The 2016 AAAI Fall Symposium Series:
Shared Autonomy in Research and Practice
Technical Report FS-16-05

Human-Initiative Variable Autonomy: an Experimental Analysis of the Interactions Between a Human Operator and a Remotely Operated Mobile Robot Which Also Possesses Autonomous Capabilities

Manolis Chiou, Goda Bieksaite, Nick Hawes
School of Computer Science
University of Birmingham, UK
{e.chiou, n.a.hawes}@cs.bham.ac.uk

Rustam Stolkin
Department of Mechanical Engineering
University of Birmingham, UK
r.stolkin@cs.bham.ac.uk

Abstract

This paper presents an experimental analysis of the Human-Robot Interaction (HRI) between human operators and a Human-Initiative (HI) variable-autonomy mobile robot during navigation tasks. In our HI system the human operator is able to switch the Level of Autonomy (LOA) on-the-fly between teleoperation (joystick control) and autonomous control (robot navigates autonomously towards waypoints selected by the human). We present statistically-validated results on: the preferred LOA of human operators; the amount of time spent in each LOA; the frequency of human-initiated LOA switches; and human perceptions of task difficulty. We also investigate the correlation between these variables; their correlation with performance in the primary task (navigation of the robot); and their correlation with performance in a secondary task, in which humans are required to perform mental rotations of 3D objects, while simultaneously trying to continue with the primary task of driving the robot.

Introduction

Some robotic systems are now able to operate autonomously and robustly for long periods of time. However, predominantly, such autonomy of the robot is still only possible for repetitive tasks in comparatively simple environments, e.g. offices and warehouses. In contrast, robots that are deployed for safety- and time-critical applications, such as emergency Search And Rescue (SAR), Explosive Ordnance Disposal (EOD), or hazardous environment inspection (e.g. for nuclear decommissioning) are not able to demonstrate the same level of self-sufficiency. Such robots are still predominantly directly teleoperated by one or more human operators. The unpredictable nature of the environments encountered in such applications, coupled with the complexity of the required tasks, e.g. critical decision making or communication with victims in SAR (Dole et al. 2015), mean that the current state-of-the-art has still not progressed beyond a human-in-the-loop paradigm (Murphy 2004).

Several field studies regarding use of robots in these domains, e.g. at the 9/11 disaster site (Casper and Murphy 2003) or during the DARPA robotics challenge (Yanco et al. 2015), suggest that these applications could benefit from robots that actively assist their human operators. Ideally, what is required is a human-robot team system that dynamically benefits from the different and complementary strengths of both human and (AI-controlled) robot agents, while also compensating for the weaknesses of each agent.

"Variable autonomy" offers a solution for blending the capabilities of humans and AIs for controlling robots. A variable autonomy system is one in which control can be traded between the human operator and the robot by switching between different Levels of Autonomy (LOAs) (Chiou et al. 2016). Issues of variable autonomy robotics comprise two different but highly coupled elements: robotics engineering and computation; and Human-Robot Interaction (HRI). The first element addresses autonomous control capabilities, and how are they integrated into a robot system. The second element refers to all aspects of the human operator interacting and cooperating with the variable autonomy robot as part of a team. It includes factors such as trust in autonomous control technologies, the operator's personal preferences, and the operator's use of the robot's autonomous abilities.

Our previous work (Chiou et al. 2016) described an experimental study to investigate these two contrasting but complementary elements of robot control. Human operators used a *Human-Initiative* (HI) variable autonomy controller to navigate a simple mobile robot vehicle through a maze-like test arena, with situational awareness (SA) provided solely by a monitor-displayed control interface, featuring a view through a robot-mounted camera, and also a 2D map created by the robot's onboard laser range-finder and SLAM algorithms. The HI controller allowed operators to switch between two different LOAs on the fly: an *autonomy* LOA, in which the robot navigated autonomously towards navigational goals instructed by the human operator by mouse-clicking on a 2D map generated by the robot's SLAM system; and a *teleoperation* LOA, in which the human operator manually controlled the robot by using a joystick, while viewing the scene around the remote robot via images transmitted from a robot-mounted camera. Participants in this experiment were tasked with maximizing overall human-robot team performance, by overcoming two performance degrading factors (periodic degradation of the robot's sensors by large amounts of additive noise, and periodic degradation of the human's ability to steer the robot by forcing the human to engage in a distracting secondary task).

