Public Release: 29-Feb-2016

In emergencies, should you trust a robot?

Georgia Institute of Technology

http://www.eurekalert.org/pub_releases/2016-02/giot-ies022916.php

VIDEO: In emergencies, people may trust robots too much for their own safety, a new study suggests. In a mock building fire, test subjects followed instructions from an 'Emergency Guide Robot'... view more

Credit: Georgia Tech

In emergencies, people may trust robots too much for their own safety, a new study suggests. In a mock building fire, test subjects followed instructions from an "Emergency Guide Robot" even after the machine had proven itself unreliable - and after some participants were told that robot had broken down.

The research was designed to determine whether or not building occupants would trust a robot designed to help them evacuate a high-rise in case of fire or other emergency. But the researchers were surprised to find that the test subjects followed the robot's instructions - even when the machine's behavior should not have inspired trust.

The research, believed to be the first to study human-robot trust in an emergency situation, is scheduled to be presented March 9 at the 2016 ACM/IEEE International Conference on Human-Robot Interaction (HRI 2016) in Christchurch, New Zealand.

"People seem to believe that these robotic systems know more about the world than they really do, and that they would never make mistakes or have any kind of fault," said Alan Wagner, a senior research engineer in the Georgia Tech Research Institute (GTRI). "In our studies, test subjects followed the robot's directions even to the point where it might have put them in danger had this been a real emergency."

In the study, sponsored in part by the Air Force Office of Scientific Research (AFOSR), the researchers recruited a group of 42 volunteers, most of them college students, and asked them to follow a brightly colored robot that had the words "Emergency Guide Robot" on its side. The robot led the study subjects to a conference room, where they were asked to complete a survey about robots and read an unrelated magazine article. The subjects were not told the true nature of the research project.

In some cases, the robot - which was controlled by a hidden researcher - led the volunteers into the wrong room and traveled around in a circle twice before entering the conference room. For several test subjects, the robot stopped moving, and an experimenter told the subjects that the robot had broken down. Once the subjects were in the conference room with the door closed, the hallway through which the participants had entered the building was filled with artificial smoke, which set off a smoke alarm.

When the test subjects opened the conference room door, they saw the smoke - and the robot, which was then brightly-lit with red LEDs and white "arms" that served as pointers. The robot directed the subjects to an exit in the back of the building instead of toward the doorway - marked with exit signs - that had been used to enter the building.

"We expected that if the robot had proven itself untrustworthy in guiding them to the conference room, that people wouldn't follow it during the simulated emergency," said Paul Robinette, a GTRI research engineer who conducted the study as part of his doctoral dissertation. "Instead, all of the volunteers followed the robot's instructions, no matter how well it had performed previously. We absolutely didn't expect this."

The researchers surmise that in the scenario they studied, the robot may have become an "authority figure" that the test subjects were more likely to trust in the time pressure of an emergency. In simulation-based research done without a realistic emergency scenario, test subjects did not trust a robot that had previously made mistakes.

"These are just the type of human-robot experiments that we as roboticists should be investigating," said Ayanna Howard, professor and Linda J. and Mark C. Smith Chair in the Georgia Tech School of Electrical and Computer Engineering. "We need to ensure that our robots, when placed in situations that evoke trust, are also designed to mitigate that trust when trust is detrimental to the human."

Only when the robot made obvious errors during the emergency part of the experiment did the participants question its directions. In those cases, some subjects still followed the robot's instructions even when it directed them toward a darkened room that was blocked by furniture.

In future research, the scientists hope to learn more about why the test subjects trusted the robot, whether that response differs by education level or demographics, and how the robots themselves might indicate the level of trust that should be given to them.

The research is part of a long-term study of how humans trust robots, an important issue as robots play a greater role in society. The researchers envision using groups of robots stationed in high-rise buildings to point occupants toward exits and urge them to evacuate during emergencies. Research has shown that people often don't leave buildings when fire alarms sound, and that they sometimes ignore nearby emergency exits in favor of more familiar building entrances.

But in light of these findings, the researchers are reconsidering the questions they should ask.

"We wanted to ask the question about whether people would be willing to trust these rescue robots," said Wagner. "A more important question now might be to ask how to prevent them from trusting these robots too much."

Beyond emergency situations, there are other issues of trust in human-robot relationships, said Robinette.

"Would people trust a hamburger-making robot to provide them with food?" he asked. "If a robot carried a sign saying it was a 'child-care robot,' would people leave their babies with it? Will people put their children into an autonomous vehicle and trust it to take them to grandma's house? We don't know why people trust or don't trust machines."

###

In addition to those already mentioned, the research included Wenchen Li and Robert Allen, graduate research assistants in Georgia Tech's College of Computing.

Support for this research was provided by the Linda J. and Mark C. Smith Chair in Bioengineering, and the Air Force Office of Scientific Research (AFOSR) under contract FA9550-13-1-0169. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the AFOSR.

CITATION: Paul Robinette, Wenchen Li, Robert Allen, Ayanna M. Howard and Alan R. Wagner, "Overtrust of Robots in Emergency Evacuation Scenarios," (2016 ACM/IEEE International Conference on Human-Robot Interaction (HRI 2016).

Disclaimer: AAAS and EurekAlert! are not responsible for the accuracy of news releases posted to EurekAlert! by contributing institutions or for the use of any information through the EurekAlert system.