

Improving Task Performance through High Level Shared Control of Multiple Robots with a Context Aware Human-Robot Interface

Simon O’Keeffe, Tomas E. Ward, Rudi Villing
Department of Electronic Engineering
National University of Ireland Maynooth
Maynooth, Ireland

Abstract—In multi-robot control, robots need the ability to perform well in the absence of valid human input. This paper describes a shared control scheme for multiple robots where control can be traded between human and autonomous agents on the fly, reducing the negative effect of robots requiring attention while the user’s attention is devoted elsewhere. The control approach is a hybrid of traded control where the control is traded between the human and autonomous agents and coordinated control where high level human user commands are mapped to low level implementation under partial control of the autonomous system. An allocation authority decides whether the user or autonomous agent is controlling the robot at a given time based on a command context. This allocation authority control scheme is compared to a coordinated control scheme in a multi-robot soccer domain. The results demonstrate that an allocation authority approach produces improved task performance and may have generalizable applicability. Specifically in the context of robotic soccer, the proposed control scheme was shown to have 18.5% more possession and 8.2% more territory than traded control while also reducing mental demands on users.

Keywords—*human-in-the-loop; shared control; multi-robot; allocation authority; human-robot interface*

I. INTRODUCTION

Autonomous robots can often function well in well-structured and static environments (such as an industrial environment) but may encounter difficulties with complex tasks in unstructured environments such as a typical home or rescue environment. Incorporating human user inputs can help the robot complete more complex tasks in less structured environments.

Shared control is based on the premise that in certain contexts humans are superior to robots in areas such as perception and decision making, while there are tasks that the robot can perform better such as precise low-level motion planning, solving an optimization problem, and operating in dirty or dangerous situations [1].

Shared control is a control scheme that causes the output or response of a system to be influenced by two or more agents. In this work, we specifically consider just two agents for the control of each robot: a human user and an autonomous agent. The human user controls each robot in a round-robin method,

interacting with each robot for a time before neglecting it to attend to other robots.

In many shared control schemes, the robot remains stationary and idle, or worse is effectively uncontrolled, in the absence of human input [2]–[6]. In situations where a stationary, idle, or uncontrolled robot results in negative consequences for robot performance or safety then it is necessary for the autonomous agent to take full control in the absence of user input. For example, consider a robot operating in a search and rescue environment. A robot would perform better with human input as humans have better perception and high level planning capabilities. However, if the user input is slow to arrive, the robot should still search for survivors as a person’s chance of survival decreases the longer they remain unbound.

When any aspect of the control is traded back and forth between the user and autonomous agent, an interesting problem arises. Human users may take a relatively long period of time (relative to the autonomous agent) to process information from the robot and decide an appropriate course of action, which they must map to the appropriate set of inputs. This delay is increased for trading control between one user and multiple autonomous agents. Whenever an operator does not devote total attention to a robot there will be a delay between when the robot requires attention and the operator intervenes [7].

Although the environment and game play challenges in standard platform league robot soccer are quite structured, it nevertheless provides an interesting and challenging environment in which to investigate shared control techniques. In particular, the game play challenges dictate that opportunities for user input are bound by time deadlines as a robot player’s performance quickly degrades if it does not act quickly in response to changes resulting from the actions of opposition players. The objective of this work is to investigate the benefit of a shared control scheme where control can be traded between human and autonomous agents on the fly, reducing the negative effect of robots requiring attention while the user’s attention is devoted elsewhere. When a control scheme allows a non-human agent to switch control from the user, they can find it irritating or frustrating. Therefore, it is not clear whether the proposed allocation authority control scheme will outperform traditional coordinated control.

This paper is organised as follows: related work is described in Section 2. Section 3 describes the development and integration of the shared control scheme implemented and the accompanying human-robot interface. Section 4 presents the experimental method used to evaluate the shared control scheme. Results are presented and in Section 5 and further interpreted and discussed in Section 6. Section 7 presents conclusion and future work.

