Experimental Analysis of a Variable Autonomy Framework for Controlling a Remotely Operating Mobile Robot

Manolis Chiou[1], Rustam Stolkin[2], Goda Bieksaite[1], Nick Hawes[1], Kimron L. Shapiro[3], Timothy S. Harrison[4]

***Abstract*—**

I. INTRODUCTION

Despite the significant advances in autonomous robotics in recent years, real-world robots deployed in high consequence and hazardous environments remain predominantly teleoperated, using interfaces that have not changed greatly in over 30 years. Examples of such applications include Explosive Ordnance Disposal (EOD), Search and Rescue (SAR) and nuclear decommissioning (e.g. robots deployed at Fukushima or at UK and US legacy nuclear sites). The reasons for continued reliance on direct teleoperation are that autonomous methods are still not robust enough to be completely self-sufficient in highly unstructured and uncertain environments.

On the other hand, several Human-Robot Interaction (HRI) field studies [1]–[3] in SAR operations identify the necessity for more autonomy to be used in such robots. Often the remote robot will be separated from its human operator by e.g. thick concrete walls (nuclear scenarios), or rubble (SAR scenarios), severely limiting communication bandwidth in situations where umbilical tethers can cause entanglement and other severe problems. Additionally, controlling a remote robot to perform precise movements with respect to surrounding objects can be extremely difficult for human operators who only have limited situational awareness (SA) (e.g. restricted views and poor depth perception using a robot-mounted camera).

It seems likely that future robot applications will therefore require some form of *variable autonomy* control. 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), such that agents can assist each other. Such a system offers the potential to assist a human who may be struggling to cope with issues such as high workload, intermittent communications or operator multi-tasking. For example, a human operator might need to concentrate on a secondary task while temporarily devolving control to an AI which can autonomously manage robot navigation.

The use of different LOAs in order to improve system performance is a challenging and open problem, raising a number of difficult questions. For example: which LOA should be used under which conditions?; what is the best way to switch between different LOAs?; and how can we investigate the trade-offs offered by switching LOAs in a repeatable manner? These questions need to be explored by conducting experiments within a rigorous multidisciplinary framework, drawing on methodologies from the fields of psychology and human factors, as well as engineering and computer science. Our previous work [4] highlighted the absence of such a framework in existing literature. Additionally it demonstrated the intrinsic complexity of conducting such experiments due to the high number of confounding factors and large variances in the results.

This paper develops from our previous work by designing and carrying out a principled experimental study to empirically evaluate the performance of a human-robot team when using a variable autonomy controller. More specifically it improves the experimental framework by: a) minimizing confounding factors, e.g. by using extensive participant training and a within-subject design; b) introducing a meaningful secondary task for human operators; and c) introducing a variable autonomy controller. We present formally-analysed, statistically-evaluated experimental evidence, to support the hypothesis that a variable autonomy system can indeed outperform teleoperated or autonomous systems in various circumstances.

In our experiments, we compare the performance of three different systems: 1) pure joystick teleoperation of a mobile robot; 2) a semi-autonomous control mode (which we refer to hereafter as the "autonomy" LOA) in which a human operator specifies navigation goals to which the robot navigates autonomously; 3) a Human-Initiative (HI) variable autonomy system, in which the human operator can dynamically switch between the teleoperation and autonomy modes using a button press. During experiments, human test subjects are tasked with navigating a differential drive vehicle around a maze-like test arena, with SA provided solely by a monitor-displayed control interface. At various points during the experiments, the robot's performance is degraded by artificially introducing controlled amounts of noise to sensor readings, and the human operator's performance is degraded by forcing them to perform a cognitively complex secondary task.

The experiments reported in this paper focus on the ability and authority of a human operator to switch LOA on the fly, based on their own judgement. We define this form of

[1]School of Computer Science, University of Birmingham, UK.
[2]School of Engineering University of Birmingham, UK.
{e.chiou, r.stolkin, n.a.hawes}@cs.bham.ac.uk
[3]School of Psychology, University of Birmingham, UK.
[4]Defense Science and Technology Laboratory, UK.

variable autonomy as Human-Initiative (HI), in contrast to Mixed-Initiative (MI) systems in which both the AI and the operator have the authority to initiate LOA changes. However, towards the end of this paper we additionally make suggestions for how the data, results and insights gathered during these experiments could be used to inform the design of a Mixed-Initiative (MI) system in future work.

II. RELATED WORK

The majority of the robotics literature is focused on describing the engineering and/or computational details of new technologies, while comparatively few studies address the issues of rigorously evaluating how well a human can use such robots to carry out a real task. Additionally, the autonomous robotics literature has historically tended to be somewhat separated and distinct from the literature investigating the issues of teleoperation, with relatively little work specifically focusing on variable autonomy systems.

A common approach to improving teleoperated systems is to enhance the user interface [5]. A carefully designed interface can often assist the operator in performing better. However, it does not alleviate them from the burden of continuous control, nor does it exploit the complementary capabilities of the robot to manage some tasks for itself.