In (Chiou et al. 2016) we provided an analysis focused on overall system performance with respect to the primary task of robot navigation through a complex environment. The results presented (for the first time - to the best of our knowledge) principled, repeatable, scientific evidence proving that variable autonomy can outperform both pure teleoperation and pure autonomy in various circumstances.

This paper provides an additional analysis of the experimental data which were gathered during the experiments of (Chiou et al. 2016). In contrast to our earlier work (Chiou et al. 2016), this paper is specifically focused on the HRI aspects of the human operator interacting with and exploiting the HI controller. More specifically this paper reports on: the human operators' preferred LOA; the frequency of human-initiated LOA switches; human perceptions of the difficultly of using HI variable autonomy systems; and the correlation between these factors and overall task performance.

Related Work

Previous research which investigates dynamic LOA switching for mobile robots is (perhaps surprisingly) very limited (Chiou et al. 2015). Furthermore, very little previous literature has attempted to rigorously evaluate variable autonomy systems which are able to switch LOA on-the-fly (Chiou et al. 2016). Consequently, human interaction with a variable autonomy system remains predominantly unexplored in the prior literature. Studies which address similar applications to ours, e.g. SAR, have evaluated how operators interact with user interfaces (Yanco, Drury, and Scholtz 2004; Baker et al. 2004). Other studies explored the human operator's interaction with the robot in order to exchange information (Fong, Thorpe, and Baur 2003), but did not explore issues of control. Other studies investigated the human operator's interaction with a robotic system, but were restricted to exploring a single LOA (Bruemmer et al. 2005), and did not explore the issues of variable LOA. Interesting studies of robotic wheelchairs, which exploited autonomous navigation capabilities by using a shared control (mixed initiative) architecture, measured the interaction of the operator with the collaborative control system based on joystick activity (Carlson and Demiris 2008). In contrast to the above literature, this paper specifically investigates issues of the interaction of a human operator with a variable autonomy (multiple LOA) system, specifically the use by human operators of LOA switching capabilities.

In (Baker and Yanco 2004) a system was presented which aids the operator's judgment by automatically suggesting potential changes in the LOA. However, unlike our work, no data were presented on the operator's interaction with this LOA switching controller, because the system was not validated experimentally. In (Shen et al. 2004), a robot was presented which could navigate autonomously to way-points specified by a human operator. This paper suggested that the performance of such robots might be improved, by enabling a human operator to teleoperatively intervene in situations such as navigating narrow corridors, where the authors anecdotally reported difficulties with autonomous navigation. However, performance and HRI of this system were not experimentally validated in (Shen et al. 2004). Marble et al. (Marble et al. 2004) conducted a SAR-inspired experiment in which participants were instructed to switch LOA in order to improve navigation and search task performance. In contrast to our work, (Marble et al. 2004) did not investigate the human operator's interaction with the a robotic system in which LOA levels can be dynamically switched. Furthermore, unlike our work, the use of a secondary tasks in (Marble et al. 2004) were opportunistic in nature, since participants were only instructed to perform such tasks optionally. Thus, in contrast to our own work, the secondary tasks of (Marble et al. 2004) were not designed to degrade human performance on the primary task (driving the robot). Additionally, (Marble et al. 2004) did not incorporate any methods into their experiments for degrading the robot's autonomous performance in a controlled way.

To the best of our knowledge, our work is the first that exploits rigorous methodologies from psychology and human factors research to carry out a principled study of variable autonomy in mobile robots. Moreover, this paper is the first work on mobile robots that reports a principled analysis of the ways in which human operators interact with, and exploit the capabilities of, a robotic system in which LOA modes can be dynamically switched.

Robotic Apparatus, Interface, and Environment

For our experiment we used a simulated robot and environment in the Modular Open Robots Simulation Engine (MORSE) (Echeverria et al. 2011), a high fidelity simulator. The robot used was a Pioneer-3DX mobile robot equipped with a laser range finder sensor and a RGB camera. The robot was controlled by the Operator Control Unit (OCU), comprising a laptop, a joystick, a mouse and a screen showing the control interface (see Figure 1).