II. RELATED WORK

Six distinct classes of robotic shared control have been identified [8]: traded control, indirect shared control, coordinated control, collaborative control, virtual constraint, and blended shared control. In traded control, the system is either under full manual control or full autonomous control and control sharing is only achieved by transitioning between the two states. With indirect shared control, the user has full control of the system and an autonomous agent attempts to modify the user input by presenting specific information to the user (for example vehicle lane departure feedback). In coordinated control, the user provides low-dimensional control inputs and the autonomous agent maps these (in possibly complex ways) to a set of high dimensional control inputs for a system. Collaborative control involves the user and the autonomous agent having control over different inputs (such as a car where the steering is under user control and the throttle is controlled by smart cruise control). In the virtual constraint approach, an autonomous agent modifies the user's input to satisfy some safety or performance constraint (for example obstacle proximity). Blended shared control allows the user and the autonomous agent to provide inputs simultaneously and a mixing function is used to combine these inputs into a unified control input for the system.

Control of multiple robots is typically achieved through coordinated control, where the operator gives high level tasks to each robot, which the robot accomplishes while the user attends to other tasks. Coordinated control is also referred to as supervisory control in the literature where robots are controlled through manual waypoints or automated path planning [9].

The shared control scheme presented in this paper implements a combination of traded and coordinated control. This hybrid approach allows for high level user inputs, when present to be mapped to lower level, more detailed control inputs. When no user input is present (perhaps after previous commands have been completed or invalidated by a change in world context) control reverts back to the autonomous agent.

Coordinated control is suited, for example, to robotic wheelchairs where users may find it difficult to provide continuous control to the wheelchair. In [4] the user could send three commands to the robot (left, right, and forward) which the robot executes at the next available occasion. Coordinated control approaches also have waypoint based inputs for both wheelchairs [10] and telepresence robots [3]. Coordinated control for multiple robots has typically involved serial command of robots through teleoperation, assigning waypoints, or monitoring and prosecuting targets using a single robot at a time [9]. In [11] the operator was able to control multiple robots using either low level waypoint commands or a more

automated control border command in a simulated capture the flag multi-robot game.

The literature for traded control has typically focused on trading control between full autonomous control and full manual (user) control [12], [13]. In [12], the autonomous agent would take control when precise movements are needed during grasping with the user in full control otherwise. Examining whether the user or the autonomous agent should initiate traded control was examined in [13]. The user could take control if they determined that the robot was struggling with a task. Alternatively, the robot could transfer control to a human user if it could not proceed with a task. Work that has examined high level inputs has typically focused on the problem of mapping these high level (low-dimensional) commands to lower level (higher dimensional) control inputs only [3].

The issue of robots waiting for human input has been considered by [14]. For effective control of the maximum number of robots: (a) robots must be capable of operating in a "degraded state" while waiting for human input, and (b) humans must be able to recognize when robots need attention so that they do not wait longer than needed. One solution is to delay the need for human attention through avoiding simultaneous demands on the user by delaying a robot's arrival at a target [15]. In [16], the authors confront this problem in an implementation using social robots at a shopping mall. A human operator supervises multiple robots and intervenes when needed. By dividing interactions into phases in which human input is critical and not critical, they were able to maintain effective supervisory control over all four robots. No solution deals with the problem of what the robot should do in time critical applications where no input is present.

The concept of allocation authority is discussed by [17]. In contrast to the above techniques, where the human is at least partially in control, an allocation authority decides to give control to either the user or the autonomous agent. An allocation authority is suited to any shared control scheme where the user and autonomous agent provide the same type of inputs to the system. The allocation authority can be part of the responsibilities of the autonomous agent or the user, or it might be an independent function that is separate from either. Allocation authority has been investigated in low level traded control where either the autonomous agent or user acted as the allocation authority [13]. They found that an approach where either the human operator can interrupt the system to take control or the system can request help if needed had a higher task completion time and success rate than an approach where only the system can request help.

Based on this related work, we propose an allocation authority acting as an independent agent where the user and autonomous agent can solely provide high-level inputs with the allocation authority responsible for deciding which to implement. Such a control scheme would be a hybrid traded coordinated control that allows for control of multiple robots while reducing the negative effects of robots waiting for valid human input. As far as the authors are aware there has been no published work describing such a shared control paradigm.

III. SHARED CONTROL IMPLEMENTATION

In this section we describe the implementation of the proposed shared control scheme for control of multiple humanoid robots that can participate in RoboCup Standard Platform League. Robot soccer was chosen as a testbed for this work as it provides a dynamic, semi-structured environment suitable for evaluation. In standard platform league robot soccer, two teams of 5 humanoid robots each play using rules based on real soccer. The existing autonomous behaviour of the robot needed to be modified to allow high level user control of certain tasks such that the shared control robot could perform behaviours that the autonomously controlled robot could not.

