Intent communication of highly autonomous robots

Hannu Karvonen and Iina Aaltonen (VTT Technical Research Centre of Finland Ltd)

INTRODUCTION

With highly autonomous robots, which are both self-directed and self-sufficient, the functioning of the technical system operating the robot can become opaque. Therefore, human beings interacting with the robot may frequently wonder on such questions about the robot as: What is it currently doing? Why is it doing that? What it is going to do next? These types of questions tell that often autonomous robots are not designed well enough from the human factors perspective to communicate their state, processes and intents in an understandable way for the involved humans interacting with the robots.

Therefore, autonomous robots should be able to signal their own status, including, for example, their current goals, stance, state of knowledge, limitations, and upcoming actions to coordinate with humans. In essence, an autonomous robot also needs to make its targets, changes, and upcoming actions externally accessible to the people who interact with them as users or manage, maintain and re-direct them as their administrators. This requirement runs counter to the advice that is sometimes given to autonomous system developers to create systems that need little as little human interaction as possible and are barely noticed. However, we see that people need to have a model of the robot as an agent participating in the joint activity. Although people can often effectively use their own thought processes as a basis for inferring the way their teammates are thinking, this self-referential heuristic does not usually apply when interacting with autonomous robots. Therefore, there is a growing concern that intelligent agents in human-agent teams should do a better job of communicating their current states, capacities, and intents (e.g., Norman, 1993).

In this short paper, we discuss examples from different domains on how robots communicate their intent to humans. By '*intent*', we refer to The Webster's New World Dictionary (1970), which defines it as '*a purpose; object; aim*'. Consequently, in communicating its intent, an autonomous robot may describe both its goals (object or aim) and the reason for pursuing these goals (purpose). This type of description provides the human interacting with the robot feedback about the robot's functioning and also a basis for adapting his or her behaviour and attitude (e.g., trust) towards the robot. Furthermore, the robot's feedback provides a framework for adapting the human's existing plans and procedures to respond to novel situations (see, e.g., Shattuck, 1995).

EXAMPLES OF INTENTION COMMUNICATION OF MACHINE-LIKE INDUSTRIAL ROBOTS

The most recent advances in human-robot interaction research in industrial environments have been made towards dispensing the physical safety barriers and allowing humans to work closer to robots. In order to ensure the safety of humans, the robots need to be either enough lightweight not to cause injuries or other safety-approved mechanisms need to detect humans and then slow down or stop completely. Either way, the work can go more smoothly if the intentions are made clear to the human interacting with the robot.

In addition, in industrial settings it has been studied that legible motion, planned to clearly express the robot's intent, leads to more fluent collaborations than predictable motion, planned to match the collaborator's expectations (Dragan, Bauman, Forlizzi, & Srinivasa, 2015). Therefore, motion planning in human-robot interaction should focus on legibility (Dragan, Lee, & Srinivasa, 2013)**.**

On the other hand, motion planning may become unnecessary. Reinforcement learning has been used to demonstrate that robots can autonomously adapt their behaviour so that humans recognize the robot's intentions early and robustly, and that robots are able to do so without a model of how humans predict motion intentions, or knowledge of the concept of legibility (Stulp, Grizou, Busch, & Lopes, 2015). Additionally, human-aware motion planners have been suggested (Alami, Clodic, Montreuil, & Sisbot, 2006)**.**

Furthermore, in the industrial setting, a robot's intentions can be visualized by projecting the intended trajectory on the objects. In a system by Augustsson, Olsson, Christiernin, & Bolmsjö (2014), the human worker first gets familiar with the robot's intended path using a simulated model. At the physical work site, the starting point of the same trajectory is projected on the real objects. The robot follows a programming rule (movement direction top to bottom, left to right), and the human workers knows where the robot is heading (Augustsson et al., 2014).

In addition, augmented reality (AR), i.e., virtual images overlaid on real objects in real time (Azuma, 1997), seems to be a general solution for displaying intentions of robots systems to operators and persons entering proximity of robotic devices. Carff, Johnson, El-Sheikh, & Pratt (2009) have elaborated this well in a 'Play-Forward' feature '*that would allow one to observe the intention of a robot by displaying the expected action of an ongoing or upcoming plan".* Carff et al., (2009) continue that *'by using mixed reality displays, we can superimpose real-time imagery with virtual processed imagery to give the operator a sense of the data, goals, and intentions of the system, improving situational awareness.*' The challenge here is, of course, how to activate outsiders and audience to watch the scene using AR equipment. We see that as AR headsets evolve and become more commonplace in the future, it will ease this challenge considerably.

In contrast, instead of displaying explicit trajectories, robot's intent has been suggested to be relayed using back-and-forth motion**.** In the study, a robot informed humans of a need to remove an obstacle on its path (Kobayashi & Yamada, 2005). Similarly, the robot can ask humans for help by first getting their attention by following a human and showing their intention through its actions, similar to how dogs go find their owners and make the owners follow them (Nicolescu & Matarić, 2001). In another example, robots status was indicated using idling movement or a wavy path (Nakata, Sato, & Mori, 1998).

