Mixed-initiative Variable Autonomy for Remotely Operated Mobile Robots

MANOLIS CHIOU, Extreme Robotics Lab, University of Birmingham, UK

NICK HAWES, Oxford Robotics Institute, Department of Engineering Science, University of Oxford, UK

RUSTAM STOLKIN, Extreme Robotics Lab, University of Birmingham, UK

This article presents an Expert-guided Mixed-initiative Control Switcher (EMICS) for remotely operated mobile robots. The EMICS enables switching between different levels of autonomy during task execution initiated by either the human operator and/or the EMICS. The EMICS is evaluated in two disaster-response-inspired experiments, one with a simulated robot and test arena, and one with a real robot in a realistic environment. Analyses from the two experiments provide evidence that: (a) Human-Initiative (HI) systems outperform systems with single modes of operation, such as pure teleoperation, in navigation tasks; (b) in the context of the simulated robot experiment, Mixed-initiative (MI) systems provide improved performance in navigation tasks, improved operator performance in cognitive demanding secondary tasks, and improved operator workload compared to HI. Last, our experiment on a physical robot provides empirical evidence that identify two major challenges for MI control: (a) the design of *context-aware* MI control systems; and (b) the *conflict for control* between the robot's MI control system and the operator. Insights regarding these challenges are discussed and ways to tackle them are proposed.

CCS Concepts: • **Human-centered computing** → **Collaborative interaction**; • **Computing methodologies** → **Control methods**; • **Computer systems organization** → **Robotics**;

Additional Key Words and Phrases: Mixed-initiative control, shared control, human-robot interaction, variable autonomy, conflict for control

ACM Reference format:

Manolis Chiou, Nick Hawes, and Rustam Stolkin. 2021. Mixed-initiative Variable Autonomy for Remotely Operated Mobile Robots. *Trans. Hum.-Robot Interact.* 10, 4, Article 37 (September 2021), 34 pages.

<https://doi.org/10.1145/3472206>

This work was funded by the British Ministry of Defence via the Defence Science and Technology Laboratory (Dstl), under their PhD bursary scheme, Contract No. DSTLX-1000074621. It was also supported by the UK's Engineering and Physical Sciences Research Council (EPSRC) under the National Centre for Nuclear Robotics Grant No. EP/R02572X/1 and related Grants No. EP/M026477/1, No. EP/P017487/1, and No. EP/P01366X/1. Rustam Stolkin was partly funded by a Royal Society Industry Fellowship.

Authors' addresses: M. Chiou and R. Stolkin, Extreme Robotics Lab, School of Metallurgy and Materials, University of Birmingham, Elms Rd, Birmingham B15 2SE, UK; emails: m.chiou@bham.ac.uk, r.stolkin@cs.bham.ac.uk; N. Hawes, Oxford Robotics Institute, University of Oxford, 23 Banbury Road, Felstead House, Oxford OX2 6NN, UK; email: nickh@robots.ox.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

2573-9522/2021/09-ART37 \$15.00

<https://doi.org/10.1145/3472206>

1 INTRODUCTION

In the aftermath of the 9/11 World Trade Center terrorist attack, several robots were used in an **Urban Search and Rescue (USAR)** operation. The robots' tasks were to inspect areas beneath debris and rubble, or enter confined spaces (e.g., dangerous voids), that humans or dog rescuers could not enter. Robot operators were working under stressful and performance-degrading conditions such as cognitive fatigue, sleep deprivation, and minimum sensor information (e.g., caused by dust or poor lighting), which led to mistakes in robot operation [8].

On March 11, 2011, a tsunami hit Fukushima Daiichi nuclear power plant causing three nuclear meltdowns and the release of radioactive material. In an effort to initially assess and contain the situation, military-grade robots were used to inspect the buildings and the reactors for any damage, radiation, and other potentially hazardous conditions. These robots had to be operated by the nuclear power plant workers, after only brief training sessions. A typical scenario involved those newly trained operators remotely controlling the robots for many hours while wearing a hazmat suit, making operation a very challenging task. Additionally, these operators had to face a number of psychologically challenging and stressful situations (e.g., dosimeter alarms going off), as famously reported by one of the robot operators in his blog.¹

Several years after these disasters, many of the initial challenges and shortcomings of robot use in disaster response remain. Despite the need to make the use of robots in safety and time critical applications (e.g., USAR, nuclear decommissioning, and hazardous environment inspection) easier, robots used are still predominately teleoperated, with little or no autonomy used to assist the human operator. There are three main reasons why autonomous systems are not used in these settings. First, these tasks often involve environments that are highly unstructured and changing, e.g., a partially collapsed building. Second, the nature of these tasks requires specific human abilities such as critical decision making based on incomplete information (e.g., determining if a victim dead or alive; and risk assessment of certain actions); or communication with victims in **Search and Rescue (SAR)** [14]. Third, high consequence industries tend to be conservative and therefore do not trust autonomous systems. For example, in the Fukushima Daiichi nuclear accident, such lack of trust led to the robots being deployed via pure teleoperation, despite having various autonomous capabilities [37]. For these reasons, current deployment of such robots always requires a human in the loop [34].

