

Researchers creating improved interfaces to help machines and humans work together to complete tasks

June 22 2016, by Jesse Delaughter



Credit: Xi Jessie Yang

As machines become more intelligent, they become embedded in countless facets of life. In some ways, they can act almost as full-fledged

members in human-machine teams. In such cases, as with any team, trust is a necessary ingredient for good performance.

But the dynamics of trust between people and machines are not yet well-understood. With a two-year project funded through the SUTD-MIT Postdoctoral Fellows Program, a collaboration between MIT and the Singapore University of Technology and Design (SUTD), postdoc Xi Jessie Yang, MIT Department of Aeronautics and Astronautics Assistant Professor Julie Shah, and SUTD Engineering Product Development Professor Katja Hölttä-Otto aim to develop greater knowledge in this area. During her fellowship, Yang spent one year as a postdoc at SUTD, and is now finishing her research during a second year at MIT. She and her faculty mentors hope their efforts will contribute to better user interface design and better human-machine team performance.

Some tasks can simply be accomplished more effectively and efficiently by a machine or with assistance from a machine than by a human alone. But machines can make mistakes, especially in scenarios characterized by uncertainty and ambiguity. As such, a proper calibration of trust towards machines is a prerequisite of good human-machine team performance.

A variety of factors intersect to determine the extent to which the human trusts the machine, and in turn, how well they can perform on a given task. These factors include how reliable the machine is, how self-confident the human is, and the concurrent workload (i.e. how many tasks the human is responsible for carrying out simultaneously). Teasing out the effect of any one factor can be challenging, but through intricate experimental design, Yang, Shah and Hölttä-Otto have been able to approach the topic in a more realistic way, dealing with the imperfection of machines and isolating various factors in the human-machine trust relationship to study their impact.

At SUTD, Yang and Hölttä-Otto designed an experiment involving a human participant performing a memory and recognition task with the assistance from an automated decision aid. In the experiment, participants memorized a series of images and later recognized them in a pool of similar images. The decision aid provided recommendations for the images to select, very reliably for some people and not so for others. Results revealed that participants tended to undertrust highly reliable aids and overtrust highly unreliable ones. The trust-reliability calibration also improved as humans gained more experience using the decision aid.

At MIT, Yang and Shah explored how interface design could facilitate trust-reliability calibration. Specifically, of interest to Yang and Shah is the potential benefits of likelihood alarm displays. In contrast to traditional binary alarms indicating only "threat" or "no threat", likelihood alarms provide additional information about the confidence level and the urgency of an alerted event. Yang and Shah believe that likelihood alarm would help mitigate the "cry wolf" effect, a phenomenon commonly observed in high-risk industries. In such industries, the threshold to trigger an alarm are oftentimes set very low in order to capture every critical event. The low threshold, however, inevitably results in false alarms which cause the human users to lose trust in the alarm system and eventually harm task performance.

In another experiment, currently underway, Yang and Shah simulate a multi-task scenario in which participants act as a soldier to perform a search and detection task. They have to maintain level flight of four unpiloted aerial vehicles (UAVs) and at the same time detect potential threats in the photo streams sent back from the UAVs. Either a binary alarm or a likelihood alarm helps the participants to complete their mission. The latter would presumably alleviate the "cry wolf" effect, as well as improve participants' attention allocation, and enhance human-machine team performance.

Ultimately, this research could help inform robotics and artificial intelligence development across many domains, from military applications like the simulations in the experiments to medical diagnostics to airport security.

"Lack of trust is often cited as one of the critical barriers to more widespread use of AI and autonomous systems," Shah says. "These studies make it clear that increasing a user's trust in the system is not the right goal. Instead we need new approaches that help people to appropriately calibrate their trust in the system, especially considering these systems will always be imperfect. "

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Researchers creating improved interfaces to help machines and humans work together to complete tasks (2016, June 22) retrieved 12 May 2024 from <https://phys.org/news/2016-06-interfaces-machines-humans-tasks.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.