Our previous work (Chiou et al. 2015) showed that conducting principled and repeatable real-world experiments on variable autonomy robots, with human test-subjects, is extremely difficult, with a variety of confounding factors that can greatly affect the repeatability of experiments and the validity of collected data. Furthermore, it is often impractical to build large-scale robot arenas. In contrast, using a high fidelity simulated robot and test-arena enabled us to carefully control our experiments (e.g. lighting conditions and other factors) while also enabling a large robot arena which would have been difficult to physically construct given the space-limitations of most modern universities.

Our system offers two LOAs: "teleoperation" in which the human operator drives the robot with the joystick, while gaining situational awareness (SA) via a video feed from the robot's onboard RGB camera, and from a laser-generated 2D map showing the robot's pose as derived from its onboard SLAM system; "autonomy" in which the human operator mouse-clicks a desired location on the laser-generated 2D map, and then the robot autonomously plans and executes a trajectory to that location, while automatically avoiding obstacles. The system is a "Human-Initiative" (HI) system, in the sense that the human operator uses his or her initiative to switch between the two possible LOAs at any time by

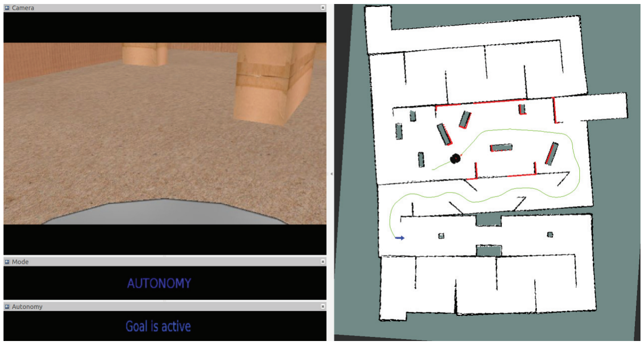

Figure 1: The control interface as presented to the operator. **Left**: video feed from the robot's RGB camera; the control mode currently being used (autonomy or teleoperation) and the status of the robot (e.g. "Goal is active"). **Right**: 2D map, generated by laser-based SLAM, showing the position of the robot, the current human-selected goal (arrow), the AI planned robot trajectory (line), the obstacles' laser reflections and the walls of the arena.

pressing a joystick button.

Experimental Design and Procedure

This paper presents new HRI findings, derived from further analysis of data gathered during the experiment presented in (Chiou et al. 2016). As such, the reader is encouraged to read the aforementioned paper for a more detailed description of the experimental design.

In summary, the experiment investigates to what extent circumstances in which the human-robot team is underperforming (i.e. low system performance on a task), can be overcome or improved by switching control between the AI and the human operator. Such circumstances may include the robot being neglected while its operator performs a secondary task (Murphy 2004), e.g. being asked to convey SA information to team mates (Murphy and Burke 2005; Burke et al. 2004), or the robot becoming stuck due to a navigation failure in autonomy mode.

Experimental Setup - Test Arena, Tasks and Tested Conditions

The experiment took place in a simulated maze with dimensions of 11×13.5 meters (see Figure 2). Operators were given the primary task of navigating from point A in Figure 2(a) (the beginning of the arena) to point B (the end of the arena) and then back again to point A.

During each experimental trial, two different kinds of performance degrading factors were introduced, one for each agent. At certain times, artificially generated sensor noise was used to degrade the performance of autonomous navigation. At other times, a cognitively intensive secondary task was used to degrade the performance of the human operator. Each of these performance degrading situations occurred twice per experimental trial, once on the way from point A to point B, and once on the way back from B to A. These

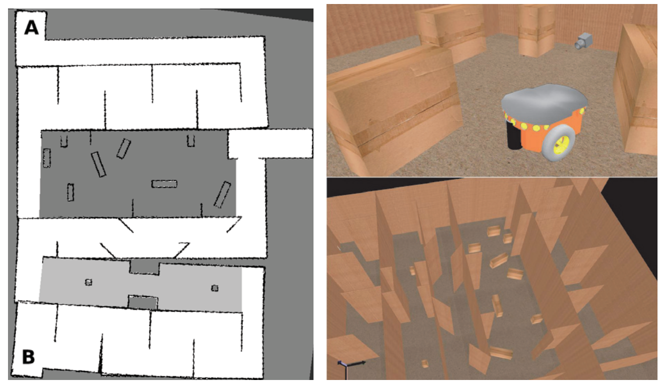

(a) (b)

Figure 2: 2(a): laser-derived SLAM map created in the simulation environment. The primary task was to drive from point A to point B and then back again to point A. The light grey shaded region is where artificial sensor noise was introduced. The dark grey shaded region is where the secondary task was presented to the operator. 2(b): the simulated arena and the robot model used in the experiment.

