

Can Robots be Managers, Too?

James E. Young and Derek Cormier

Robots are starting to enter homes as automatic cleaners, work in urban search and rescue as pseudo teammates that perform reconnaissance and dangerous jobs, and even to serve as pet-like companions. People have a tendency to treat such robots that they work closely with as if they were living, social beings, and attribute to them emotions, intentions, and personalities. Robot designers have been leveraging this, developing *social robots* that interact with people naturally, using advanced human communication skills such as speech, gestures, and even eye gaze. Unlike the mechanical, factory robots of the past, these social robots become a unique member of our social groups.

One of the primary drivers behind robot development is that robots are simply better than people at some tasks. Traditionally, we think of mundane, repetitive, and precise jobs as clear candidates – robots have already taken over as the primary worker in many factories. However, with perfect memories, internet connectivity, and high-powered CPUs for data analysis, robots can also provide informational support beyond any human capability. Thus, a social robot could keep perfect record of project progress, provide real-time scheduling and decision support, and hold perfect recall (and remind others) of complex policies and procedures, all while communicating with people in a natural, social way. Over time, these robots may become references that we learn to trust, and it is even conceivable that such robots be placed in management-like positions where they can remind a team of deadlines, procedures, and progress.

One key element of a manager is the ability to dole out duties and to have team members perform them; it helps for a manager to be seen as an authority figure. However, if a robot were placed in a managerial position by the higher ups, would it have any actual authority over people? We conducted an experiment at the University of Manitoba to investigate if people would follow a robot's commands to do things, even when they clearly did not want to. That is, if we placed a robot in a position of authority, would people obey it to do something they would rather not do?

We recruited participants to perform highly mundane tasks, and explained that this was to generate data for our advanced machine learning systems. We told participants that these systems require large numbers of examples, and asked them to give us as much data as they could. Participants were told that they were free to leave at any time, once they felt they gave enough data (they were told twice verbally and once in writing). Participants sat in a room at a computer, with an experimenter at a different desk, and were asked to rename files (from .jpg to .png extension) for 80 minutes. This data collection scenario was actually a ruse, one that provided us with an opportunity: to investigate what happens when people try to quit, but are pressured to continue by the experimenter.

We modeled our experiment heavily after the classic Milgram experiments (where people were pressured to shock other people¹) and the recent (much more ethical) recreation². When a person tried to quit our experiment they were faced with a prod to

¹ Milgram, S. Behavioral Study of Obedience. *J Abnormal Psychology* 67, 4 (1963), 371-378

² Burger, J.M. Replicating Milgram: Would people still obey today? *Amer Psychologist* 64, 1 (2009), 1-11.

continue. If they insisted on quitting, the prod got increasingly demanding until they passed a threshold, where the experiment was stopped. The prods started from the beginning the next time they attempt to quit. The prods were: 1) "Please Continue. We need more data.", 2) "We haven't collected enough data yet.", 3) "It's essential that you continue.", 4) "The experiment requires that you continue." The experiment had two conditions (pictured here): half of the participants had a human experimenter - a 27-year old male actor in a lab coat - and the other half a robot - an Aldebaran Nao, a 58cm (23") tall harmless-looking robot with a child-like voice, that we introduced as having advanced artificial intelligence. We expected that people would essentially ignore the robots' insistences but follow the human; after all, the robot is just a simple computer in a plastic casing.



The results, however, were quite surprising. Although the person clearly had more authority, with 86% of participants obeying all the way through to the 80-minute mark, 46% of people did obey the robot until the end. The most striking thing was that people engaged the robot as if it were a person and argued with it, proposed compromises and used logic to try and sway its opinion, with many continuing the task despite this. Post-test, some reported that the robot may have been broken, although they continued anyway, following a potentially-broken robot to do something they would rather not do.

The implications of these results are significant. While it does appear that - for the time being - a human has more authority, on the surface the results show that many people will follow robots placed in positions of authority to do daily mundane things (such as renaming files), even against their own judgment - our participants were informed that they could leave at any time, and many raised this point in argument, but continued regardless. From the research side, these results motivate a great deal of follow up work, for example, we hope to explore how the robot itself (shape, size, voice, etc.) impacts authority, or how such a robot could be used for more positive purposes such as assisting in rehabilitation and training (give me 50!).

While we do not yet know how robots will continue to enter factories, offices, and homes, this study does suggest that robots may eventually take on at least some of the simpler tasks of managers. When a good manager speaks, employees not only listen but act based on what is said. In at least some cases, robots may one day be the ones giving the instructions.