Research which focuses on investigating dynamic LOA switching on mobile robots is fairly limited. Furthermore, the investigation of MI systems to address this dynamic switching is even more limited, as highlighted by Jiang and Arkin [6] and Chiou et al. [4]. A large part of the literature, e.g. [7], [8], is focused on comparing the relative performance of separate LOAs, and does not report on the value of being able to switch between LOAs. In contrast, our work specifically addresses the issues of dynamically changing LOA on-the-fly (i.e. during task execution) using either a MI or HI paradigm.

Baker and Yanco [9] presented a robotic system in which the robot aids the operator's judgement by suggesting potential changes in the LOA. However, the system was not validated experimentally. Marble et al. [10] conducted a SAR-inspired experiment in which participants were instructed to switch LOA in order to improve navigation and search task performance. However, [10] was intended to be a usability study which explored the ways in which participants interacted with each different LOA. In contrast, our own work is focused on evaluating and demonstrating the overall task-performance when LOA levels can be dynamically switched. As in our own work, [10] also incorporate secondary tasks into their experiments. However, in contrast to our work, the use of these secondary tasks was opportunistic in nature because participants were only instructed to perform them optionally. Hence, the secondary tasks in [10] do not degrade human performance on the primary task (steering the robot). Also, unlike our work, [10] did not incorporate any methods into their experiments for degrading the robot's autonomous performance in a controlled way.

Much of the published experimental work does not carefully control for possible confounding factors. These factors can vary from partially uncontrolled test environments (as in [10]), up to the absence of standardized training for human test-subjects as in [8], [11], [12]. It is particularly important to control for the training and experience of human test-subjects, as these factors are known to affect overall robot operating performance [13], [14]. Additional confounding factors include the robot having different speed limits in the different conditions tested [11], or different navigation strategies of human operators [4]. In contrast to our work, Nielsen et al. [15] report no significant primary task results due to large measurement variances, but they do present a method for systematically categorizing the different navigational strategies of human operators.

All of the papers discussed above make important contributions in their own right, and we do not intend to devalue such work in any way. However, across the related literature we note a deficiency of: a) rigorous statistical analysis; b) clarity on assumptions and hypotheses; c) precise and detailed descriptions of the experimental protocol followed; d) a formalized, coherent and repeatable experimental paradigm. In contrast, in disciplines such as psychology and human factors, the above criteria constitute standard practice.

An excellent example of related work, which does provide a rigorous protocol, statistical analysis and detailed description, is the work of Carlson et al. [16]. They validate an adaptive shared control system, while degrading task performance with the use of a secondary task. However, their work is focused on the use of a Brain-Computer Interface for robot control. Because this field is relatively young, and the problems are extremely difficult, [16] used a robot navigation task which was comparatively simplified, i.e. operators only control left-right movement of a robot using a keyboard.

Lastly, variable autonomy research in the field of multiple robots being controlled by a single operator, provides similar experimental studies. However much of this research (e.g. [17], [18]) is focused on higher levels of abstraction than our work, e.g. planning or task allocation. Other experiments, e.g [19], [20], are focused on human factors issues such as gaining SA when controlling multiple robots, or how the operator interacts with as many robot as possible.

In contrast to the above work, to the best of our knowledge, our paper 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; the first mobile robot experiments that combine quantifiable and repeatable degradation factors for both human and robot; and the first work which formally and systematically evaluates the benefits of combining the capabilities of both human and autonomous control in a dynamically mode-switching system.

III. APPARATUS AND ROBOTIC SOFTWARE

Our robot and environment were simulated in the Modular Open Robots Simulation Engine (MORSE) [21], which is 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 is controlled by the Operator Control

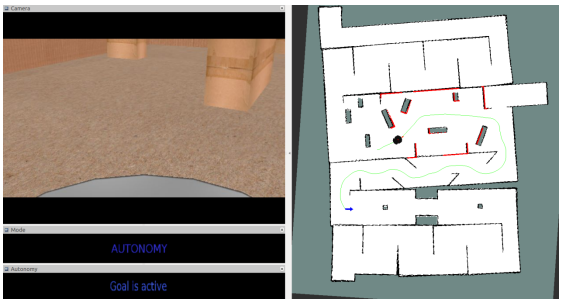

Fig. 1. The control interface as presented to the operator. **Left**: video feed from the camera, the control mode in use and the status of the robot. **Right**: The map showing the position of the robot, the current goal (blue arrow), the AI planned path (green line), the obstacles' laser reflections (red) and the walls (black).

Unit (OCU), composed of a laptop, a joystick, a mouse and a screen showing the control interface (see FIG. 1).

In our previous work [4] we built a large maze-like test arena (see FIG. 2(b) and FIG. 3(b)), and carried out human-subject tests using a real Pioneer-3DX robot fitted with camera, laser scanner and WiFi communication to the remote Operator Control Unit. While demonstrating new methods on real robots is important, we observed that this can introduce difficult confounding factors, which can detract from the repeatability of experiments and the validity of collected data. For example, tests at different times of day or different weather, mean that daylight levels inside the lab change, affecting the video images observed by each test-subject. Different amounts of battery charge can cause top speed of the robot to vary slightly between different test-subjects. These and other factors led us to design the experiments reported in this paper using a high fidelity simulated robot and test-arena. As can be seen in FIG. 2 and FIG. 3, and comparing the real and simulated video feeds (FIG. 1 and FIG. 4), the simulation environment creates very similar situations and stimuli for the human operators as experienced when driving the real robot, but with a much higher degree of repeatability.