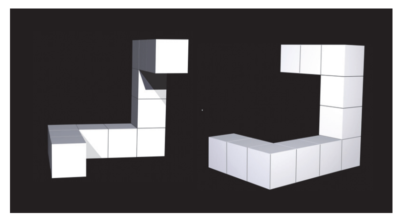

Figure 3: Example of one of the rotated 3D objects cards, which were used for the secondary task.

degrading factors occurred separately from each other, as shown in Figure 2(a).

More specifically, autonomous navigation was degraded by adding Gaussian noise to the laser scanner range measurements, thereby degrading the robot's localization and obstacle avoidance abilities. To ensure experimental repeatability, this additional noise was instantiated whenever the robot entered a predefined area of the arena, and was deactivated when the robot exited that area. Note that such region-specific noise can in fact happen in real-world applications, e.g. when a robot travels through a highly radioactive region during nuclear decommissioning or the exploration and remediation of nuclear disaster sites such as Fukushima.

To degrade the performance of the human, a secondary task of mentally rotating 3D objects was used (Ganis and Kievit 2015). Whenever the robot entered a predefined area in the arena, the operator was presented with a series of 10 cards, each showing images of two 3D objects, see Figure 3. On five cards, the objects were identical but rotated, and on the other five cards the objects were mirror images with opposite chiralities. The operator was required to state (yes or no) whether or not the two objects were identical.

For each human test-subject, three different control modes were tested. In teleoperation mode, the human operator was restricted to using only direct joystick control to drive the robot. Conversely, in autonomy mode, the operator was only allowed to guide the robot by clicking desired destinations on the robot's laser-generated 2D map. In Human-Initiative (HI) mode, the operator was allowed to switch LOA at any time (using a button on the joy-pad) according to their judgment, in order to maximize performance.

Participants and Procedure

A total of 24 test subjects participated in a within-groups experimental design (i.e. every human test-subject performed all three trials). Usable data were obtained from 23 participants. A prior experience questionnaire revealed that the majority of the participants were experienced in driving, playing video games or operating mobile robots. Each participant underwent extensive and standardized training before the experiment. This ensured that all participants had adequate understanding of the system and had attained a common minimum skill level. Counterbalancing was used in the experimental trials (i.e. the order of the tested conditions was rotated differently for different participants) in order to prevent both learning and fatigue effects from biasing the results. Additionally, all participants were required to perform the secondary task separately, prior to the experiment (i.e. while not driving the robot) in order to establish baseline performance on the secondary task.

Participants were asked to perform the primary task (robot navigation) as quickly and as possible while minimising collisions. Participants were also instructed that, when presented with the secondary task, they should do it as quickly and as accurately as possible. They were explicitly told to prioritise the secondary task over the primary task, and only to perform the primary task if the workload allowed. In contrast, our previous work (Chiou et al. 2015) showed that serious confounding factors are introduced if participants are not clearly told which task to prioritise.

The human operators could only acquire SA information via the OCU, which displays real-time video feed from the robot's front-facing camera, and the estimated robot location on the 2D SLAM map. All participants were given an identical and complete 2D map, generated by SLAM prior to the trials. In contrast, our previous work (Chiou et al. 2015) showed that serious confounding factors are introduced if each participants are allowed to independently create their own SLAM map by explorative navigation of the robot during the experiment. Such confounding factors can only be eliminated by constraining as many experimental variables as possible, even if this means that certain experiences within the experiment are somewhat different from those encountered e.g. in a real emergency robot deployment, where SLAM maps must obviously be created exploratively on-the-fly.

Participants completed a NASA Task Load Index (TLX) questionnaire (Sharek 2011) at the end of each experimental run. NASA-TLX is a questionnaire rubric system which enables the perceived workload and difficulty of using a system to be numerically quantified.