Equipping robots with autonomous capabilities can potentially tackle some of the intrinsic difficulties of teleoperation. Several field studies have pointed out the need for robots that actively assist operators (e.g., in the 9/11 World trade center [8]; the DARPA robotics challenge [40, 50]; and other studies, e.g., References [35, 36]). What is required is a human-robot team that benefits from the capabilities of both agents, whilst counteracting the weaknesses of each. Such systems offer the potential to assist a human operator who may be struggling to cope with issues such as high workload, intermittent communications, operator multitasking, fatigue, and sleep deprivation. For example, a human operator might need to concentrate on a secondary task while temporarily delegating control to the robot to navigate autonomously. This is something very common as robot operators have to convey situational awareness information, e.g., to SAR task force team mates [6, 35].

Our research addresses the use of *variable autonomy* as an approach to blending the capabilities of humans and robots. A variable autonomy system is one in which control can be traded between a human operator and a robot by *switching between different Levels of Autonomy* [11]. **Levels of Autonomy (LOAs)** refer to the degree to which the robot, or any artificial agent, takes its own

¹IEEE Spectrum, Accessed: 2020-06-26: <http://spectrum.ieee.org/automaton/robotics/industrial-robots/fukushima-robot-operator-diaries>.

decisions and acts autonomously [48]. In the case of robotics, LOAs can vary from the level of pure teleoperation (human has complete control of the robot), to the other extreme, which is full autonomy (robot has control of every capability), within a single robot.

This article addresses the problem of dynamically changing LOA on-the-fly (i.e., during task execution) using either **Human-Initiative (HI)** or **Mixed-Initiative (MI)** control. HI refers to a system where the human operator is solely responsible for switching LOA based on their judgement. MI refers to a system where both the human operator and the robot's control switcher have authority to initiate actions and to change the LOA. How best to dynamically switch LOAs to improve system performance is a challenging and open problem that remains largely unexplored in the literature.

The hypothesis explored in this article is that timely switching between different LOAs during task execution (e.g., during navigation) can improve task performance and will enable the system to overcome various performance-degrading factors. This is compared to robotic systems in which LOAs cannot switch on-the-fly. More specifically, we investigate the capacity of the human operator (based on their judgement) and the robot's MI control switcher (based on an online performance metric) to use HI and MI control to improve performance compared to pure teleoperation or autonomy. In the rest of this article, when we use the term "robot," we use it for the physical robot hardware. The term **Expert-guided Mixed-Initiative Control Switcher (EMICS)** will refer to the robot's MI control switcher, i.e., the autonomous software system that initiates the LOA switching. The terms autonomy and teleoperation will refer to the individual LOA in terms of navigation.

1.1 Scope of Current Research and Contributions

Previous work [12] demonstrated that conducting variable autonomy experiments is challenging due to the high number of intrinsic confounding factors that can lead to large variances in the results. Such confounding factors include individual differences in personality traits; experience in operating robots or playing video games; and map exploration strategies. Subsequent work contributed a systematic variable autonomy experimental framework to avoid these issues [11]. In Reference [11], human operators used HI control to navigate a remotely controlled mobile robot in a maze-like test arena. The HI control switcher allowed operators to switch between two different LOAs: autonomous navigation and teleoperation. The results showed that that HI variable autonomy can overcome various performance-degrading factors and outperform both pure teleoperation and pure autonomy in various circumstances where using the complementing capabilities of humans and autonomous agents is beneficial. Additionally, in Reference [10] the HRI aspects of the human operator interacting with and exploiting the HI control switcher were explored. It provided evidence that operators' interactions with the variable autonomy control switcher (e.g., frequency of LOA switching; preferred LOA) are not necessarily influenced only by task performance.

In this article, we advance variable autonomy research by moving from HI to MI robotic systems. We also aim to provide a complete picture of research regarding robots which dynamically switch LOA and make the case that HI or MI control provides advantages over systems with single modes of operation (e.g., pure teleoperation). This is done through the design and evaluation of an MI control switcher (i.e., the EMICS).

This article makes the following contributions:

- It presents statistically validated empirical evidence that HI systems outperform single mode of operation systems (e.g., pure teleoperation) in navigation tasks, based on real robot USAR-inspired experiments. The evidence refer to circumstances where using the complementing capabilities of humans and autonomous agents is beneficial.

- It presents a framework and guidance for designing robotic MI control systems. More specifically, the design of *expert-guided MI control switchers* is proposed.
- It proposes an MI control system in which both the operator (based on judgement) and the robot's EMICS (based on an online performance metric) are able to switch between different LOAs. Our approach in designing the EMICS makes use of: expert knowledge; an online performance metric; and simplified contextual knowledge to infer if a LOA switch is needed.
- It identifies and presents empirical evidence and insights into two major challenges for MI control: the design of context-aware MI control systems; and the conflict for control between the operator and the MI control switcher.
- It provides empirical evidence regarding the potential advantages of the MI control in various circumstances, as compared to HI and teleoperation in our simulated robot experiments.

