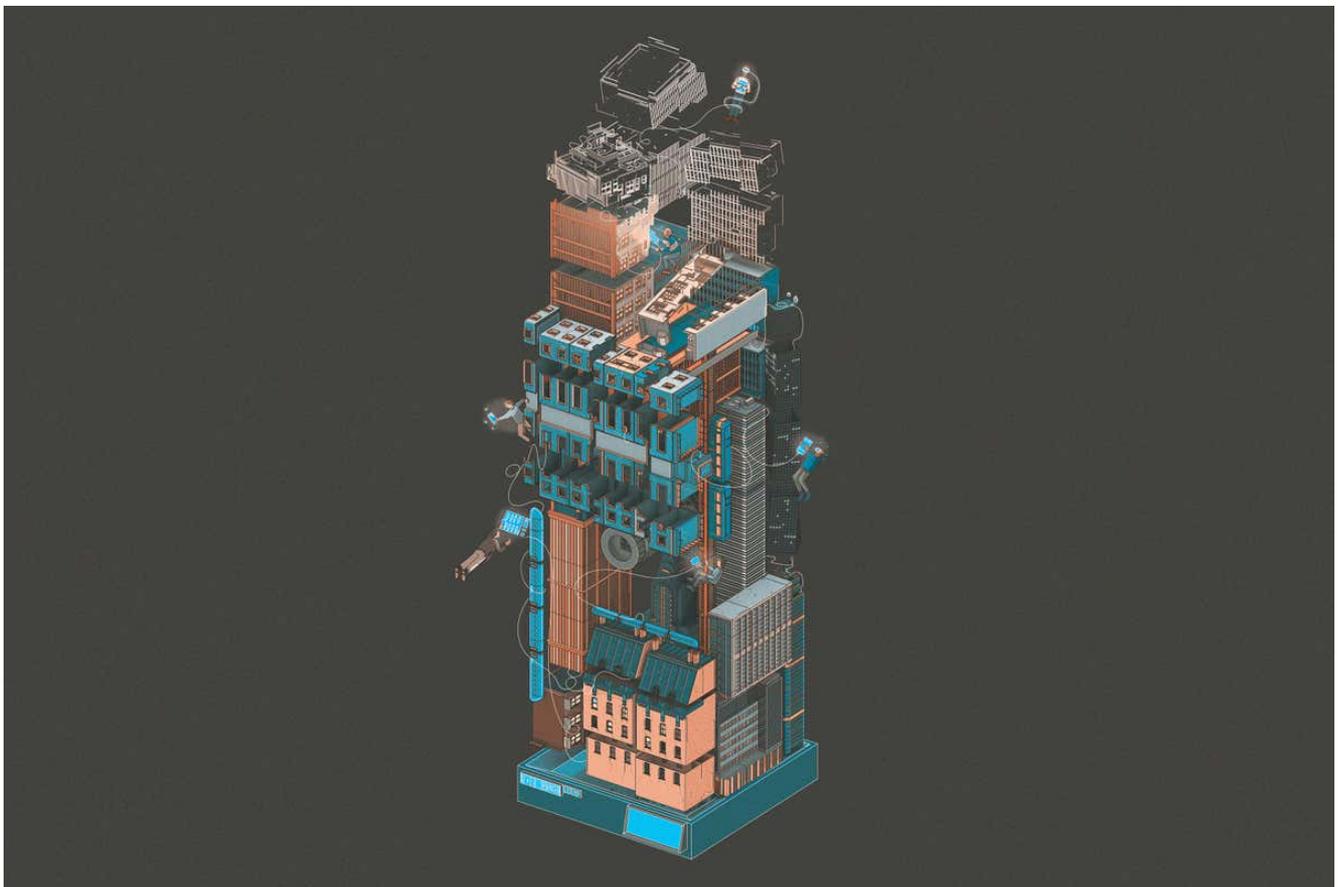
In the past eight years, [a million Syrian refugees have fled to Europe](#), and some 20,000 of them settled in Norway. The influx of mostly Muslim immigrants into a relatively ethnically homogeneous, secular country with Christian roots has stirred up tensions. Harmonious integration is an urgent issue. The third largest political party in Norway is the right-wing, anti-immigration Progress Party. “You want to have a society that is not full of radicalisation,” says Shults.

The old-fashioned way to achieve this is to design and implement policies that you believe will work. “Everybody sits around the table and argues about the right policy: should we invest a lot of money on making immigrants feel economically secure? Should we invest in teaching them the language or culture? Should we spend it on education? Should we spend it on places for the young men to play soccer with Norwegians? Everyone has their idea,” says Shults.

The stakes are high: if you make the wrong call, the outcome could be catastrophic and irreversible. “If 10 years later you have economic collapse and terrorism, you can’t hit the reset button,” says Shults.

But with a computer simulation, you can try out all sorts of interventions. If a policy backfires in the model, you can hit reset.

That is the goal of a model being developed by Diallo, Shults and others, which simulates a typical Norwegian city with a sudden influx of refugees. It is a relatively small model with just 50,000 agents but will run for three generations to test the long-term outcomes of various policies. Models such as this take between hours and days to complete a run, depending on the number of parameters involved. “It allows you to do experiments that are impossible in the real world,” says Shults.