Our system offers two LOAs. **Teleoperation:** the human operator drives the robot with the joystick, while gaining SA via a video feed from the robot's onboard RGB camera. Additionally a laser generated 2D map is displayed on the OCU. **Autonomy:** the operator clicks on a desired location on the 2D map, then the robot autonomously plans and executes a trajectory to that location, automatically avoiding obstacles. The system is a **Human-Initiative (HI)** system as the operator can switch between these LOAs at any time by pressing a joystick button. The software used was developed in Robot Operating System (ROS) and is described in more detail in [4].

IV. EXPERIMENTAL DESIGN AND PROCEDURE

This experiment investigates to what extent circumstances in which the robot is under-performing, can be overcome or improved by switching control between the AI and the human operator. Such circumstances may include idle time,

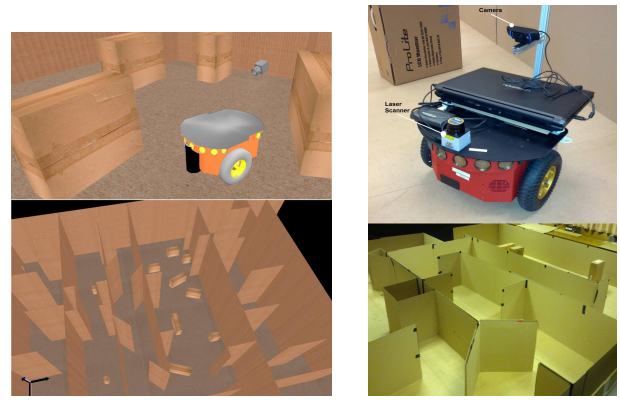

(a) (b)

Fig. 2. 2(a): the simulated arena and the robot model used in the experiment. 2(b): the real arena and robot used in our previous experiment. Note that the simulation recreates the real environment with a good degree of fidelity.

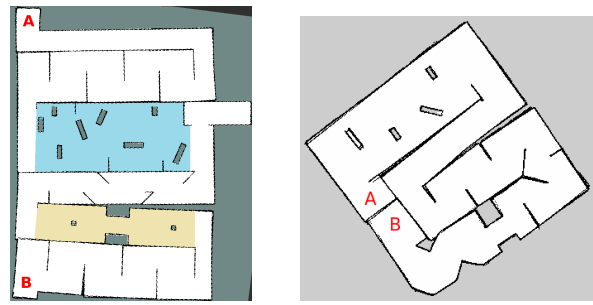

(a) (b)

Fig. 3. 3(a): laser-derived SLAM map created in the simulation environment. Primary task was to drive from point A to B and back again to A. The yellow shaded region is where artificial sensor noise was introduced. The blue shaded region is where the secondary task was presented to the operator. 3(b): laser-derived SLAM map generated by real robot in our previous experiment. Note the similarities between the real and simulated data.

which is the time passed without any progress towards achieving a goal [4]. For example a robot being neglected by its operator when in teleoperation mode, or stuck due to a navigation failure in autonomy mode. Similar situations are quite common in real world robotics deployments [22]. For example, consider the case in which a robot operator must interrupt their control of the robot, to provide information to the SAR team leader or EOD team commander. Our hypothesis is that in such circumstances, trading control to another agent will improve the overall task performance of the system.

A. *Experimental setup - operator control unit and robot test arena*

In the work described in this paper, we used an identical OCU (see FIG. 4(b)) as that used in our previous experiments with a real robot [4]. A simulated maze was designed with dimensions of 11×13.5 meters (see FIG. 2(a) and FIG. 3(a)). It approximates a yellow coded National Institute of Standards and Technology arena [23]. As can be seen in FIG.

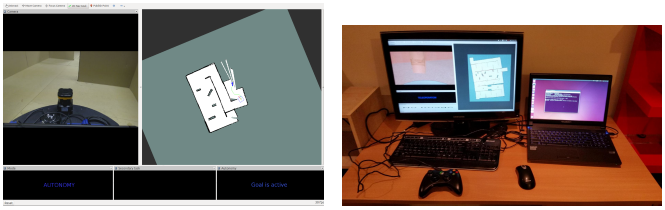

(a) (b)

Fig. 4. 4(a): the control interface as presented to the operator in our previous real world experiment. 4(b): the Operator Control Unit (OCU), composed of a laptop, a joystick, a mouse and a screen showing the control interface. The same OCU was used in both experiments.

3(b) and FIG 4(a), the data presented to the human operator via the OCU is almost identical to that experienced by human test subjects operating the real robot in a real arena in our prior work.

