New artificial intelligence system automatically evolves to evade internet censorship

November 13 2019

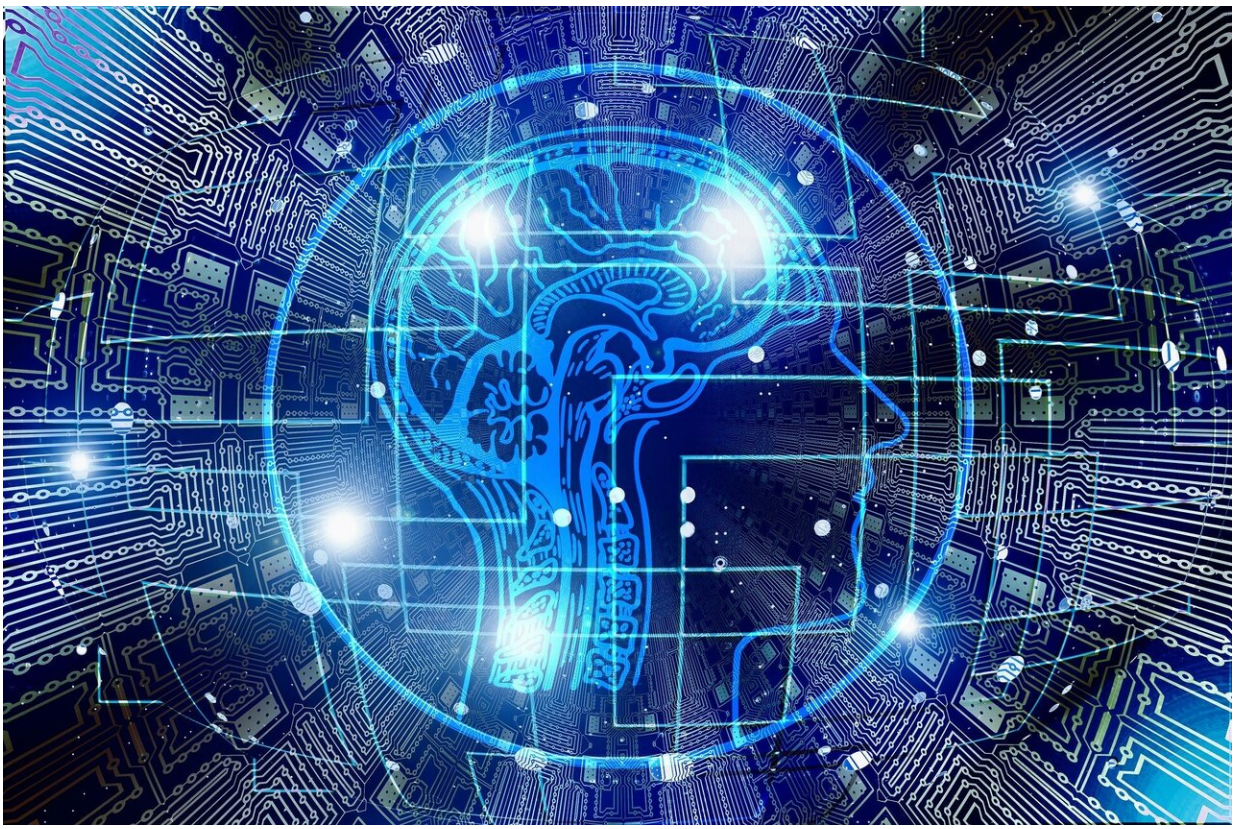

Credit: CC0 Public Domain

Internet censorship by authoritarian governments prohibits free and open access to information for millions of people around the world. Attempts

to evade such censorship have turned into a continually escalating race to keep up with ever-changing, increasingly sophisticated internet censorship. Censoring regimes have had the advantage in that race, because researchers must manually search for ways to circumvent censorship, a process that takes considerable time.

New work led by University of Maryland computer scientists could shift the balance of the censorship race. The researchers developed a tool called Geneva (short for Genetic Evasion), which automatically learns how to circumvent censorship. Tested in China, India and Kazakhstan, Geneva found dozens of ways to circumvent censorship by exploiting gaps in censors' logic and finding bugs that the researchers say would have been virtually impossible for humans to find manually.

