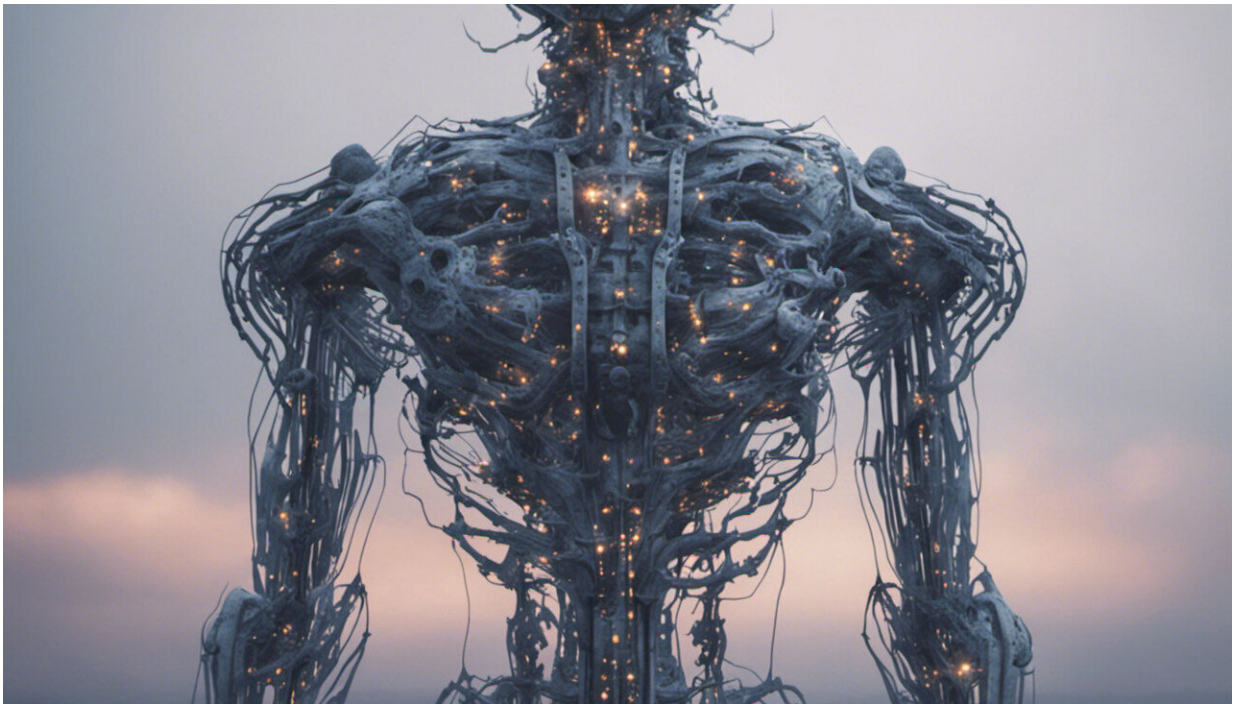# Humans and machines can improve accuracy when they work together

March 11 2019, by Davide Valeriani

Credit: AI-generated image (disclaimer)

Whether artificial intelligence systems steal humans' jobs or create new work opportunities, people will need to work together with them.

In my research I use sensors and computers to monitor how the brain itself processes decision-making. Together with another brain-computer

interface scholar, [Riccardo Poli](#), I looked at one example of possible human-machine collaboration – situations when police and security staff are asked to keep a lookout for a particular person, or people, in a crowded environment, such as an airport.

It seems like a straightforward request, but it is actually really hard to do. A security officer has to monitor several [surveillance cameras](#) for many hours every day, looking for suspects. Repetitive tasks like these are [prone to human errors](#).

Some people suggest these tasks should [be automated](#), as machines do not get bored, tired or distracted over time. However, [computer vision algorithms](#) tasked to recognize faces [could also make mistakes](#). As my research has found, [together, machines and humans](#) could do much better.

Two types of artificial intelligence

We have [developed two AI systems](#) that could help identify target faces in crowded scenes. The first is a facial recognition algorithm. It analyzes images from a [security camera](#), identifies which parts of the images are faces and compares those faces with an image of the person that is sought. When it identifies a match, this algorithm also reports how sure it is of that decision.

People and computers were asked to look at images like this briefly and then identify whether they had seen a particular face. Credit: ChokePoint data, NICTA

The second system is a brain-computer interface that uses sensors on a person's scalp, looking for neural activity related to confidence in decisions.

We conducted an experiment with 10 human participants, showing each of them 288 pictures of crowded indoor environments. Each picture was shown for only 300 milliseconds – about as long as it takes an eye to blink – after which the person was asked to decide whether or not they

had seen a particular person's face. On average, they were able to correctly discriminate between images with and without the target in 72 percent of the images.

When our entirely autonomous AI system performed the same tasks, it correctly classified 84 percent of the images.

Human-AI collaboration

All the humans and the standalone algorithm were seeing the same images, so we sought to improve the decision-making by [combining the actions of more than one](link) of them at a time.