2 BACKGROUND AND RELATED WORK

In this section, we discuss research on strategies for switching LOAs. More specifically, we identify gaps in the literature regarding three key aspects: (a) conducting rigorous experiments on LOA switching; (b) LOA switching initiated by the human operator (namely, HI); and (c) LOA switching initiated by the robot's MI control switcher or the operator (namely, MI). The focus of this article is MI control and thus the literature relevant to MI robotic systems will be reviewed in detail. Literature on HI control is briefly discussed for completeness. For further information on experimental frameworks and HI please refer to our previous work [10–12].

2.1 Human-Initiative Variable Autonomy

A **Human-Initiative (HI)** variable autonomy system is a system in which the human operator can dynamically switch between different LOAs (e.g., between teleoperation and autonomy) [11]. In such systems only the human operator has the authority to initiate LOA switches based on their judgement. Many of the systems found in literature, such as Reference [22], are restricted to an initial LOA choice as they cannot change LOA on the fly. Other systems (e.g., References [2, 17, 47]) aid the operator's judgement by suggesting potential changes in the LOA. Other SAR-inspired studies (e.g., Reference [32]) are mostly focused on evaluating the usability of such HI systems.

2.2 Human-Robot Interaction with LOA Switching Robots

Human-Robot Interaction (HRI) in a LOA switching system remains mainly unexplored in the prior literature. Previous variable autonomy studies investigated the human operator's interaction with a robotic system, but were restricted to exploring a single LOA [4, 7]. Other studies [2, 32, 47] did not present any data on the operator's interaction with the LOA switching controller. In Reference [10], an analysis of the ways in which human operators interact with, and exploit the capabilities of, a HI robotic system with dynamic LOA capabilities is reported. In this article, we report on the HRI between the operator and the EMICS.

2.3 Definition and Taxonomy of Mixed-Initiative Control

In Reference [27] Jiang and Arkin define MI control in the context of human-robot teams as:

“A collaboration strategy for human-robot teams where humans and robots opportunistically seize (relinquish) initiative from (to) each other as a mission is being executed, where initiative is an element of the mission that can range from low-level motion control of the robot to high-level specification of mission goals, and the initiative is mixed only when each member is authorized to intervene and seize control of it.”

They also present the first taxonomy for MI robotic systems. Their taxonomy has three dimensions. The first dimension is **span-of-mixed-initiative**, which characterizes the control elements (initiatives) in which both agents are capable of initiating actions. The system we present in Section 3.2 is *mostly joint* (a term from this dimension in Reference [27]), since both agents have initiative over two of the control elements: navigation execution and LOA switch. The second taxonomy dimension is **initiative reasoning capacity**, which characterizes the ability of an agent to reason about taking the initiative. Our system is *deliberative*, since it has the ability to reason about initiating actions deliberately based on an online performance metric and simplified context awareness. The final dimension of the taxonomy is **initiative hand-off coordination**, which characterizes the strategies used by the system when shifting initiative from one agent to the other. Our system is characterized as *explicitly coordinated* in this dimension.

2.4 Mixed-Initiative Control Systems

Given the comprehensive taxonomy in Reference [27], in this survey, we review systems that have some form of MI control (i.e., both human and operator can initiate actions or LOA switches). Shared control is a term often used generally to describe systems in which the human cooperates at some level with the robot's **Artificial Intelligence (AI)** controller; or is used specifically to describe systems in which some form of input mixing or policy blending between the robot's AI controller and/or the operator's commands is used [15]. In this article, to avoid confusion, we will use this term to explicitly refer to research involving some form of input mixing. Shared control from our perspective is a control mode that can fall under the banner of MI control. Any mixed initiative behaviour is restricted within the shared control mode (or LOA). This means that the robot's AI controller can only take the initiative to blend its navigation control input with the one from the operator, to improve the control output. Similarly, in safeguard teleoperation (another form of shared control), the AI controller reactively takes the initiative to prevent collisions. In both cases no LOA switching takes place. Also in most cases the human operator does not have any initiative in choosing how to blend the control inputs. In Section 3.3 further discussion can be found on how MI control and our system relates to shared control and input blending.

Nielsen et al. [39] conduct experiments using multiple LOAs. The LOA is chosen during the initialization of the system and cannot change on the fly. Moreover, similar to shared control, the AI controller has reactive initiative inside a specific LOA to prevent collisions. Multiple LOAs (teleoperation, safe mode, shared mode, autonomy) are tested in Reference [4]. In shared mode the robot's AI controller drives autonomously while accepting interventions from the operator. In safe mode the AI controller takes initiative to prevent collisions. Few et al. [16] present a control mode in which the operator gives directional commands using a joystick to adjust the robot's navigation. The system offers initiative, which depends on the frequency of the operators interaction. Rigger et al. [43] propose a contextual multi-armed bandit framework for learning policies that enable a control switcher to choose between LOAs with the goal to minimize human interventions, e.g., recovering the robot from failures.

Gombolay et al. [20, 21] use an industrial assembly scenario to investigate how varying degrees of task allocation authority in physical collaboration between the humans and the robot affect the human-robot team's task performance, SA, and task work-flow. Their system has three different LOAs that can be chosen in the beginning of the scenario: (a) the human decides how to allocate the tasks; (b) the human decides which tasks they will perform, while the AI controller allocates the remaining tasks to itself and to other humans; (c) the AI controller has full authority on task allocation.