The autonomous software platform used as the basis for development was the software platform that had been developed by RoboEireann team members over a number of years. For this experiment we decided to work with only three outfield robots and no goalkeepers. The goalkeeper is a specialist role with special behaviour that relies heavily on reactions and less so on strategic play. Outfield players could either be a striker or a non-striker. One robot is designated as a striker which is the robot that is determined to be in the best position to play the ball. The other robots are tasked with positioning themselves strategically to support the striker for both attacking and defending.

The behaviours are implemented in a hierarchical manner so that lower level skills may be composed to realize higher level behaviours. In fact, a relatively small set of skills may be used to realize a large number of behaviours. Nevertheless, high level behaviours also depend on the availability of appropriate percepts and decision making. Therefore, the robot might have the fundamental skills required for a wide range of new behaviours without necessarily having the necessary perceptual or decision making implementation available.

For this work the set of skills used were the ability to locomote to a particular location and orientation (“move to”) on the playing field and the ability to kick a ball towards a particular target point (“kick to”). With these two commands, sophisticated behaviours can be developed.

A number of components were needed to implement the proposed control scheme. First coordinated control capabilities needed to be established, which exposed the skill commands to the user and determined when they were achieved. The allocation authority was required to determine when to trade control between the user and the autonomous agent. Finally, a human robot interface was required so that user input could be provided to multiple robots in a simple way.

A. Coordinated Control

Coordinated control involves the mapping of lower-dimensional inputs to high-dimensional realisation through potentially complex ways. The “move to” command is a low level skill in the behaviour system. The implementation of the skill is hidden from the user such that the user is not required to consider obstacles in the robot’s path. The implementation of the “kick to” command does not require the user to consider which of the robot’s types of kicks or how much kick power is

needed to achieve a particular type of kick. When the robot is not processing a command from the user, it does nothing.

B. Allocation Authority

The allocation authority was implemented as a distinct function, responsible for implementing either a human command or an autonomous agent command. The allocation authority can be considered a finite state machine with two states, human control and autonomous control. The allocation authority is illustrated in Figure 1. A context change triggers a transition from human control to autonomous control. A new and valid human command then triggers a transition from autonomous control to human control. Furthermore, a new human command overrides previous commands when under human control.

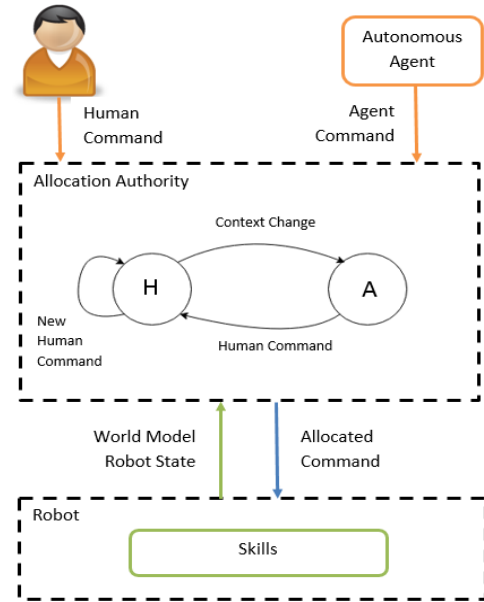


Figure 1: Allocation Authority is responsible for allocating control depending on context

The context describes the world model and robot state. When a user command is issued, the associated context is saved with the command. This context is continuously evaluated, and upon detecting a change in context the allocation authority invalidates the human command and allocates control back to the autonomous agent. For this work the context is heavily influenced by the ball position and visibility. For example, if a human command is issued when the robot not the closest robot to the ball and while executing the command the ball suddenly comes close to the robot, the robot is now the closest robot to the ball. Then the command context changes and the human’s command is invalidated, allocating control to the autonomous agent. Although this is a fairly simple allocation strategy it is sufficient for investigating the efficacy of the allocation authority shared control scheme.