To merge several decisions into one, we weighted individual responses by decision confidence – the algorithm's self-estimated confidence, and the measurements from the humans' brain readings, transformed with a machine-learning algorithm. We found that an average group of just humans, regardless of how large the group was, did better than the average human alone – but was less accurate than the algorithm alone.
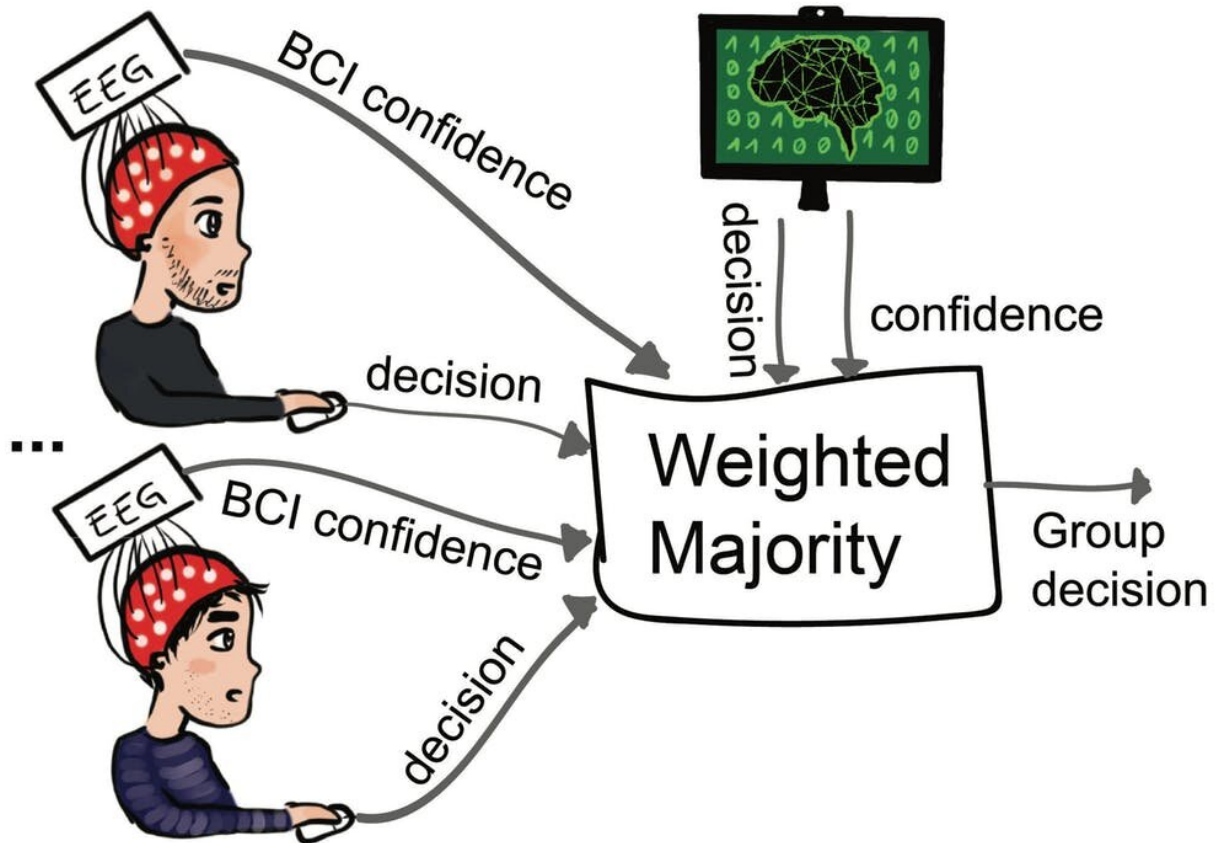
Factoring humans' decisions, and confidence level in those choices, together with algorithmic judgments, yields a more accurate result than people or machines can deliver independently. Credit: Davide Valeriani and Eleonora Adami, CC BY-ND

However, groups that included at least five people and the algorithm were statistically significantly better than humans or machine alone.

Keeping people in the loop

Pairing people with computers is getting easier. Accurate computer vision and image processing software programs are common in airports and other situations. Costs are dropping for consumer systems that read

brain activity, and they provide reliable data.

Working together can also help address concerns about the ethics and bias of algorithmic decisions, as well as legal questions about accountability.

In our study, the humans were less accurate than the AI. However, the brain-computer interfaces observed that the people were more confident about their choices than the AI was. Combining those factors offered a useful mix of accuracy and confidence, in which humans usually influenced the group decision more than the automated system did. When there is no agreement between humans and AI, it is ethically simpler to let humans decide.

Our study has found a way in which machines and algorithms do not have to – and in fact should not – replace humans. Rather, they can work together with people to find the best of all possible outcomes.

This article is republished from The Conversation under a Creative Commons license. Read the original article.

Provided by The Conversation