Bruemmer et al. [5] present a theoretical, multiple LOA, MI system. This system is based on theories of robot behavior (human understanding of the robot) and human behavior (robot

understanding of the human). The latter proposes gathering readily available non-intrusive workload cues from the operator as an indication of poor performance. More specifically, it proposes the use of the frequency of human input and the number and kind of dangerous commands issued by the operator as performance indicators. This can provide the robot's MI control system with the capacity to initiate switches between the different LOAs. Adams et al. [1] propose a MI robot control architecture that relies on the detection of the operator's emotional state. Initiative is mixed in all the levels of the system, i.e., in setting goals and constraints, planning and execution. Changes in control are initiated based on the operator's sensed state (e.g., boredom, stress, drowsiness, engagement).

For multi-robot systems, variable autonomy often lies at a higher level of abstraction compared to single robot systems. Manikonda et al. [31] describe a multi-robot MI control system and testbed for human-robot teams in tactical operations. The agents in the system share information via a common world model. Based on this information they are able to initiate modifications to their goals and associated roles in the team. In Reference [23] a MI approach is proposed in a multi-robot search task. Robots are equipped with the ability to initiate changes in their respective search areas (e.g., size of search area). These changes are reactively triggered by specific events, e.g., the human operator has identified an item of interest. In References [24, 29] a hierarchical reinforcement learning architecture is proposed for semi-autonomous control of multiple robots by one operator. It is focused on task allocation between robots. Additionally, when a task is failing, the system asks for operator's help.

In summary, MI robotic systems found in the literature often offer limited initiative inside a pre-defined LOA. Some of the MI systems proposed are either purely theoretical, or not experimentally evaluated. Moreover, initiative actions from the robot's AI controller are often based on reacting to sensor input (e.g., obstacles) rather than online task performance metrics. In the case of multi-robot systems, the MI lies in a higher level of abstraction, making assumptions about other layers, e.g., navigation.

The work presented in this article aims to address some of the gaps in the literature by focusing on the problem of switching LOA during task execution. This is achieved by designing a MI control switcher that relies on an online task performance metric.

Last, some of the literature lack the use of a rigorous experimental framework, e.g., absence of standardized training for the human operator; absence of statistical analysis; protocols not explained. In this article, we aim to address this shortcoming by using a rigorous experimental framework that consists of: (a) appropriate statistical analysis; (b) clarity on assumptions and hypotheses; (c) precise and detailed descriptions of the experimental protocols followed; and (d) a formalized and repeatable experimental paradigm. Because of the nature of this research (i.e., human-robot systems in complex problems and environments) it is very difficult for this or other related studies to fully control for every possible confounding factor. Here, we aim for a balance between realism (i.e., realistic tasks and scenarios) and meaningful scientific inference (i.e., controlled experiment).

3 AN EXPERT-GUIDED FRAMEWORK FOR DESIGNING MIXED-INITIATIVE CONTROL SYSTEMS

The fundamental problem in MI control is how to allow LOA switching by either the operator or the robot's MI control switcher to improve the human-robot system's performance. Without loss of generality, in the rest of this article, we assume a human-robot system that has two LOAs: one in which a human has most of the control (i.e., some form of teleoperation), and one in which an AI controller has most of the control (i.e., some form of autonomous control). However the descriptions in the following can be extended to the case where other LOAs exist, e.g., alternate autonomous or shared control LOAs.

In MI control the robot, meaning the hardware, can be seen as a resource with two different agents having control rights: the human operator and the MI control system. At any given moment, the most capable agent should take control. Hence, of particular importance is the ability for each agent to diagnose the need for an LOA change, and to take or relinquish control successfully.

An operator's LOA switches are based on judgement. As demonstrated by the HI results in Reference [11], given sufficient understanding of the system and the situation humans are able to determine when they need to change LOA. More specifically, analysis of the HI results from our previous work [10] revealed that operators switch LOA based on three factors: (a) preferred LOA; (b) context; and (c) performance degradation. The preferred LOA is the LOA the participants tend to instinctively return to. Their preference may be based on a number of task- or system-specific factors (e.g., trust in the control software) and personal traits (e.g., preference to be in control). In context-sensitive LOA switches, operators are able to evaluate the current context and infer if a LOA switch is needed. For example, they can change control preemptively as they predict performance degradation in a given situation.

To enable the MI control system to automatically take control when it is needed, and to relinquish control when under-performing, it is necessary to provide it with a mechanism to detect when the performance of the human-robot system is degraded. For example, the MI control switcher could initiate a switch to a LOA that offers increased autonomous capabilities in situations where the human operator is too preoccupied with the primary cause of their performance degradation to voluntarily switch control.