Finally, advanced warning systems have been suggested for informing and warning the humans collocated with robots. One of the suggested systems includes a human active interface based on vibrotactile stimuli, where the vibrations were relayed to the human using a wrist-mounted band with vibrators (Ogorodnikova, 2008). The stimuli should match the human expectations. In the case of flashing lights, the rates of about 3-10 per second (with duration at least 0.05 s) have been recommended for attracting attention. (Ogorodnikova, 2008). Furthermore, it has been discovered that the most noticeable audio signals are: beep with frequency 425 Hz and yeow (descending change in frequency from 800 to 100 Hz every 1.4 s). Reaction time to these signals also decreases with increased signal intensity (Ogorodnikova, 2008).

EXAMPLES OF INTENTION COMMUNICATION OF TELEOPERATED ROBOTS

Teleoperated robots can move on ground, air, and water. An operator can control multiple robots, if they have enough autonomity and the operators receive support from the interface. Several authors have studied interfaces for controlling multiple robots simultaneously. An extensive set of several interface guidelines can be found in (Chen, Barnes, & Harper-Sciarini, 2011), and guidelines for multi-unmanned vehicle system design based on a proposed taxonomy of operator task complexity is available in (Lewis, 2013). A review on models on how the operator's interaction with the unmanned vehicles is divided as a sequence of control episodes is also available (Lewis, 2013). Retargeting attention on the whole team after having controlled one unmanned system, and the adaptivity of systems and switch-off are discussed in (Chen et al., 2011).

A specific case of teleoperated robots that have been nowadays studied vastly are unmanned aerial vehicles (UAVs). The control of UAVs can be either management by consent or management by exception (Goodrich & Cummings, 2015). Latter has problems, e.g., related to automation bias, meaning that given an unreliable system, humans have been found to still approve computer-generated recommendations. According to Ruff, Narayanan, & Draper (2002) management of multiple UAVs, management by consent (in which a human must approve an automated solution before execution) was superior to management by exception (where the automation gives the operator a period of time to reject the solution). Management by consent appeared to provide the best situation awareness ratings, the best performance scores, and the most trust for controlling up to four UAVs. In general, there is some evidence that human operators may not be limited in their ability to control multiple vehicles that need navigation and payload assistance, especially with unreliable automation, but the control of two UAVs might be feasible (see e.g., a review by Goodrich & Cummings, 2015).

For the coordination of unmanned vehicles (UVs), the use of planners has been suggested (Lewis, 2013). Planners are programs that generate plans to satisfy user-specified criteria. Plan libraries are useful when complex, variable, interdependent behaviour is needed, such as missions in which multiple UVs assume different roles at different times. Cooperative planners are useful when intent can be expressed succinctly in an objective function and a near-optimal solution is desired.

EXAMPLES OF INTENTION COMMUNICATION OF HUMANOID ROBOTS

Humanoid robots resemble humans to some extent in their appearance, and they typically have a head or a display and upper extremities, which makes it more natural to convey intentions to humans using gaze or facial expressions or hand gestures. For example, a robot can use eye movements, or body poses (e.g., leaning back or craning its 'neck') to establishing personal space around it (Breazeal & Fitzpatrick, 2000). These types of robots often observe human-human collaboration and use it as a basis for implementing corresponding movements on their robotic arms (Gleeson, MacLean, Croft, & Alcazar, 2013; Haddadi, Croft, Gleeson, MacLean, & Alcazar, 2013). In addition, verbal speech is often included. A review on verbal and non-verbal human-robot interactive communication has been published by Mavridis (2015).

Also, in a study using animated models of simplified humanoid robots, it was found that showing the robot's forethought makes people more sure of their interpretations of robot behaviour, and make the robot seem more appealing and approachable (Takayama, Dooley, & Ju, 2011)**.** Similarly, anticipatory variants of robot's gestures and poses, e.g., beckoning, stopping, pointing, allowed humans to discern the robot's motion intent sooner than motions without anticipation (Gielniak & Thomaz, 2011).

CONCLUSIONS

Without clear explicit or implicit communication, humans are likely to ignore the requests of the robots. As can be seen from the examples in this paper, several approaches, such as legible motion, reinforcement learning, augmented reality, advanced warning systems, or human-like gestures and facial expressions, have been taken to address the issues in human-robot intent communication. In addition to these examples, we suggest that a novel approach could be to use haptic feedback. Also, a promising approach could be gaze-direction detection or vocalizations from the robot or human to determine the intent of either party. However, this doesn't account for when the human has become distracted or if they have lost interest in interacting with the robot.