B. Primary and secondary tasks, and experimental test modalities

Each human test subject was given the primary task of navigating from point A in FIG. 3(a) (the beginning of the arena) to point B (the end of the arena) and back to point A. The path was restricted and one way, i.e. no alternative paths existed.

Two different kinds of performance degrading factors were introduced, one for each agent: artificially generated sensor noise was used to degrade the performance of autonomous navigation; and a cognitively intensive secondary task was used to degrade the performance of the human test subject. In each experimental trial, each of these performance degrading situations occurred twice, once on the way from point A to point B, and a second time on the way from point B back to point A. The two different kinds of degradations occurred separately from each other, as shown in FIG. 3(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. For every experimental trial this additional noise was instantiated when the robot entered a pre-defined area of the arena, and was deactivated when the robot exited that area.

To degrade the performance of the human operator, their cognitive workload was increased via a secondary task of mentally rotating 3D objects. Whenever the robot entered a predefined area in the arena, the test subject was presented with a series of 10 cards, each showing images of two 3D objects (see FIG. 5). In half of the cards, the objects were identical but rotated by 150 degrees. In the other half the objects were mirror image objects with opposite chiralities. The test subject was required to verbally state whether or not the two objects were identical (i.e. yes or no). This set of 3D objects was previously validated for mental rotation tasks in [24].

For each human test subject, three different control modes were tested. In *teleoperation* mode, the operator was restricted to using only direct joystick control to steer the robot,

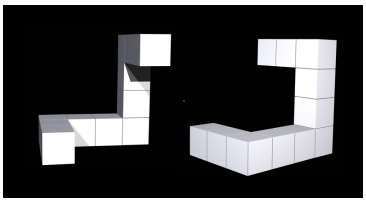

Fig. 5. A typical example of a rotated 3D objects card.

and no use of the robot's autonomous navigation capabilities was allowed at any time. In *autonomy* mode, the operator was only allowed to guide the robot by clicking desired destinations on the 2D map. The only exception was in the case of critical incidents such as the robot becoming stuck in a corner. Under such circumstances the experimenter would instruct the human operator to briefly revert to joystick control in order to free the robot so that the experiment could continue. In *Human-Initiative (HI)* mode, the operator was given freedom to switch LOA at any time (using a push-button on the joy-pad) according to their judgement, in order to maximize performance.

C. Participants and procedure

A total of 24 test subjects participated in a within-groups experimental design (i.e. every test subject performed all three trials), with usable data 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 test subject underwent extensive training before the experiment. This ensured that all participants had attained a common minimum skill level (which otherwise might lead to a confounding factor in later data analysis). Participants were not allowed to proceed with the experimental trials until they had first demonstrated that they could complete a training obstacle course three times, within a specific time limit, with no collisions and while presented with the two degrading factors (i.e. the secondary task and sensor noise). Each of the three training trials used a different control mode. Additionally, all participants were required to perform the secondary task separately (i.e. without driving the robot) in order to establish baseline performance.

During the actual experimental trials (testing the three different control modes), counterbalancing was used, i.e. the order of the three control modes was rotated (through six different possible permutations) for different participants. The purpose of this counterbalancing measure was to prevent both learning and fatigue effects from introducing confounding factors into the data from a within-groups experiment. Ideally, counterbalancing should have been done using 24 test-subjects (i.e. a multiple of 6). Unfortunately, due to technical reasons, only 23 out of our 24 human test-subjects yielded usable data, however our slightly imperfect counterbalancing over 23 subjects should still have eliminated most learning and fatigue effects from our statistical results. For the secondary task, different cards, but of equal difficulty

[24], were used for each control mode, again to eliminate learning as a confounding factor in the test data.

Participants were instructed to perform the primary task (controlling the robot to reach a destination) as quickly and safely (i.e. minimizing collisions) as possible. Additionally they were instructed that, when presented with the secondary task, they should do it as quickly and as accurately as possible. They were explicitly told that they should give priority to the secondary task over the primary task and should only perform the primary task if the workload allowed. Also they were told that there would be a score penalty for every wrong answer. This experimental procedure was informed by initial pilot study tests, with pilot participants, which showed that when people are instructed to "do both tasks in parallel to the best of your abilities", they either a) ignore the secondary task or b) choose random answers for the secondary task to alleviate themselves from the secondary workload, so that they can continue focusing on the primary task of robot driving. Lastly, participants were informed that the best performing individuals in each trial (using a weighted performance score based on both primary and secondary tasks) would be rewarded with a gift voucher. The purpose of this prize was to provide an incentive for participants to achieve the best score possible on both primary and secondary tasks.

The human operators can only acquire situational awareness information via the Operator Control Unit (OCU) which displays real-time video feed from the robot's front-facing camera, and displays the estimated robot location (derived from laser scanner and SLAM algorithm) on the 2D SLAM map.

Our previous work [4] showed that a difficult confounding factor can be introduced by the fact that different test subjects may explore in different directions, thus revealing different information about the test arena at different times, as the robot's onboard laser SLAM progressively performs mapping. Additionally, real-time SLAM can produce maps of varying accuracy between trials. To overcome this confounding factor, all participants were given an identical and complete 2D map, generated offline prior to the trials by driving the robot around the entire arena and generating a complete SLAM map.

