

Playing nicely with your robot

Jeremy M Wolfe^{1,2}, Makaela S Nartker¹

¹ Surgery, Brigham & Women's Hospital

² Radiology & Ophthalmology, Harvard Medical School

When humans and AI agents collaborate on visual decisions, it does not always go well. In breast cancer screening, for example, the combination of expert radiologists with expert AI does not produce as much benefit as signal detection theory might predict. We study this interaction with an artificial 2AFC task where observers classify textures as positive or negative based on average color. Signal and noise color mixtures are drawn from overlapping normal distributions ($d'=2.2$). Observers can be helped by a simulated AI (also $d'=2.2$). We tested three conditions. 1) AI acts as a “second reader”. After the human response, AI states if it disagrees with the human decision. Humans can then choose to change or not change their response. 2) AI could “triage” trials before human response, showing humans only those stimuli that might plausibly contain signals. The criteria for making these AI decisions is a parameter that can be manipulated. For example, for triage, it makes sense to set a ‘liberal’ criterion so that AI makes very few “miss” errors. Method 3 combines 1 & 2 on the same trial. When these AI rules are used in a block of trials with 50% signal prevalence, AI helps the observer. However, when signal prevalence is 10%, second reader AI actually hurts performance and triage does little or nothing to help. Interestingly, if the same AI information is used both for triage and as a second reader on the same trial, performance is improved at both 10% and 50% prevalence. Asking your AI about its opinion in two different ways may be useful. This example comes from just one set of parameters. The method is flexible and can reveal rules of collaborative perceptual decision making. These experiments illuminate circumstances where intelligent algorithms can be helpful to human experts.

Acknowledgement: NIH EY017001, CA207490