To this end, we propose an *expert-guided approach* to the design of MI control switchers. For a given task, we assume the existence of a *task expert* that can provide the expected task performance for the human-robot system in the absence of performance-degrading factors. Based on this assumption, the core idea is that the **Expert-guided Mixed-Initiative Control Switcher (EMICS)** should compare the current performance of the system with the expert performance. This comparison yields an online *task effectiveness* metric, expressing how effective the system is performing its given task relative to the performance of the expert. If this metric indicates that the current task performance is worse than the expert expects, then the EMICS should initiate a LOA switch. Such experts can be created in many ways, e.g., through simulation, assuming no noise, learning by demonstration, and so on.

The application of our expert-guided MI approach to a given human-robot team and task relies on the following assumptions: (a) the human operator is willing to be handed control and to hand over control based on the initiative of the EMICS; (b) the agent to which the control will be handed (i.e., either the human or the MI control system) is capable of correcting the task effectiveness degradation [12]; and (c) the system is equipped with an expert controller providing information on how (given the assumptions made in the design of the expert) the system should be performing in a given task.

3.1 An Expert-guided Mixed-initiative Control Switcher for Navigation

Many of the applications for remotely operated robots involve navigation. Therefore, for the remainder of this article, we focus on the problem of MI control for navigation, specifically the task of the moving the robot to a goal location as quickly and as safely as possible.

Creating an EMICS for navigation requires an online task effectiveness measure for navigation. To this end, we propose the *goal-directed motion error* (referred to as "error" for the rest of the article). This is the difference between the robot's current motion and the motion of the robot required to achieve its goal (reach a target location).

To provide more environment context to the EMICS (i.e., to account for environment structure), we extract an expert performance measure from a concurrently active navigation system, which

Fig. 1. Block diagram of the expert-guided MI control system. Given a navigation goal; the current pose of the robot; and the map (or a known surrounding area); the expert planner yields an ideal suggested speed. This speed denotes how fast the robot should be moving toward achieving the navigation goal. Then this speed is compared with the current speed of the robot toward that goal to calculate the raw performance error. The MI control switcher (i.e., the EMICS) decides on switching LOA based on the filtered error.

is given an idealised (obstacle- and noise-free) view of the robot's world. In a real scenario (e.g., the robot has to inspect a building for damage) such view or map can be provided to the system as prior knowledge, e.g., taken from the building's plan. This expert uses the navigation planner framework and algorithms (i.e., Dijkstra's algorithm and the dynamic window approach) from the **Robot Operating System (ROS)** navigation stack [33]. Using this, the expert continually reports the velocity with which the robot should be moving toward its goal. In essence, the expert provides an idealised model of possible robot behavior.

From the expert-reported velocity, we extract the linear speed (i.e., in the x axis denoting the forward/backward motion of the robot). This denotes the speed that the robot should be moving along the optimal path toward the goal. To create an error measure, the expert's linear speed is constantly compared with the current linear speed of the robot. The difference of these values is the *goal directed motion error* that we use in the EMICS. Please refer to Figure 1 for the block diagram of the system.

To implement this approach, the goal-directed motion error must satisfy the following requirements: (a) frequent, brief changes in error should be treated as noise; and (b) the final error signal to be used should express the accumulated error over time. The second requirement comes from our model, since it assumes that the performance drop is observed over a short period of time before any initiative takes place. The **exponential moving average (EMA)** algorithm [3] applied to the raw error (see Equation (1)), satisfies both of these requirements:

$$E_t = \alpha \cdot e_t + (1 - \alpha) \cdot E_{t-1}. \quad (1)$$

The term E_t refers to the final accumulated and smoothed error. The term e_t refers to the current error observation. The smoothing factor α controls the weight of the most recent (i.e., current) observation and thus the time window in which the error will be accumulated.

We assume that human operators switch LOA when the filtered error E_t reaches a certain threshold. This is for the cases in which operators initiate a performance-based LOA switch. We determined the values of the parameters (i.e., error threshold and the smoothing factor α) by performing a grid search. For each parameter setting, we replayed the robot motion for all operator trials from the HI experiment in Reference [11] through a simple threshold control switcher. For each trial, this yielded a set of LOA changes proposed by the control switcher, which aim to model the human LOA switches in HI. We compared this set of LOA switches to the LOA switches

initiated by the operators in the previous HI data. We then used the cost function in Equation (2) to calculate the cost of a particular parametrisation of the control switcher for each individual trial. The final parameters chosen were the parameters that minimized the sum of costs for all participants. The cost function for each individual trial was

$$j = \sum_{k=1}^n |\hat{t}_k - t_k| + c \cdot p. \quad (2)$$

Put simply, Equation (2) expresses the LOA switch prediction error from the threshold control switcher on HI data, with the addition of a penalty term for every prediction that does not match (i.e., false positives) where $\hat{t}$ is the time-stamp of the predicted LOA switch and t is the time-stamp of the operator's actual switch. The subscript k denotes a specific time-stamp, where n is the total number of time-stamps, c is the number of predictions that do not match operator's switches, and p is a cost penalty. A non-matching prediction is defined as a prediction that falls outside of a small time window (± 5 s) around any of the operator's LOA switches. The assumption is that these predictions do not correspond to any actual operator's LOA switches and thus they get penalized.