Although a variety of data were collected (primary task completion time; total number of collisions; secondary task completion time; number of secondary task errors), in this paper we focus on HRI related data from the HI trials.

Results

Firstly, we summarize the main results from (Chiou et al. 2016) for completeness. In *Primary task completion time* (ANOVA *F(1.275, 28.057) = 34.567,* $p < .01$*,* $power > .9$*,* $\eta^2 = .61$, HI ($M = 413.6\ sec$) performed significantly ($p < .05$) better than teleoperation ($M = 429.6\ sec$) and pure autonomy ($M = 483.9\ sec$).

Secondary task completion time refers to the average time per trial (in seconds), that the participants took to complete one series of the 3D object cards. ANOVA with $F(1.565, 34.420) = 7.821, p < .01, power > .85$, $\eta^2 = .26$, and pairwise comparisons suggests that participants performed significantly ($p < .05$) better in the baseline trial ($M = 33.2$) compared to their performance during robot operation; HI mode ($M = 39.3$); autonomy ($M = 39.5$); and teleoperation ($M = 41.7$).

In *NASA-TLX* scores, ANOVA ($F(2, 44) = 11.510, p < .01, power > .9$, $\eta^2 = .34$) and pairwise comparisons showed that autonomy ($M = 35.2$) was perceived by participants as having the lowest difficulty, as compared to HI ($M = 41.4$) with $p < .05$ and teleoperation ($M = 47.8$) with $p < .01$.

Here, we present data and analyses regarding the ways in which human operators interacted with the dynamic LOA switching capabilities of the HI system. The percentage of time spent in each LOA (i.e. teleoperation or autonomy) during the HI trials was measured. We report mostly on the time spent in the autonomy LOA, because: firstly, it was the dominant LOA chosen by human test-subjects during the HI trials; secondly, everything that is correlated (or not) with percentage of time in autonomy is also inversely correlated (or not) with percentage of time in teleoperation LOA. This is due to the fact that these two modes combine together to form almost 100% of the time of each trial. The remaining small percentage of time corresponded to a "stop mode" which allows operators to perform emergency canceling of navigational goals and robot movement.

The average percentage of time spent in teleoperation mode was $M = 37.17\%$) and the average time spent in autonomy was $M = 62.7\%$, see Figure 4(a). The standard deviation, across trials, on these percentages was $std = 23.7$. The remaining 0.13% of time was spent in stop mode. As can be seen in the histogram (see Figure 4(b)) the majority of the human test-subjects spent more than 60% of the time using autonomy. Additionally there were two other smaller groups of operators based on Figure 4(b). One group equally split their time between autonomy and teleoperation, while the third group mostly (i.e. more than 60% of the time) chose to use the teleoperation mode.

The number of LOA switches operators performed in each trial denotes the frequency in which they make use of the HI controller capabilities. The mean number of LOA switches per trial was $M = 9.78$ with a standard deviation of $std = 7.367$. As can be seen from the histogram of Figure 5(a), the vast majority of the operators (more than 74%)

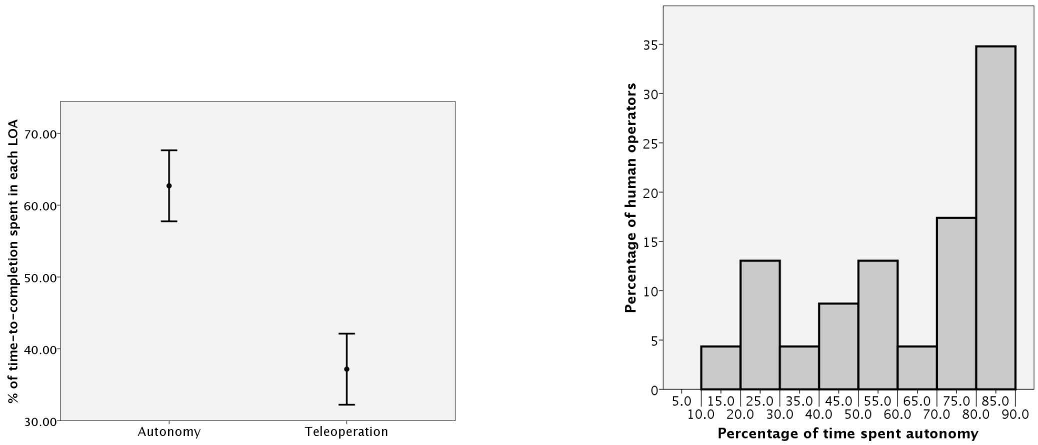

(a) (b)