During each trial, a variety of data and metrics were collected: primary task completion time (time taken for the robot to travel from point A to point B and back again to point A (see FIG.1); total number of collisions; secondary task completion time; number of secondary task errors.

At the end of each experimental run, participants had to complete a NASA Task Load Index (NASA-TLX) [25] questionnaire. NASA-TLX is a widely-used, subjective questionnaire tool. It rates perceived workload in order to assess a technology or system. The total workload is divided into six subscales: Mental Demand, Physical Demand, Temporal Demand, Performance, Effort, and Frustration.

V. RESULTS

Statistical analysis was conducted on a number of metrics gathered during the experiments. A repeated measures one-way ANOVA was used, with a Greenhouse-Geisser correction in the cases that sphericity assumption was violated (i.e. that the variances of the differences between conditions/levels are not equal). The independent variable was the control mode with three levels. Fisher's least significant difference (LSD) test was used for pairwise comparisons given the a) clear hypothesis; b) predefined post-hoc comparisons; c) small number of comparisons. LSD is typically used after a significant ANOVA result to determine explicitly which conditions differ from each other through pairwise comparisons. Here we consider a result to be significant when it yields a p value less than 0.05, i.e. when there is less than a 5 percent chance that the observed result occurred merely by chance. We also report on the statistical power of the results. Power denotes the probability that a statistical significant difference will be found, if it actually exists. It is generally accepted that greater than 80 percent chance to find such differences constitutes a good power value. Lastly η^2 is reported as a measure of effect size.

ANOVA for *primary task completion time* (see FIG. 6(a)) showed overall significantly different means with *F(1.275, 28.057) = 34.567,* $p < .01$, $power > .9$, $\eta^2 = .61$ between HI variable-autonomy ($M = 413.6$), autonomy ($M = 483.9$) and teleoperation ($M = 429.6$). Pairwise comparison reveals that pure autonomy performed significantly worse than the other two modes of operation with $p < .01$. Also HI variable autonomy performed significantly better than teleoperation ($p < .05$).

The effect of control mode on the number of *collisions* (see FIG. 6(b)) was significant, *F(1.296, 28.507) = 9.173,* $p < .05$, $\eta^2 = .29$ with a $power > .85$. Pure autonomy mode led to significantly ($p < .05$) fewer collisions ($M = .61$) than teleoperation ($M = 2.43$). HI variable autonomy mode ($M = .57$) also led to fewer collisions ($p < .01$) than teleoperation. HI and autonomy had no significant difference. Playback of the recorded trials revealed that in teleoperation most of the collisions occurred during the time of the secondary task. This was true for the participants that attempted to perform both tasks in parallel.

It is useful to be able to rank each trial according to an overall performance metric, which we refer to as the *primary task score*. This overall score is needed to be able to compare e.g. one human operator who achieves a very fast task completion time, but with many collisions, against another operator who achieves a slower time but with few collisions. We generate the primary task score by adding a time penalty, of $10\ sec$ for every collision, onto the primary task completion time for each participant. This is inspired by the performance scores used in the RoboCup competitions [26]. FIG. 6(a) shows the mean primary task scores for each robot control mode. ANOVA analysis confirmed that control mode had a significant effect on the primary task score, *F(1.336, 29.403) = 19.342,* $p < .01$, $power > .95$,

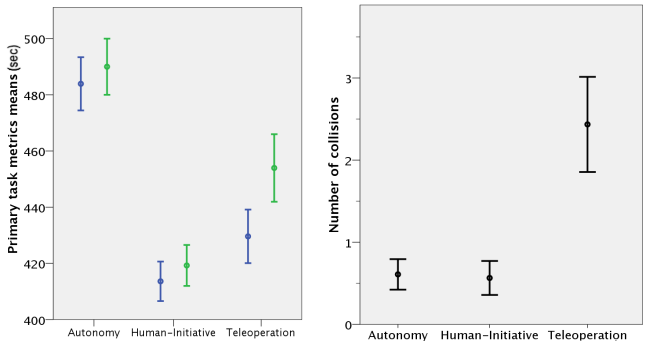

(a) (b)

Fig. 6. Primary task results. 6(a): average time to completion (blue) and score (green) combining time and collisions penalty. 6(b): average number of collisions. In all graphs the error bars indicate the standard error.

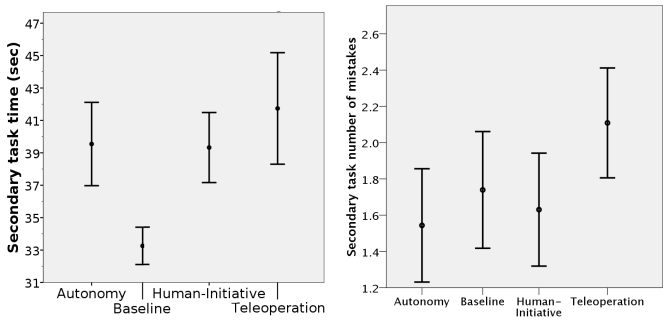

(a) (b)

Fig. 7. Secondary task performance. 7(a): average time to completion for one series of 3D objects. 7(b): average number of errors for one series of 3D objects.