The parameters calculated with the grid search were $\alpha = 0.06$ and error *threshold* = 0.07 ms^{-1} . The α calculated here represents a trade-off between fast reactions to large current errors and filtering noise (i.e., sudden instantaneous errors). These two extremes (i.e., fast reactions versus react cautiously to accumulated error) are both included in realistic navigation tasks similar to the ones in our experiments, as they can both occur during task execution. Hence, the α value does not need to be re-derived for every scenario as long as the scenario is of the same nature and not a specialized case of navigation. For this reason, the sensitivity of the control switcher's ability to react and switch LOA is primarily governed by how the error value is managed.

3.2 Fuzzy Mixed-initiative Control Switcher

The LOA switching problem under investigation is complex and its underlying dynamics are not precisely known. Thus, a fuzzy control switcher offers several advantages. Primarily it allows the efficient management of the real world fuzziness through the use of expert knowledge coming from human operators. Additionally, the control switcher can be extended to have more inputs and output states to facilitate a more complex system. Last, fuzzy logic enables more transparency and better understanding of how the control switcher works, something important given the current state of the research.

The EMICS was designed as a bang-bang (also known as on-off controller) Mamdani type fuzzy controller [30, 38]. In this case, the bang-bang control switcher has two states: (a) switch LOA; (b) do not switch LOA. The fuzzy lite [42] C++ library was used for the task. The repository² containing the ROS code for the MI control switcher and any code necessary to replicate the experiments described in this article, is provided under MIT license.

In fuzzy logic a linguistic variable associates a linguistic concept and the fuzzy set representing that concept to a numerical value. Our fuzzy control switcher has two variables as inputs. The first linguistic input variable is "error" that denotes the filtered error E_t . The smoothing factor α used in the EMA filter, was the one found by the parameters search for the threshold control switcher. The second input variable is "speed" that denotes the speed at which the robot is currently moving.

For the linguistic input variable $x_1 = \text{"error"}$ the universe of discourse (i.e., the range of all possible values for an input) is $X_1 = [0, 0.1] \text{ ms}^{-1}$. A value of 0 ms^{-1} means that no goal directed motion error exists, i.e., the robot is making progress toward the goal without any performance

²Extreme Robotics Lab GitHub repository: https://github.com/uob-erl/fuzzy_mi_controller.

Fig. 2. **2(a)**: Membership functions for the linguistic input variable “error.” **2(b)**: Membership functions for linguistic input variable “speed.”

degradation. The value of 0.1 ms^{-1} is the maximum error, meaning that the robot is not progressing toward the goal (e.g., robot is idle). This is despite the maximum speed of the robot being 0.4 ms^{-1} . Due to physical and mechanical constraints (e.g., acceleration limits) the maximum new speed that can be given to the robot at any time, cannot have a difference with the current speed more than 0.1 ms^{-1} . For the linguistic input variable $x_2 = \text{“speed”}$ the universe of discourse is $X_2 = [-0.4, 0.4] \text{ ms}^{-1}$. The value of 0.4 ms^{-1} is the maximum speed of the robot. The value of -0.4 ms^{-1} is the maximum reverse speed of the robot.

The input variable “error” (i.e., x_1) is mapped into three linguistic values (i.e., three fuzzy set membership functions; see Figure 2(a)). The set A_1^i for input x_1 can be defined for the following linguistic values: $A_1^i = [A_1^1 = \text{small}, A_1^2 = \text{medium}, A_1^3 = \text{large}]$. The error threshold calculated from the operators’ HI data was used to denote the fuzzy linguistic value “large.” In essence what operators consider to be a large error to justify a LOA switch, is encoded into the fuzzy control switcher. The values and membership functions for “small” and “medium” were heuristically chosen to smoothly overlap throughout the universe of discourse (see Table 4). This is a common practice when designing fuzzy controllers.

The input variable “speed” (i.e., x_2) is mapped into three linguistic values (see Figure 2(b)). The set A_2^i for input x_2 can be defined for the following linguistic values: $A_2^i = [A_2^1 = \text{reverse}, A_2^2 = \text{zero}, A_2^3 = \text{forward}]$. The value “reverse” denotes that the robot’s speed is negative, which means the robot is reversing. The value “zero” denotes that the robot is idle and not moving. The value “forward” denotes that the robot is moving forward (see Table 4).

The fuzzification process is achieved using the fuzzy membership functions described in the appendix (see Appendix Table 4). Then, a set of fuzzy rules is applied to the fuzzy inputs (see Table 1). The fuzzy rules used in the control switcher were constructed using expert knowledge from HI data. Note that if the “error” is “large” and the robot is reversing, the assumption is that the agent in control is trying to extricate the robot from an error situation. Hence, by taking into account this simplified contextual knowledge, the control switcher will not switch LOA.

Similar to the work of Nagi et al. [38], we follow the **fuzzy bang-bang relay controller (FB-BRC)** approach. This means that the **Largest of Maxima (LOM)** defuzzification method is used. The output’s universe of discourse $Y = [-1, 1]$ represents the bang-bang output (see Figure 3). The value of $y = -1$ means that no LOA switching initiative will take place. The value of $y = 1$ means that the control switcher will initiate a LOA switch.