Figure 4: 4(a): Percentage of time-to-completion spent in each of the two LOAs during the HI trials. Error bars indicate the standard error. 4(b): Histogram shows the proportion of human operators who spent various different proportions of their time in the autonomy LOA during HI experiments.

changed LOA less than 11 times. A much smaller number of operators choose to switch control very often (more than 16 times per trial).

Correlation analysis using a two-tailed Pearson's r was conducted to investigate any relationships between human-robot team performance and other variables. Firstly the *number of LOA switches* and the *percentage of time spent in autonomy* were not correlated $r(21) = .24, p > .05$. No correlation was found between the *number of LOA switches* and performance of the system in terms of *primary task time-to-completion* (duration of navigation from A to B and back), $r(21) = -.43, p > .05$. Also there was no correlation between the *number of LOA switches* and performance in the *secondary task completion time*, $r(21) = .016, p > .05$. The *percentage of time spent in autonomy* and performance in the *primary task time-to-completion*, were not correlated $r(21) = .012, p > .05$. Lastly, the *percentage of time spent in autonomy* was not found to be correlated with the *secondary task completion time*, $r(21) = .16, p > .05$.

Correlation analysis showed that *NASA-TLX scores* were inversely and significantly correlated with the *percentage of time spent in autonomy* during the HI trials, $r(21) = -.446, p < .05$. The greater the proportion of time spent in autonomy mode, the easier the task was perceived to be by the human test-subjects. Analysis found no correlation between the *number of LOA switches* and *NASA-TLX scores*, $r(21) = -.016, p < .05$.

Discussion

The analysis showed that the majority of the participants chose to use mostly autonomy mode. Furthermore, most participants switched LOA 9.78 times on average. The fact that there is no correlation between either of those two factors and performance in primary or secondary tasks, suggests that operators used LOA switching for reasons that are not purely related to task performance. However, our previous work (Chiou et al. 2016) did clearly demonstrate that HI variable autonomy can outperform pure autonomy or pure teleoperation in various situations. Thus, it can be inferred that part of the time spent in autonomy, and a number of the observed LOA switches, were crucial in improving overall task performance. This in turn, suggests that all participants attained a minimum skill level to successfully exploit the HI LOA switching capabilities, further reinforcing the findings in (Chiou et al. 2016).

The remaining use of autonomy and number of LAO switches (i.e. those in excess of what were needed to achieve good performance), may simply reflect personal preferences. These individual preferences can be driven by several factors. The level of trust in the autonomous control can lead operators to use more or less autonomy. For example, a test-subject who does not trust the robot's autonomous capabilities, may choose rely on direct teleoperation more than is necessary. Since, in our experiments, the autonomy mode was highly used by most test-subjects, this suggests that human operators did indeed trust that the robot's autonomous navigation AI will perform at least as well as a human teleoperator in most situations.

Operators' personality traits may also play an important role. Possibly some humans prefer to be more in control of a situation or show a more hands-on attitude, and are therefore likely to use more teleoperation. Others may prefer the role of a supervisor in conjunction with a more laid back attitude. Those individuals are likely to use more autonomy. Additionally, note that a number of LOA switches can be traced to the general alertness of human operators, triggered by their anticipation of events. For example, some operators may switch preemptively to autonomy while anticipating the appearance of a secondary task. Other operators may switch preemptively to teleoperation if they anticipate the robot getting stuck in an awkward situation in the test arena. The latter was observed by the experimenters in a number of par-