$\eta^2 = .47$. LSD test suggests that HI variable autonomy ($M = 419.2$) significantly ($p < .01$) outperforms both the pure autonomy mode ($M = 490$) and the pure teleoperation mode ($M = 453.9$). Note also that teleoperation appears to outperform autonomy ($p < .05$) in these experiments.

Secondary task completion time (see FIG. 7(a)) refers to the average time per trial, that the participants took to complete one series of the 3D object cards. ANOVA with $F(1.565, 34.420) = 7.821, p < .01, power > .85,\ \eta^2 = .26$, suggests that there is a significant difference between the mean secondary task completion times with and without also performing the primary task of controlling the robot. Participants performed significantly ($p < .05$) better in the baseline trial ($M = 33.2$) compared to their performance during robot operation. During robot operation, HI variable autonomy mode ($M = 39.3$), pure autonomy mode ($M = 39.5$) and teleoperation mode ($M = 41.7$) did not show statistical differences.

No significant differences were observed between the different robot control modes with respect to numbers of secondary task errors (see FIG. 7(b)) according to ANOVA with $F(3, 66) = 1.452, p > .05, power < .8,\ \eta^2 = .06$. Participants had $M = 1.7$ errors during baseline tests without operating the robot, $M = 1.6$ during HI variable autonomy mode, $M = 1.5$ in pure autonomy mode, and $M = 2.1$ in pure teleoperation mode.

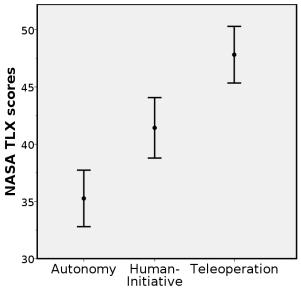

Fig. 8. NASA-TLX score showing the overall trial difficulty as perceived by the operators.

Control mode had a significant effect on *NASA-TLX scores* (see FIG. 8) as suggested by ANOVA ($F(2, 44) = 11.510, p < .01, power > .9,\ \eta^2 = .34$). Pairwise comparisons showed that autonomy ($M = 35.2$) was perceived by participants as having the lowest difficulty, as compared to HI variable autonomy mode ($M = 41.4$) with $p < 0.05$ and teleoperation mode ($M = 47.8$) with $p < 0.01$. HI variable autonomy is perceived as being less difficult than teleoperation ($p < 0.05$).

A. Discussion

In terms of overall primary task performance, HI variable autonomy control significantly outperformed both pure teleoperation and pure autonomy. This confirms our hypothesis that a variable autonomy system with the capability of on-the-fly LOA switching can improve overall performance of the human-robot team. In essence, it does so by being able to overcome situations in which a single LOA may struggle to cope. For example, external distractions to the operator such as the secondary task can be overcome by the operator switching from teleoperation to autonomy. In contrast, when autonomous control struggles to cope with noisy sensory information, the situation can be ameliorated by switching to teleoperation. From the Human-Robot Interaction (HRI) perspective, operators were able to successfully change LOA on-the-fly in order to maximize the system's performance. Since the LOA change was based on the operator's judgement, these experiments suggest that, given sufficient training, operators make efficient use of the variable autonomy capability. Additionally, note that autonomy generates significantly fewer collisions than teleoperation, however HI variable autonomy generates equally few collisions. This reinforces the conclusion that human operators can efficiently exploit autonomy by making smart decisions about switching between autonomy and teleoperation when most appropriate.

Regarding the secondary task, when performed in isolation from the primary task (during baseline testing), participants perform better. Since participants were instructed to focus on the secondary task whenever it was presented, this suggests that even having the primary task waiting on standby was enough to impair their performance on the secondary task. The absence of statistical differences across control modes

in the secondary task time to completion and errors, suggests that a) the choice of control mode did not have any effect on secondary task performance; b) participants had the same level of engagement with the secondary task across trials.

NASA-TLX showed that autonomy is perceived as the easiest control mode, while HI is perceived as being easier than teleoperation. The fact that HI is perceived as more difficult than autonomy might perhaps reflect the cognitive overhead imposed on the operator by having to make judgements about switching LOA. This suggestion was further reinforced by observations made during trials and from informal conversations with participants. Most participants demonstrated a more laid-back attitude while using autonomy. However, participants stated that, while HI variable autonomy mode was "more stressful and demanding", it was also "more fun" due to a perception of increased engagement. For this reason, many participants expressed strong preference for HI variable autonomy over the other control modes. These observations are perhaps related to those of [27] which suggests that humans' "sense of agency" is improved when they interact more actively with a system.