C. Human-Robot Interface (HRI)

To implement this control scheme, a suitable interface is required. The HRI has two functions: it provides robot information to the user so that they may make control decisions

and it accepts commands from the user and communicates them to the robot. For this initial work the HRI supports two commands. The robot can be commanded to go to a specific target location and orientation on the field, or else the kick target point can be set.

The main display contains the layout of a soccer field as shown in Figure 2. The interface shows the position and orientation of all robots under control, it shows the current ball position (relative to the selected robot), and the selected robots target position and orientation. Target positions next to the ball indicate the position the robot will take to execute a kick.

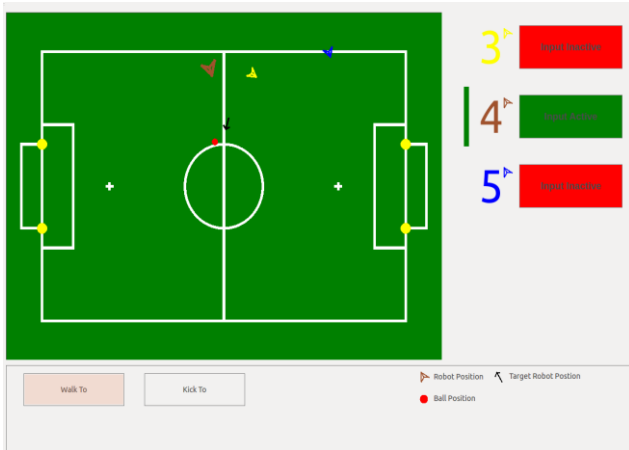


Figure 2: User interface for controlling multiple robots

The right hand side of the interface shows all the robots available to control. The user can select which robot to control by clicking one of the robot displays that appears on this side. This robot display shows the current status of each robot to show the authority over each robot. In contrast when using the coordinated control scheme the status shows when the robot is currently executing a human command or waiting for a command.

To command a target position for a robot the user selects the robot to command, selects the “move to” command, clicks on the field to set field position followed by a second click, relative to the first click, which gives orientation. The “kick to” command is issued the same way however it only requires one click on the field.

IV. EXPERIMENTS

The experiments were designed to investigate whether the proposed allocation authority method of control could outperform traditional coordinated control, where in the absence of human input the robot remains stationary or idle. This coordinated control approach is used in many shared control schemes [2], [3] and is appropriate when the human should have full control over the robot’s actions and the robot should not make high level decisions in the absence of human input.

For the experiment, eight naïve users were asked to control a team of three robots playing a robot soccer game against a team of autonomous robots, using the RoboEireann team codebase that competed in RoboCup 2015. This provides a

standardized opponent. The group consisted of five male and three female users aged between 19 - 26.

The users used both coordinated control and allocation authority control methods in two separate trials. The first control method used was alternated between participants. We decided to limit the number of robots under control to three as previous research found that operator performance dropped off when controlling more than three robots using coordinated control [9].

Each trial lasted 10 minutes, which is one half of a RoboCup Standard Platform League match. Three outfield robots were on each team, with no goalkeepers. Users were given a brief description of the interface and got 5 minutes to practice using the commands through the interface on one robot.

The metrics we used to evaluate task performance were the result of the game, the amount of possession, and the amount of territory enjoyed by the user’s team. The result of the game is an obvious metric as a better result against an autonomous team indicates a better performance. Possession and territory metrics give us richer information regarding the trial, and allow us to analyse the game even if the game ends in a draw. A robot was deemed in possession of the ball if the robot was clearly the closest to the ball up to a distance of 1 meter. Territory was calculated by the percentage of time the ball spent in the opponent’s half.

The average number of commands users sent per second was also calculated. This metric determines how much interaction is required between both modes of control. Users were asked to rate their performance on a NASA TLX scale [18]. This scale measures operator perceived workload, mental demand, temporal demand, and performance.

V. RESULTS

The results of the games are displayed in Table 1. In general participants performed better using the allocation authority control scheme over coordinated control. Three users managed to win against the autonomously controlled robots with the allocation authority control scheme. Two users managed to win with the coordinated control approach. These implemented a defensive strategy taking long shots to score goals. These users still had more possession while using allocation authority control.