The researchers will introduce Geneva during a peer-reviewed talk at the Association for Computing Machinery's 26th Conference on Computer and Communications Security in London on November 14, 2019.

"With Geneva, we are, for the first time, at a major advantage in the censorship arms race," said Dave Levin, an assistant professor of computer science at UMD and senior author of the paper. "Geneva represents the first step toward a whole new arms race in which artificial intelligence systems of censors and evaders compete with one another. Ultimately, winning this race means bringing free speech and open communication to millions of users around the world who currently don't have them."

All information on the internet is broken into data packets by the sender's computer and reassembled by the receiving computer. One prevalent form of internet censorship used by authoritarian regimes works by monitoring the data packets sent during an internet search. The censor blocks requests that either contain flagged keywords (such as "Tiananmen Square" in China) or prohibited domain names (such as

"Wikipedia" in many countries).

When Geneva is running on a computer that is sending out web requests through a censor, Geneva modifies how data is broken up and sent, so that the censor does not recognize forbidden content or is unable to censor the connection.

Known as a genetic algorithm, Geneva is a biologically inspired type of artificial intelligence that Levin and his team developed to work in the background as a user browses the web from a standard internet browser. Like biological systems, Geneva forms sets of instructions from genetic building blocks. But rather than using DNA as building blocks, Geneva uses small pieces of code. Individually, the bits of code do very little, but when composed into instructions, they can perform sophisticated evasion strategies for breaking up, arranging or sending data packets.

Geneva evolves its genetic code through successive attempts (or generations). With each generation, Geneva keeps the instructions that work best at evading censorship and kicks out the rest. Geneva mutates and crossbreeds its strategies by randomly removing instructions, adding new instructions, or combining successful instructions and testing the strategy again. Through this evolutionary process, Geneva is able to identify multiple evasion strategies very quickly.

"This completely inverts how researchers typically approach the problem of censorship," said Levin, who holds a joint appointment in the University of Maryland Institute for Advanced Computer Studies. "Ordinarily we identify how a censorship strategy works and then devise strategies to evade it. But now we let Geneva figure out how to evade the censor, and then we learn what censorship strategies are being used by seeing how Geneva defeated them."

The team tested Geneva in the laboratory against mock censors and in

the real world against real censors. In the lab, the researchers developed censors that functioned like those known from previous research to be deployed by autocratic regimes. Within days, Geneva identified virtually all the packet-manipulation strategies that had been discovered by previously published work.

To demonstrate that Geneva worked in the real world against undiscovered censorship strategies, the team ran Geneva on a computer in China with an unmodified Google Chrome browser installed. By deploying strategies identified by Geneva, the user was able to browse free of keyword censorship. The researchers also successfully evaded censorship in India, which blocks forbidden URLs, and Kazakhstan, which was eavesdropping on certain social media sites at the time. In all cases, Geneva successfully circumvented censorship.

"Currently, the evade-detect cycle requires extensive manual measurement, reverse engineering and creativity to develop new means of censorship evasion," said Kevin Bock (B.S. '17, M.S. '18, computer science), a computer science Ph.D. student at UMD and lead author of the paper. "With this research, Geneva represents an important first step in automating censorship evasion."

The researchers plan to release their data and code in the hopes that it will provide open access to information in countries where the internet is restricted. The team acknowledges that there may be many reasons why individuals living under autocratic regimes might not want or be able to install the tool on their computers. However, they remain undeterred. The researchers are exploring the possibility of deploying Geneva on the computer supplying the blocked content (known as the server) rather than on the computer searching for blocked content (known as the client). That would mean websites such as Wikipedia or the BBC could be available to anyone inside countries that currently block them, such as China and Iran, without requiring the users to configure anything on

their computer.

"If Geneva can be deployed on the server side and work as well as it does on the client side, then it could potentially open up communications for millions of people," Levin said. "That's an amazing possibility, and it's a direction we're pursuing."

More information: The paper "Geneva: Evolving Censorship Evasion Strategies," Kevin Bock, George Hughey, Xiao Qian and Dave Levin will be presented at the 26th Association for Computing Machinery Conference on Computer and Communications Security in London, England, on November 14, 2019.

Provided by University of Maryland