VI. THEORETICAL FRAMEWORK FOR DESIGNING A MI CONTROLLER

The results of these experiments yielded several insights for how to design a MI controller. The robot can be seen as a resource with two different agents having control rights: one agent is the human operator and the other is the robot's autonomous control system. At any given moment, the most capable agent should take control. Of particular importance is the ability of each agent to diagnose the need for a LOA change, and to take control (or hand over control) successfully. We assume that humans are able to diagnose when they need to intervene, given sufficient understanding of the system and the situation. On the other hand, it is not obvious how to enable the autonomous controller to detect when the human operator's performance is degraded, enabling the AI to robustly and automatically take control when it is needed. Automatic switching of control to an autonomous LOA would be important in situations where the human operator is too preoccupied with the primary cause of his or her performance degradation to voluntarily switch control to the robot.

In future work we propose to develop, test and analyse such an MI system. To make initial progress, it may be necessary to at first rely on naive assumptions, such as operators being willing to give control and the context and timing of a LOA change being appropriate [4]. We propose to carry out initial validation of our MI system using the same experimental design as reported in this paper, so that the MI can be compared against the HI system reported here. To be useful, the MI algorithm should provide the same level of performance or better, in terms of primary task completion, as compared to the simpler HI system.

Two different approaches are being investigated for the design of such MI algorithms. The first is focused on *task effectiveness*. The second is focused on using *machine learning techniques* on the HI data gathered during the experiments described in this paper.

In the first approach and more specifically in a navigation task, an online metric could express the effectiveness of *goal directed motion*. In the simplified case, this could be a function of speed towards achieving a desired goal position [28] or the number of collisions inside a thresholded time window. The general idea is that the metric should compare the current speed towards achieving a goal, with the optimal speed towards achieving the same goal.

Such proposals are limited, in that they rely on a variety of assumptions: the full map is known in advance, or the navigational goal lies inside an already known region; the robot's AI possesses a planner which is capable of reliably computing both the optimal path, and also the optimal velocities, from the current pose of the robot towards the goal; the agent to which the control will be traded, is capable of coping with the cause of performance degradation in the other agent.

An alternative approach is one of exploiting machine learning techniques in order to learn patterns of how human operators efficiently change LOA. The HI variable autonomy experiments reported in this paper, have enabled the collection of a variety of measurements that might be used as the training features of such a learning system. Such data includes: current mode of control at each time-step, positions and times of each change of LOA; time-stamped joystick logs; time-stamped series of velocity commands given to the robot; complete robot trajectories and information about periods of robot idle time.

VII. CONCLUSION

This paper presented a principled and statistically validated empirical analysis of a variable autonomy robot control system. Previously, a comparatively small part of the robotics literature has addressed the issues of variable control. Previous studies have focused on the engineering and computer science behind building such systems; or on enhancing the human-robot interface; or investigated the ways in which humans interact with the system.

In contrast, this paper has made a variety of new contributions, including: showing how to carry out a principled performance evaluation of the combined human-robot system, with respect to completing the overall task; presenting clear empirical evidence to support the notion that variable autonomy systems may have advantages over purely autonomous or teleoperated systems for certain kinds of tasks; using rigorous methodologies, transferred from the fields of psychology and human factors research, to inform experimental design, eliminate confounding factors, and yield results that are statistically validated; demonstrates that human operators, when appropriately trained, make successful decisions about switching LOA, which efficiently exploit the contrasting strengths of both teleoperation and autonomous controllers. We must note here that our hypothesis and experimental paradigm are intended to be a starting point, from which more complex hypotheses and scenarios can be formulated.

We believe this is the first study which has used truly scientifically repeatable experiments to support the continued development of variable autonomy mobile robots. Additionally, this paper has discussed the difficult issues involved in extending notions of variable autonomy from Human-Initiative (HI) to Mixed-Initiative (MI) robotic systems, and makes several suggestions for different approaches for building an autonomous MI switching algorithm. Developing such an MI system forms the subject of our ongoing research.

ACKNOWLEDGMENT

This research is supported by the British Ministry of Defence and the UK Defence Science and Technology Laboratory, under their PhD bursary scheme, contract no. DSTLX-1000074621. It was also supported by EPSRC grant EP/M026477/1.

REFERENCES

[1] J. L. Casper and R. R. Murphy, "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*, vol. 33, no. 3, pp. 367–385, 2003.

[2] R. R. Murphy and J. L. Burke, "Up from the Rubble: Lessons Learned about HRI from Search and Rescue," *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*, vol. 49, no. 3, pp. 437–441, sep 2005.

[3] R. R. Murphy, J. G. Blitch, and J. L. Casper, "AAAI/RoboCup-2001 Urban Search and Rescue Events: reality and competition," *AI Magazine*, vol. 23, no. 1, pp. 37–42, 2002.

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

[5] J. Y. C. Chen, E. C. Haas, and M. J. Barnes, "Human Performance Issues and User Interface Design for Teleoperated Robots," *Systems, Man and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 6, pp. 1231–1245, nov 2007.

[6] S. Jiang and R. C. Arkin, "Mixed-Initiative Human-Robot Interaction : Definition , Taxonomy , and Survey," in *Systems, Man, and Cybernetics (SMC), IEEE International Conference on*, 2015, pp. 954–961.

[7] E. Krotkov, R. Simmons, F. Cozman, and S. Koenig, "Safeguarded teleoperation for lunar rovers: From human factors to field trials," in *IEEE Planetary Rover Technology and Systems Workshop*, 1996.

