One Robot and Two Humans: Some Notes on Shared Autonomy in the Case of Robotic Telepresence

André Potenza, Alessandro Saffiotti
Center for Applied Autonomous Sensor Systems, Örebro University, Sweden
andre.potenza@oru.se, asaffio@aass.oru.se

Abstract

Telepresence robots, similar to other teleoperated robots, can benefit strongly from shared autonomy as a way to enhance ease of use for the operator. With ever-increasing capabilities of autonomous robots, it is crucial to understand what can be automated and under which circumstances. We argue that within a dynamic environment, the allocation of tasks between human and robot should not be fixed, but rather adaptable, taking into account the current state of the environment.

1 A holistic view of shared autonomy

Any given task, cooperative or not, can be seen as embedded in the context of the system in which it is performed. In human-robot interaction, where human and robot are agents working towards a common goal, we usually view the system as the total of the human, the robot as well as the environment enclosing them. Following this view, the system is denoted as S = H + R + E, where H and R are both embedded in the same environment E. We can further assume that the robot and the human both are capable of performing a given set of functionalities, F_H as the human's capabilities and F_R for the robot, that may contribute in different parts to carrying out the task assigned to the human-robot team. Each set can comprise both physical and cognitive functionalities, ranging from sensing and acting to understanding and decision making.

The entire system S then has to perform a given task T in E, which requires the use of a set of functionalities F_T. If S is able to perform T, each item of F_T occurs in the specific F of at least one agent in S. It is important to note that this is a necessary, though not sufficient condition. Of course, in practice there are additional factors to consider, such as temporal and spatial constraints.

2 Problem landscape in shared autonomy

Using this view, one can identify a number of interesting problems in shared autonomy for a given application domain or case:

(a) What are the overall functionalities F_T required for a given task T?

(b) What are the F_H that can be performed by H?

(c) What are the F_R that can be performed by R?

(d) What is the most efficient way to distribute F_T between F_H and F_R when also taking into account additional aspects, such as for instance interfaces and interdependencies between functionalities?

(e) What new problems are introduced if the allocation of functionalities is performed dynamically, as in adjustable autonomy?

A few examples of the above problems in some concrete cases might look like the following:

- Telemanipulation, e.g., remote surgery using the Da Vinci system. Here, F_R are physical actuation and sensing; F_H comprises interpretation of sensor data, planning and control; R interacts with E directly, while H interacts with R directly and with E only via R. In this scenario, E is the patient.
- Telemanipulation with local autonomy, e.g., underwater manipulation using a smart ROV. Here, F_R are physical actions and physical sensing, plus low-level local navigation control like posture stabilization; F_H entail interpretation of sensor data, planning, high-level navigation control like deciding set points, manipulator control; R interacts with E directly, H interacts with R directly and with E only via R.
- Smart wheelchair. Here, F_R are physical action and low-level motion control like throttle compensation for inclines, and collision avoidance; F_H are all decision making and navigation functionalities plus low-level set points for velocity and orientation; in some cases, F_R may include higher-level motion control and navigation functions, e.g., if the human cannot or does not want to use F_H for those; both R and H interact with E directly, plus H interacts with R.

3 The case of mobile robotic telepresence

Now, consider mobile robotic telepresence. Here, the overall system includes (at least) one more human [Kristoffersson *et al.*, 2013b]. How does this change our picture? In this scenario, the system takes on the form S = H1 + H2 + R + E. Typically, H2 is denoted as the local user, and together with

R embedded in E. H1, whose role is generally described as the remote user or operator, is remotely interacting with R. The task T commonly involves a social interaction between H1 and H2, which may consist in simply a casual social exchange, or, in more special cases, health assessment or consultation.

In mobile robotic telepresence, there are at least two (often concurrent) tasks, the social interaction (T_{Soc}) and robot actuation (T_{Act}). The purpose of the latter is generally to facilitate and support the former. The required functionalities for (T_{Soc}) are almost exclusively located within F_H, whereas the contributing functionalities in T_{Act} can be provided to varying degrees, dependent on the particular H1 and R, by F_H and F_R.