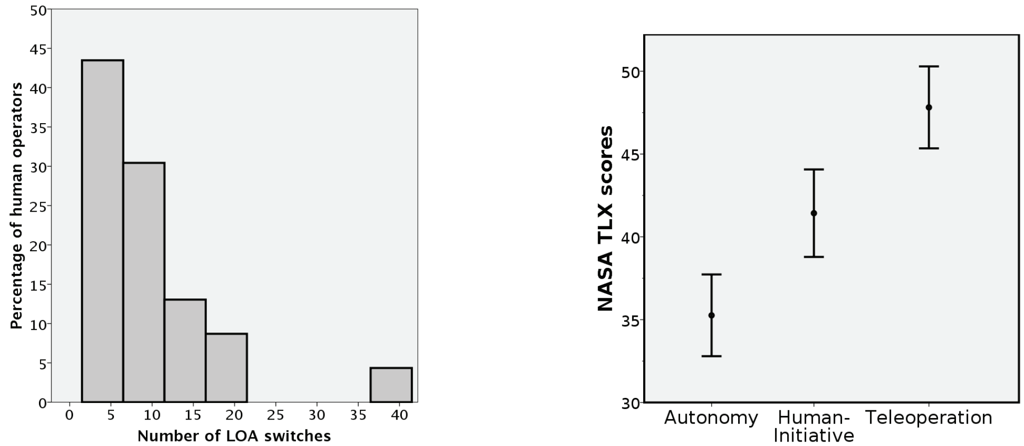

(a) (b)

Figure 5: 5(a): Histogram showing the percentages of human operators who chose to make various different numbers of LOA switches during HI trials. 5(b): NASA-TLX score showing the overall trial difficulty as perceived by the human operators, for three different modes of operation (pure autonomy, pure teleoperation, and HI variable autonomy.

ticipants, who would switch momentarily to teleoperation in parts of the arena where they felt the robot was performing sub-optimally, e.g. while navigating around sharp corners. These participants tend to switch LOA more frequently and thus fall into the group of operators with a very high number of LOA switches, see Figure 5(a).

Correlation between NASA-TLX scores and the percentage of time spent in autonomy, shows that participants who mostly used the autonomy mode perceived the overall task and operational difficulty to be easier than those who used mostly teleoperation. Firstly, this further validates the NASA-TLX results in (Chiou et al. 2016) that found autonomy to be easier than teleoperation. However, it is still not clear if the causal reason that operators preferentially chose autonomy more than teleoperation, is because they perceived it as an easier and more effective control mode. Further investigation is needed to confirm cause and effect. Secondly, it is a possible indication of the extra cognitive overhead that switching LOA based on judgment may impose on the operator. In particular while teleoperating or while occupied with a secondary task, the operator can find himself in an overwhelming situation. Thus, he/she may find it difficult to judge which is the appropriate LOA for the situation or he/she may be too overwhelmed to actually perform the LAO-switching action (i.e. making a cognitive decision to switch, followed by executing a button press). This problem of cognitive overloading of the human operator, can be tackled by implementing a Mixed-Initiative (MI) robotic system in which both the AI and the human operator have the authority to initiate LOA changes. AI can assist by taking control from the operator or switching LOA while he/she is occupied, thus alleviating him/her from the burden of manual control. We have already implemented such an MI system, and it will form the subject of our next paper.

Conclusion

This paper presented a principled and statistically validated empirical analysis of human operators interactions with a Human-Initiative variable autonomy robot. More specifically, our work was focused on the use of the dynamic mode-switching capabilities of the system by the operator. We believe this is the first study which has quantitatively reported on metrics such as percentage of time spent in each LOA, frequency of LOA switches, and perceived workload and difficulty in regard to such metrics.

The findings reported in this paper suggest that operators choose to mostly use the autonomy control mode rather than teleoperation. The fact that the amount of time spent in autonomy mode was not correlated with overall performance, possibly indicates that operators used autonomy for reasons beyond considerations of only performance. However, the possibility of operators overestimating AI abilities while underestimating their own human abilities, cannot yet be excluded as a possible explanation for this lack of correlation. Other explanations may include the fact that autonomy was perceived as easier to use, as demonstrated by the high correlation between NASA-TLX and the amount of time spent in autonomy mode.

In this paper, we also presented evidence that the frequency of LOA switching does not correlate with performance. In contrast, our previous work (Chiou et al. 2016) reported strong evidence that HI variable autonomy outperforms both pure autonomy and pure teleoperation in various situations. Hence, it can be inferred that operator's switched LOA more frequently than the minimum necessary to maximizes performance.