Table 1: Results for Shared Control and Coordinated Control versus autonomous robots (W- Win, L – Loss, D - Draw)

User	1	2	3	4	5	6	7	8
Coordinated Control	0-4 (L)	0-2 (L)	1-2 (L)	0-4 (L)	2-1 (W)	0-4 (L)	2-1 (W)	1-2 (L)
Allocation Authority	2-1 (W)	1-1 (D)	0-2 (L)	1-2 (L)	1-3 (L)	2-1 (W)	2-1 (W)	2-2 (D)

Possession levels are shown in Figure 3. For all users possession increased when using allocation authority. The

mean possession increased from 35% with coordinated control to 53.5% with allocation authority control.




Figure 3: User Possession (%) for coordinated control and allocation authority

Territory levels are shown in Figure 4. Although the results are not quite as consistent, 5 of the 8 participants gained more territory, keeping the ball in the opponent’s half for longer than they did in the coordinated control scheme. The mean territory increased from 34.9% with coordinated control to 43.1% with allocation authority control.




Figure 4: User Territory (%) for coordinated control and allocation authority

The number of commands users sent per second was also calculated. This can be seen in Figure 5. Users typically sent slightly fewer commands using the allocation authority control method. The mean number of commands reduced from 0.1018 commands per second to 0.0804 commands per second.

Regarding the NASA TLX score, users generally rated a reduced mental demand, with an average reduction of 10%. Users also perceived 11.25% lower temporal demand and 14.38% lower frustration when controlling the robots with the allocation authority scheme.

VI. DISCUSSION

Users won more games using the presented hybrid shared control scheme over traditional coordinated control. Users also had on average 18.5% more possession and 8.2% more territory using the allocation authority control scheme. Along with performance improvements, users also perceived a reduced mental demand, a reduced temporal demand, and reduced frustration levels when using the allocation authority control scheme. This frustration measure is interesting as it might be expected that users could be more inclined to lose

situational awareness when robots switch to autonomous behaviour, disobeying the user’s commands. The ability to monitor the progress of all robots through the interface potentially reduced the effects of loss of situational awareness.




Figure 5: Commands per Second for coordinated control and allocation authority

Although the shared control scheme implemented for this work represents an early proof of concept a number of interesting issues have been uncovered for further exploration. One of these is the rule which impacts the allocation of control. Control was allocated to the autonomous agent whenever the associated command context for the current command changed. At some instances during the trials, users’ commands were invalidated by an immediate context switch at which point the users proceeded to re-issue the original command again. This suggests that the user may have felt there were more advantageous possibilities available than those the allocation authority was pursuing. In such a situation it may be worth exploring an approach where the allocation authority backs off if the user repeatedly gives the same command. Whenever control was allocated back to the autonomous agent, the robot switches to executing its default autonomous behaviour. However, perhaps a future autonomous agent could use the user commands to alter its autonomous behaviour.

In this work we implemented a hybrid approach of traded and coordinated control and compared it to traditional coordinated control. However, we did not compare it to a control scheme where control could be traded back to the autonomous agent whenever the users command is completed as traditional coordinated control is typical to shared control schemes. We expect our proposed allocation authority control scheme, which constantly monitors human commands for validity, to outperform such a traded control scheme.

Due to the structured environment of robot soccer, the autonomous agent has a high neglect tolerance and should be able to operate for an entire game without the user input. Although such a high neglect tolerance is not present in all robotic domains it is likely to become more common as the state of the art in autonomous robotics continues to advance. The allocation authority control approach presented could be expected to operate best in environments with a high neglect tolerance but we anticipate the scheme would also be useful in any situation where a human could outperform the autonomous perception and decision making capabilities of a robot, at least temporarily. Specifically, it would have utility in a dynamic assisted living environment, either operating multiple robots or one robot while also attending to other tasks.

VII. CONCLUSION

This work presented in this paper describes a hybrid coordinated traded control approach for multiple robots. In this approach, the user and the autonomous agent provide inputs into the system with an independent allocation authority responsible for trading control between the user and the autonomous agent. This approach ensures that in the presence of time pressures in a dynamic environment, robots constantly have valid input to prevent the negative consequences of robots remaining stationary or executing unsuitable commands.