Since the act of controlling a telepresence robot while at the same time engaging in a social exchange with another person has been reported to be challenging [Desai *et al.*, 2011], efforts are being made to reduce the burden on the operator by enabling the robot to navigate semi-autonomously. However, this is evidently no simple endeavor, as environments are highly dynamic and, as of yet, autonomous robots cannot be taught to deal with any conceivable situation. Moreover, given that the robot essentially represents a remote embodiment of themselves, users are likely to favor being in control of it as long as they are capable of doing so comfortably. We therefore have a task which can be carried out by both the human and the robot, though neither of them can do so perfectly and their performance level may vary throughout the execution.

Thus, with the problem outlined above, our goal is to maximize the performance of the system by dynamically leveraging the capabilities of both agents. On top of this, to avoid burdening the operator with keeping track of yet another task, we aim at having the robot monitor the performance of the system and take the initiative to shift control whenever the agent currently in control (robot *or* human) is having trouble in the present situation.

Just how can this be accomplished? Indeed, there are a number of observable metrics that can conceivably serve as indicators of system performance. We distinguish between direct and indirect observations. Direct ones are task-specific and describe how well this task is performed in isolation. For navigation, this can, for instance, be the occurrence (rate) of collisions or adequate positioning of the robot relative to local users. In the case of social interaction there have been different approaches to estimate the interaction quality from a variety of non-verbal cues and modalities [Bensch *et al.*, 2017]. Indirect metrics are more difficult to measure accurately, though they have the advantage of being, to the greatest extent, task-agnostic. In humans, these metrics have been subsumed under the term human factors [Parasuraman *et al.*, 2008] and describe a variety of dynamic characteristics of a person. As an example, mental workload is concerned with the degree to which somebody is occupied by the sum of their current tasks. If the workload is high, this could be used as a reason for the robot to take over a part of the functions being performed. Likewise, if it is low, the user could be reassigned a task. Although the robot is expected to take the initiative, as a practical consideration, we argue that the user should be warned ahead of a task shift in either direction and given the chance to confirm or cancel it. Attempting to estimate the quality of a social interaction is arguably a great deal more difficult. Direct measures may involve speech analysis for engagement or turn taking [Kristoffersson *et al.*, 2013a]. On the other side of the equation, the telepresence robot cannot be expected to perform equally well in any scenario either. We might encounter a situation in which the sensor readings are noisy and inconclusive, or where there are many people standing around and no clear path is discernable. As a result it might even be desirable for the robot to return control in spite of high workload measurements in the operator.

Finally, there is the role of the local user, denoted above as H2. In a social interaction, H1 and H2 collaborate towards achieving a productive exchange, as they would when meeting in person. Hence, they both are expected to provide F_{Soc} and, similar to the operator, the local user's performance may vary. If we set out to try and measure the system's performance in the social exchange task T_{Soc}, it is important to take into account the local user and examining how their satisfaction with the interaction can be measured.

4 Conclusions

The above considerations show that the study of shared autonomy in the case of robotic telepresence systems presents new facets and challenges compared to the study of shared autonomy in the more customary case of teleoperation. The study of these specific challenges constitutes our current line of research within the Socrates ETN project.

References

[Bensch *et al.*, 2017] Suna Bensch, Aleksandar Jevtic, and Thomas Hellström. On interaction quality in human-robot interaction. In *ICAART 2017 Proceedings of the 9th International Conference on Agents and Artificial Intelligence, vol. 1*, pages 182–189. SciTePress, 2017.

[Desai *et al.*, 2011] Munjal Desai, Katherine M Tsui, Holly A Yanco, and Chris Uhlik. Essential features of telepresence robots. In *Technologies for Practical Robot Applications (TePRA), 2011 IEEE Conference on*, pages 15–20. IEEE, 2011.

[Kristoffersson *et al.*, 2013a] Annica Kristoffersson, Silvia Coradeschi, Kerstin Severinson Eklundh, and Amy Loutfi. Towards measuring quality of interaction in mobile robotic telepresence using sociometric badges. *Paladyn, Journal of Behavioral Robotics*, 4(1):34–48, 2013.

[Kristoffersson *et al.*, 2013b] Annica Kristoffersson, Silvia Coradeschi, and Amy Loutfi. A review of mobile robotic telepresence. *Advances in Human-Computer Interaction*, 2013:3, 2013.

[Parasuraman *et al.*, 2008] Raja Parasuraman, Thomas B Sheridan, and Christopher D Wickens. Situation awareness, mental workload, and trust in automation: Viable, empirically supported cognitive engineering constructs. *Journal of Cognitive Engineering and Decision Making*, 2(2):140–160, 2008.