[8] D. J. Bruemmer, D. A. Few, R. L. Boring, J. L. Marble, M. C. Walton, and C. W. Nielsen, "Shared Understanding for Collaborative Control," *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, vol. 35, no. 4, pp. 494–504, jul 2005.

[9] M. Baker and H. A. Yanco, "Autonomy mode suggestions for improving human-robot interaction," in *IEEE International Conference on Systems, Man and Cybernetics (SMC)*, vol. 3, 2004, pp. 2948–2953.

[10] J. L. Marble, D. J. Bruemmer, D. A. Few, and D. D. Dudenhoeffer, "Evaluation of supervisory vs. peer-peer interaction with human-robot teams," in *Proceedings of the 37th Hawaii International Conference on System Sciences*, vol. 00, no. C, 2004, pp. 1–9.

[11] D. A. Few, D. J. Bruemmer, and M. Walton, "Improved Human-Robot Teaming through Facilitated Initiative," in *IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, sep 2006, pp. 171–176.

[12] D. J. Bruemmer, R. L. Boring, D. A. Few, L. Julie, and M. C. Waiton, " I Call Shotgun !: An Evaluation of Mixed-Initiative Control for Novice Users of a Search and Rescue Robot *," in *IEEE Enternational Conference on Systems, Man and Cybernetics (SMC)*, 2004, pp. 2847–2852.

[13] D. J. Bruemmer, C. W. Nielsen, and D. I. Gertman, "How training and experience affect the benefits of autonomy in a dirty-bomb experiment," in *The 3rd ACM/IEEE international conference on Human robot interaction (HRI)*. New York, New York, USA: ACM Press, 2008, pp. 161–168.

[14] M. E. Armstrong, K. S. Jones, and E. A. Schmidlin, "Tele-Operating USAR Robots : Does Driving Performance Increase with Aperture Width or Practice ?" in *59th Annual Meeting of the Human Factors and Ergonomics Society*, no. 2011, 2015, pp. 1372–1376.

[15] C. W. Nielsen, D. A. Few, and D. S. Athey, "Using mixed-initiative human-robot interaction to bound performance in a search task," in *IEEE International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, dec 2008, pp. 195–200.

[16] T. Carlson, R. Leeb, R. Chavarriaga, and J. D. R. Millán, "Online modulation of the level of assistance in shared control systems," in *IEEE International Conference on Systems, Man and Cybernetics (SMC)*, no. 2, 2012, pp. 3339–3344.

[17] B. Hardin and M. A. Goodrich, "On Using Mixed-Initiative Control : A Perspective for Managing Large-Scale Robotic Teams," in *The 4th ACM/IEEE international conference on human robot interaction (HRI)*, 2009, pp. 165–172.

[18] M. A. Goodrich, T. W. McLain, J. D. Anderson, J. Sun, and J. W. Crandall, "Managing autonomy in robot teams: observations from four experiments," in *Proceedings of the ACM/IEEE international conference on Human-robot interaction (HRI)*, 2007, pp. 25–32.

[19] J. M. Riley and L. D. Strater, "Effects of Robot Control Mode on Situation Awareness and Performance in a Navigation Task," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 50, 2006, pp. 540–544.

[20] A. Valero-Gomez, P. de la Puente, and M. Hernando, "Impact of Two Adjustable-Autonomy Models on the Scalability of Single-Human/Multiple-Robot Teams for Exploration Missions," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 53, no. 6, pp. 703–716, oct 2011.

[21] G. Echeverria, N. Lassabe, A. Degroote, and S. Lemaignan, "Modular open robots simulation engine: MORSE," in *Proceedings - IEEE International Conference on Robotics and Automation*, 2011, pp. 46–51.

[22] R. R. Murphy, "HumanRobot Interaction in Rescue Robotics," *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)*, vol. 34, no. 2, pp. 138–153, may 2004.

[23] A. Jacoff, E. Messina, B. Weiss, S. Tadokoro, and Y. Nakagawa, "Test arenas and performance metrics for urban search and rescue robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3. IEEE, 2003, pp. 3396–3403.

[24] G. Ganis and R. Kievit, "A New Set of Three-Dimensional Shapes for Investigating Mental Rotation Processes : Validation Data and Stimulus Set," *Journal of Open Psychology Data*, vol. 3, no. 1, 2015.

[25] D. Sharek, "A Useable, Online NASA-TLX Tool," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 55, no. 1, pp. 1375–1379, 2011.

[26] A. Jacoff, E. Messina, B. Weiss, S. Tadokoro, and Y. Nakagawa, "Test arenas and performance metrics for urban search and rescue robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3. Ieee, 2003, pp. 3396–3403.

[27] W. Wen, A. Yamashita, and H. Asama, "The sense of agency during continuous action: Performance is more important than action-feedback association," *PloS one*, vol. 10, no. 4, p. e0125226, 2015.

[28] D. R. Olsen and M. A. Goodrich, "Metrics for evaluating human-robot interactions," in *PERMIS*, 2003.