3.3 Relevance of the Mixed-initiative Control Switcher to Shared Control

As explained in Section 2.4, we consider shared control to be a specific control mode that only loosely falls under the banner of MI control. This is because in most cases of shared control only

Fig. 3. Output membership functions.

Table 1. Fuzzy Rule Base

No.	Rules
1	IF error is small OR error is medium THEN LOA is no change
2	IF error is large AND speed is not reverse THEN LOA is change
3	IF speed is reverse AND error is large THEN LOA is no change

the robot’s AI controller has the initiative to adjust operator’s control input, e.g., to blend robot commands with operator’s joystick navigation commands. According to Dragan and Srinivasa [15] shared control is a policy blending. More specifically, “an arbitration of two policies, namely, the user’s input and the robot’s prediction of the user’s intent. At any instant, given the input, U , and the prediction, P , the robot combines them using a state-dependent arbitration function $a \in [0, 1]$.”

In the framework of Reference [15], two extreme cases exist: (a) the operator’s input U is used as the output command directly; or (b) the AI controller fully blends/changes U and hence the AI controller’s commands are used as output. These two extreme cases roughly correspond to our MI approach with the equivalent of the arbitration function being the EMICS. Case (a) corresponds to the teleoperation LOA (operator fully in control); and case (b) corresponds to the autonomy LOA (robot’s autonomous navigation capabilities fully in control).

There are fundamental differences between the above formalism and our EMICS. In shared control the arbitration function decides how the operator’s input will be blended with the robot’s AI commands in a continuous range from teleoperation to autonomy, and only in a very specific control element. This control element usually lies in low level control, e.g., velocity navigation commands. Also the initiative lies within the arbitration function. In contrast, the EMICS proposed here allows switches between different discrete LOAs with both the operator and the EMICS having the authority to initiate or override each other’s actions. Last, the framework and the control switcher presented in this article (see Figure 1) are applicable and can be extended to also include multiple LOAs without changing their fundamental principles and/or include higher levels of abstraction.

4 EXPERIMENT 1: EVALUATION USING A SIMULATED ROBOT AND TEST ENVIRONMENT

To allow direct comparisons between MI and HI control, we first evaluate our EMICS using the same framework as in our previous work [11]. To be useful, the MI algorithm should provide the same level of performance or better in terms of primary task completion time or score when compared to the HI system. The reason that the MI control switcher can be useful despite potentially having the same level of performance as HI is that there are situations in which a human-initiated LOA switch may not be possible (e.g., loss of communication or incapacitated operator);

Fig. 4. **4(a)**: Simulated arena and the robot model used in the experiment. **4(b)**: 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.

therefore, the MI control switcher has the ability to provide additional system redundancy, which is not evaluated in these experiments.

The aim was to make an initial evaluation of the EMICS and compare its performance with that of the HI control switcher. More specifically, the experiment described in this section evaluates: (a) the human's and EMICS's ability to switch LOA between teleoperation and autonomy to overcome circumstances in which an MI system is under-performing; (b) how the EMICS performs compared to the HI one of Reference [11]; (c) the unfolding Human-Robot Interaction (HRI) of MI control.

4.1 Experimental Setup – Test Arena, Tasks and Tested Conditions

In the previous work [11], we carried out experiments in a high fidelity simulated arena with dimensions of 11×13.5 m (see Figure 4) using **Human-initiative (HI)** control. The robot was controlled by an **Operator Control Unit (OCU)**, composed of a laptop, a joystick, a mouse and a screen showing the control interface. We used a simulated Pioneer-3DX mobile robot fitted with a camera and a laser scanner. We conducted the experiment described in this section using an identical setup to Reference [11], i.e., identical robot arena; system; experimental paradigm and procedures. This is to facilitate direct comparison between the control switchers.

Operators were given the primary task of navigating from point A in Figure 4(b) (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 degrading factors occurred separately from each other, as shown in Figure 4(b).

Fig. 5. **5(a)**: Typical example of a rotated 3D objects card. **5(b)**: 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**: Map showing the position of the robot, the current goal (blue arrow), the navigation planner path (green line), the obstacles' laser reflections (red) and the walls (black).

More specifically, autonomous navigation was degraded by adding Gaussian noise to the laser scanner range measurements, thereby degrading the robot's AI 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; or when a robot travels through an area that is dusty.

To degrade the performance of the human, a secondary task of mentally rotating 3D objects was used [18]. 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 5(a)). These cards were put by the experimenter on the OCU desk, on the right hand side of the operator. 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 whether or not the two objects were identical.

In Reference [11], for each human test-subject, three different control modes were tested: (a) **teleoperation** mode, in which the human operator was restricted to using only direct joystick control to drive the robot; (b) **autonomy** mode, in which the operator was only allowed to guide the robot by clicking desired destinations on the robot's laser-generated 2D map; (c) **HI** mode, in which the operator was allowed to switch LOA at any time (using a button on the joystick) according to their judgement, to maximize performance.