Because of this power, MAAI technology has the potential to tackle the world’s most complex problems. This month, Shults and his colleagues are sitting down with experts on climate, energy and conflict to start modelling a [refugee crisis triggered](#)

[by climate change](#). “Most experts think that climate was a big factor in the Syrian refugee crisis,” says Shults. “A million people flowed into Europe. As sea levels rise over the next 20 to 30 years, we’re talking at least 100 million. Where are they going to go? There will be massive human suffering. Our goal is to come up with policy initiatives to change behaviours and avoid conflict.”

Other modellers are working on preventing ethnic conflict and breaking up protection rackets and sex trafficking rings. Shults also sees applications in politics: “I’d like to understand what is driving populism – under what conditions do you get Brexit, or Le Pen?”

Of course, it isn’t possible to capture the full complexity of human behaviour and social interactions. “We still don’t really know how people make decisions, which is a major weakness,” says Bruce Edmonds, director of the Centre for Policy Modelling at Manchester Metropolitan University, UK. “In most cases, there are some bits of the model that are empirically validated but some bits that are guesses.” MAAI is still so new that we don’t yet know how accurate it will be. Shults says the outputs of the model are still valid – if you get your inputs right: “One of the common phrases you hear is ‘all models are wrong, but some are useful’.”

The first step is to decide what to model, then bring in the best expertise available. For the refugee model, for example, Shults and his colleagues will call on social scientists who have theoretical models and empirical data on religious conflict and

social integration.

Stage one is to “formalise the theory”, which means nailing down exactly how the theoretical models apply to people in the real world and describing it mathematically. At this point, the modellers start to build agents.

Every conceivable social interaction can be modelled: between family, friends, bosses, colleagues, subordinates and religious leaders, and from economic dealings to social media engagement. Through these interactions, the agents learn, altering their future behaviour. Summed across the whole simulation, they can alter the trajectory of the society.

Once the simulation is built, it has to be validated. That means plugging in data from the real world and seeing whether it recapitulates what actually happened, and if not, tweaking it accordingly. In the refugee assimilation model, Shults and his team will use data from social surveys carried out by the Norwegian government, plus a decade of data on assimilation in London and Berlin.

Unsurprisingly, it isn't a trivial undertaking, taking about a year. But once validated, you are ready to play God. That might just mean setting initial conditions and watching how things pan out. It might mean testing an intervention that you think might help – say, pumping resources into a [deradicalisation programme](#). Or it might mean asking the simulation to find a pathway to a desirable future state.

Don't be evil

The power of the technology is that you can do all of these things at once. “The simulation is running hundreds, thousands, millions of parameter sweeps to see under what conditions agents are going to move, change and do different things,” says Shults.

The power brings great responsibility. “The ethical question bothers me,” says Shults. “Could this technology be used for evil?” We already know the answer. Shults’s team modelled a society with a majority religious group in conflict with a minority one. They found that such societies easily spiral into deadly violence. When they ran the simulation to find the most efficient way to restore peace, the answer that popped out was deeply troubling: genocide.

There is also a very real fear of the technology being exploited, as many feel happened with Cambridge Analytica. “They used AI to trick people into believing something so they would vote a certain way,” says Shults. He and his colleagues fear that something even more manipulative could be done with MAAs. The US election scenario is hypothetical, but plausible. Using simulation technology, theoretical insights could be weaponised for electoral gain. “Yes, it can be used for doing bad,” says Diallo. “It could be used to psychologically target people or groups and work out how to influence them.”

Or worse. A group at the Center for Mind and Culture in Boston has created an MAAs to test ways to break up child sex

trafficking rings. Team leader Wesley Wildman points out that the traffickers could hire someone to build a rival simulation to disrupt the disrupters in a technological arms race. “It could already be happening. As far as I know, we’re ahead of them, but they will catch up,” he says.

The Society for Modeling and Simulation International, of which Diallo is president, takes these threats so seriously that it is drawing up a set of ethical guidelines for modellers. They are forbidden from working with criminals, but what if a politically motivated group asks for help? “At that point, you’re face to face with a conundrum,” says Wildman. He says that Cambridge Analytica didn’t do anything wrong, except for not telling people what they were up to. “If that is the way political campaigns are going to be run, fine, but be clear about it.” The only ethical requirement that could be placed on modellers working for political campaigns is transparency. Does that make you feel secure?

The complexity and obscurity of MAAI mean it is unlikely that anyone is manipulating you – yet. Outside the small community of modellers, the existence of MAAI remains largely unknown. “I think it is possible that bad actors are using it,” says Wildman, “but I don’t think they’d be very far along.”

Gilbert says some policy analysts are becoming aware of it, but most politicians are in the dark. According to Edmonds, Dominic Cummings, special advisor to UK prime minister Boris Johnson, is aware and interested.

It is only a matter of time. You know about it now, and maybe Trump does too. For Wildman, the genie will soon be out of the bottle: “This is coming, whether we’re ready for it or not.”