References

Baker, M., and Yanco, H. A. 2004. Autonomy mode suggestions for improving human-robot interaction. In *IEEE*

International Conference on Systems, Man and Cybernetics (SMC), volume 3, 2948–2953.

Baker, M.; Casey, R.; Keyes, B.; and Yanco, H. A. 2004. Improved interfaces for human-robot interaction in urban search and rescue. In *IEEE International Conference on Systems, Man and Cybernetics*, volume 3, 2960–2965.

Bruemmer, D. J.; Few, D. A.; Boring, R. L.; Marble, J. L.; Walton, M. C.; and Nielsen, C. W. 2005. Shared Understanding for Collaborative Control. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 35(4):494–504.

Burke, J. L.; Murphy, R. R.; Coovert, M. D.; and Riddle, D. L. 2004. Moonlight in Miami: Field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise. *HumanComputer Interaction* 19(1-2):85–116.

Carlson, T., and Demiris, Y. 2008. Collaborative Control in Human Wheelchair Interaction Reduces the Need for Dexterity in Precise Manoeuvres. In *Robotic Helpers: User Interaction, Interfaces and Companions in Assistive and Therapy Robotics, a Workshop at ACM/IEEE HRI 2008*, number March, 59–66.

Casper, J. L., and Murphy, R. R. 2003. HumanRobot Interactions During the Robot-Assisted Urban Search and Rescue Response at the World Trade Center. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 33(3):367–385.

Chiou, M.; Hawes, N.; Stolkin, R.; Shapiro, K. L.; Kerlin, J. R.; and Clouter, A. 2015. Towards the Principled Study of Variable Autonomy in Mobile Robots. In *Systems, Man, and Cybernetics (SMC), 2015 IEEE International Conference on*, 1053–1059.

Chiou, M.; Stolkin, R.; Bieksaite, G.; Hawes, N.; Shapiro, K. L.; and Harrison, T. S. 2016. Experimental Analysis of a Variable Autonomy Framework for Controlling a Remotely Operating Mobile Robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

Dole, L. D.; Sirkin, D. M.; Murphy, R. R.; and Nass, C. I. 2015. Robots Need Humans in the Loop to Improve the Hopefulness of Disaster Survivors. In *IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN)*, 707–714.

Echeverria, G.; Lassabe, N.; Degroote, A.; and Lemaignan, S. 2011. Modular open robots simulation engine: MORSE. In *Proceedings - IEEE International Conference on Robotics and Automation*, 46–51.

Fong, T.; Thorpe, C.; and Baur, C. 2003. Robot, asker of questions. *Robotics and Autonomous Systems* 42(3-4):235–243.

Ganis, G., and Kievit, R. 2015. A New Set of Three-Dimensional Shapes for Investigating Mental Rotation Processes : Validation Data and Stimulus Set. *Journal of Open Psychology Data* 3(1).

Marble, J. L.; Bruemmer, D. J.; Few, D. A.; and Dudenhoeffer, D. D. 2004. Evaluation of supervisory vs. peer-peer interaction with human-robot teams. In *Proceedings of the 37th Hawaii International Conference on System Sciences*, volume 00, 1–9.

Murphy, R. R., and Burke, J. L. 2005. Up from the Rubble: Lessons Learned about HRI from Search and Rescue. *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting* 49(3):437–441.

Murphy, R. R. 2004. Human-Robot Interaction in Rescue Robotics. *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)* 34(2):138–153.

Sharek, D. 2011. A Useable, Online NASA-TLX Tool. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 55(1):1375–1379.

Shen, J.; Ibanez-Guzman, J.; Ng, T. C.; and Seng Chew, B. 2004. A collaborative-shared control system with safe obstacle avoidance capability. In *IEEE Conference on Robotics, Automation and Mechatronics*, volume 1, 119–123.

Yanco, H. A.; Norton, A.; Ober, W.; Shane, D.; Skinner, A.; and Vice, J. 2015. Analysis of Human-robot Interaction at the DARPA Robotics Challenge Trials. *Journal of Field Robotics* 32(3):420–444.

Yanco, H.; Drury, J.; and Scholtz, J. 2004. Beyond Usability Evaluation: Analysis of Human-Robot Interaction at a Major Robotics Competition. *Human-Computer Interaction* 19(1):117–149.