Clearly, one key issue in human-robot interaction, especially with highly autonomous robots, is trust. When a person doesn't trust another person, they have a significantly low probability of believing what they are being told and therefore, also interacting positively with each other. This concept applies to humans listening to robots as well. People are sceptical of a robot as they are, but when a robot begins operating autonomously along-side them, they become much more critical to the requests of the robot. Therefore, the challenge with building appropriate trust between robots and humans lies in the way that they communicate their intent to each other.

REFERENCES

Alami, R., Clodic, A., Montreuil, V., & Sisbot, E. (2006). Toward human-aware robot task planning. In *AAAI Spring Symp. To boldly go where no human-robot team has gone before* (pp. 39–46).

Augustsson, S., Olsson, J., Christiernin, L. G., & Bolmsjö, G. (2014). How to transfer information between collaborating human operators and industrial robots in an assembly. In *Nordic Conference on Human-Computer Interaction* (pp. 286–294).

Azuma, R. T. (1997). A survey of augmented reality. *Presence*, *6*(4), 355–385.
Breazeal, C., & Fitzpatrick, P. (2000). That Certain Look: Social Amplification of Animate Vision. In *AAAI 2000 Fall Symposium* (Vol. Technical, pp. 18–22).
Carff, J., Johnson, M., El-Sheikh, E. M., & Pratt, J. E. (2009). Human-robot team navigation in visually complex environments. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009* (pp. 3043–3050).
Chen, J. Y. C., Barnes, M. J., & Harper-Sciarini, M. (2011). Supervisory control of multiple robots: Human-performance issues and user-interface design. *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews*, *41*(4), 435–454.
Dragan, A. D., Bauman, S., Forlizzi, J., & Srinivasa, S. S. (2015). Effects of Robot Motion on Human-Robot Collaboration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15* (Vol. 1, pp. 51–58). Portland, Oregon, USA: ACM.
Dragan, A. D., Lee, K. C. T., & Srinivasa, S. S. (2013). Legibility and predictability of robot motion. *ACM/IEEE International Conference on Human-Robot Interaction*, 301–308.
Gielniak, M. J., & Thomaz, A. L. (2011). Generating anticipation in robot motion. *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, 449–454.
Gleeson, B., MacLean, K., Croft, E., & Alcazar, J. (2013). Gestures for Industry: Intuitive Human-Robot Communication from Human Observation. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction* (pp. 349–356). IEEE.
Goodrich, M. A., & Cummings, M. L. (2015). Human Factors Perspective onNext Generation Unmanned Aerial Systems. In K. P. Valavanis & G. J. Vachtsevanos (Eds.), *Handbook of Unmanned Aerial Vehicles* (pp. 2405–2423). Springer Science+Business Media Dordrecht.
Haddadi, A., Croft, E. A., Gleeson, B. T., MacLean, K., & Alcazar, J. (2013). Analysis of task-based gestures in human-robot interaction. In *Proceedings - IEEE International Conference on Robotics and Automation* (pp. 2146–2152). Karlsruhe, Germany.
Kobayashi, K., & Yamada, S. (2005). Informing a User of Robot's Mind by Motion. In *Proceedings of International Conference on Computational Intelligence, Robotics and Autonomous Systems* (p. SS4A-3).
Lewis, M. (2013). Human Interaction With Multiple Remote Robots. In *Reviews of Human Factors and Ergonomics* (Vol. 9, pp. 131–174).
Mavridis, N. (2015). A review of verbal and non-verbal human–robot interactive communication. *Robotics and Autonomous Systems*, *63*, 22–35.
Nakata, T., Sato, T., & Mori, T. (1998). Expression of emotion and intention by robot body movement. In *5th Conference on Intelligent Autonomous Systems* (pp. 352–359).
Nicolescu, M. N., & Matarić, M. J. (2001). Learning and Interacting in Human–Robot Domains. *Systems and Humans*, *31*(5), 419–430.
Norman, D. A. (1993). *Turn signals are the facial expressions of automobiles*. Basic Books.
Ogorodnikova, O. (2008). Creating an active awareness system for humans in robotic workcell. *Acta Polytechnica Hungarica*, *5*(2), 113–126.
Ruff, H., Narayanan, S., & Draper, M. H. (2002). Human Interaction with Levels of Automation and Decision-Aid Fidelity in the Supervisory Control of Multiple Simulated Unmanned Air Vehicles. *Presence: Teleoperators and Virtual Environments*, *11*(4), 335–351.
Shattuck, L. G. (1995). *Communication of intent in distributed supervisory control systems*. Ohio State University.
Stulp, F., Grizou, J., Busch, B., & Lopes, M. (2015). Facilitating intention prediction for humans by optimizing robot motions. In *IEEE International Conference on Intelligent Robots and Systems* (Vol. 2015–Decem, pp. 1249–1255).
Takayama, L., Dooley, D., & Ju, W. (2011). Expressing thought: improving robot readability with animation principles. In *Proceedings of the 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 69–76).
The Webster's New World Dictionary. (1970).