The hybrid shared control approach was shown to outperform coordinated control with users winning more games using the allocation authority scheme while also enjoying 18.5% more possession and 8.2% more territory. Users also perceived 10% lower mental demand and 11.25% lower temporal demand along with reduced frustration levels of 14.38% when using the presented control scheme. Generalization of this result to other operational contexts is not trivial, but the robot soccer domain studied contains many of the challenges facing collaborative robotic systems in semi-structured domains with dynamic contexts. Consequently, it is a result that should encourage further research in this area of human-robot control systems.

ACKNOWLEDGMENT

The authors would like to gratefully acknowledge funding provided by the Irish Research Council under their Government of Ireland Postgraduate Scholarship 2013.

REFERENCES

- [1] V. Dimitrov and T. Padir, "A shared control architecture for human-in-the-loop robotics applications," in *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, 2014, pp. 1089–1094.
- [2] R. Jarvis, "A semi-autonomous robotic walk-chair," in *Autonomous Robots and Agents, 2009. ICARA 2009. 4th International Conference on*, 2009, pp. 33–37.
- [3] B. Coltin, J. Biswas, D. Pomerleau, and M. Veloso, "Effective semi-autonomous telepresence," in *RoboCup 2011: Robot Soccer World Cup XV*, Springer, 2011, pp. 365–376.
- [4] J. R. Millan, F. Renkens, J. Mouriño, and W. Gerstner, "Noninvasive brain-actuated control of a mobile robot by human EEG," *Biomedical Engineering, IEEE Transactions on*, vol. 51, no. 6, pp. 1026–1033, 2004.
- [5] S. S. Nudehi, R. Mukherjee, and M. Ghodoussi, "A shared-control approach to haptic interface design for minimally invasive telesurgical training," *Control Systems Technology, IEEE Transactions on*, vol. 13, no. 4, pp. 588–592, 2005.
- [6] L. Riano, C. Burbridge, and T. McGinnity, "A Study of Enhanced Robot Autonomy in Telepresence," in *Proceedings of Artificial Intelligence and Cognitive Systems, AICS*, 2011, pp. 271–283.
- [7] A. Steinfeld, T. Fong, D. Kaber, M. Lewis, J. Scholtz, A. Schultz, and M. Goodrich, "Common metrics for human-robot interaction," in *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, 2006, pp. 33–40.
- [8] A. Enes and W. Book, "Blended Shared Control of Zermelo's navigation problem," in *Proceedings of the 2010 American Control Conference*, 2010, pp. 4307–4312.
- [9] M. Lewis, "Human interaction with multiple remote robots," *Reviews of Human Factors and Ergonomics*, vol. 9, no. 1, pp. 131–174, 2013.
- [10] L. Montesano, M. Diaz, S. Bhaskar, and J. Minguez, "Towards an intelligent wheelchair system for users with cerebral palsy," *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, vol. 18, no. 2, pp. 193–202, 2010.
- [11] P. Squire, G. Trafton, and R. Parasuraman, "Human control of multiple unmanned vehicles: effects of interface type on execution and task switching times," in *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, 2006, pp. 26–32.
- [12] J. Kofman, X. Wu, T. J. Luu, and S. Verma, "Teleoperation of a robot manipulator using a vision-based human-robot interface," *Industrial Electronics, IEEE Transactions on*, vol. 52, no. 5, pp. 1206–1219, 2005.
- [13] B. Sellner, F. W. Heger, L. M. Hiatt, R. Simmons, and S. Singh, "Coordinated multiagent teams and sliding autonomy for large-scale assembly," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1425–1444, 2006.
- [14] M. L. Cummings and P. J. Mitchell, "Predicting controller capacity in supervisory control of multiple UAVs," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 38, no. 2, pp. 451–460, 2008.
- [15] M. L. Cummings and P. Mitchell, "Automated scheduling decision support for supervisory control of multiple UAVs," *Journal of Aerospace Computing, Information, and Communication*, vol. 3, no. 6, pp. 294–308, 2006.
- [16] D. F. Glas, T. Kanda, H. Ishiguro, and N. Hagita, "Teleoperation of multiple social robots," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 42, no. 3, pp. 530–544, 2012.
- [17] T. B. Sheridan, "Adaptive automation, level of automation, allocation authority, supervisory control, and adaptive control: distinctions and modes of adaptation," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 41, no. 4, pp. 662–667, 2011.
- [18] S. G. Hart and L. E. Staveland, "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research," *Advances in psychology*, vol. 52, pp. 139–183, 1988.