In the experiment presented here each operator was tested in **Robot-Initiative (RI)** and **MI** control. The EMICS, as described in Section 3.2, gives the MI system and the operator the capacity and authority to switch dynamically (i.e., during task execution) between teleoperation and autonomy. The EMICS initiates LOA switches based on *goal directed motion* performance (i.e., effectiveness) and simplified (i.e., limited) context. The operator switches LOA using a button on the joystick, according to their judgement. The RI control switcher uses identical mechanisms with the EMICS for the switching behavior. However, it removes the operator's ability to trigger LOA switches and thus it restricts them to using the RI control switcher's dictated LOA. When a LOA switch occurs, the control switchers alert operators in three different ways using: (a) an alarm sound identical to the one denoting "engine failure" in aircrafts; (b) synthetic speech expressing the LOA the control switcher have switched to; (c) a GUI notification.

4.2 Participants and Procedure

The 24 participants of our previous HI study were asked to participate in this experiment. From these, 16 were available at the time of the experiment and participated. This allowed us, along with the identical setup, to directly compare the results from MI and RI with the ones of the same participants from the previous experiment to minimize between-participants variation. The majority of participants were recruited from the university's student population with the majority been males in their early 20s. We used a within-groups experimental design with every participant performing two trials: (a) one using the EMICS; (b) one using the RI control switcher.

Each participant underwent extensive standardized training before the experiment, similar to our previous work [11, 13], that lasted approximately 20 to 30 min depending on the participant. Each system aspect was introduced gradually and participants were allowed to try them in practice. Additionally, participants were not allowed to proceed with the experimental trials until they had first demonstrated that they could complete a training obstacle course within a specific time limit and with no collisions. This ensured that all participants had adequate understanding of the new system and had attained a common minimum skill level. Counterbalancing was used in the experimental trials (i.e., the order of the tested control switchers was rotated for different participants) to prevent both learning and fatigue effects from biasing the results.

Participants were asked to perform the primary task (robot navigation) as quickly as possible while minimizing 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 prioritize the secondary task over the primary task, and only to perform the primary task if the workload allowed. This is to prevent different operators having different priorities (e.g., some of them focusing on driving while others focusing on the secondary task) and thus minimizing a potential confounding factor. Additionally, as explained in our previous work [11], when people are instructed to do both tasks in parallel to the best of their 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.

The 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 (for the interface see Figure 5(b)). All participants were given an identical and complete 2D map, generated by SLAM prior to the trials.

At the end of each experimental run participants completed a NASA **Task Load Index (TLX)** questionnaire [46]. NASA-TLX is a questionnaire system that enables the perceived workload and difficulty of using a system to be numerically quantified.

4.3 Results

A repeated measures one-way ANOVA was used to compare RI, MI and HI. An ANOVA with Greenhouse-Geisser correction was used for the cases where the sphericity assumption (i.e., that the variances of the differences between trials are not equal) was violated. For HI, the subset of data from Reference [11] corresponding to the 16 participants was used. Although a large period of time (7 months) had passed since their initial participation, we identified learning as a factor that might have affected performance (see Section 4.4). Fisher's **least-significant difference (LSD)** test was used for pairwise comparisons. Post hoc adjustments, such as Bonferroni, were not used in this article, given the (a) clear hypothesis, (b) predefined post-hoc comparisons, (c) small number of comparisons, and (d) early stage of research in this domain (i.e., to avoid missing important findings or interesting insights due to inflated type II error [41, 45, 49]).

We consider a result to be significant when it yields a p value less than 0.05. Last, we report on the statistical power of the results and the effect size. The detailed statistical calculations are

Table 2. Table Showing the ANOVA Results and the Descriptive Statistics for the Metrics Used in the Experiment

metric	ANOVA control mode effect	descriptive statistics
primary task completion time	$F(2, 30) = 19.116, p < 0.01,$ $power = 1, \eta^2 = 0.56$	HI: $M = 418.5$ s, $SD = 34.9$ MI: $M = 374.2$ s, $SD = 15.6$ RI: $M = 392.3$ s, $SD = 22.4$
collisions	$F(1.469, 22.036) = 1, p > 0.05,$ $power < 0.8, \eta^2 = 0.062$	HI: $M = .63$, $SD = 1.1$ MI: $M = 0.38$, $SD = 0.6$ RI: $M = .81$, $SD = 1.5$
primary task score (lower is better)	$F(2, 30) = 16.774, p < 0.01,$ $power > 0.9, \eta^2 = 0.528$	HI: $M = 424.8$, $SD = 36.2$ MI: $M = 378$, $SD = 18.7$ RI: $M = 400.4$, $SD = 22.9$
secondary task completion time	$F(3, 45) = 10.344, p < 0.01,$ $power > 0.95, \eta^2 = 0.408$	HI: $M = 40$ s, $SD = 11.1$ MI: $M = 34.5$ s, $SD = 10.4$ RI: $M = 33.2$ s, $SD = 8.6$ baseline: $M = 34$ s, $SD = 6$
secondary task errors	$F(3, 45) = 0.891, p > 0.05,$ $power < 0.8, \eta^2 = 0.056$	HI: $M = 1.7$, $SD = 1.67$ MI: $M = 1.4$, $SD = 1.4$ RI: $M = 2$, $SD = 1.67$ baseline: $M = 1.7$, $SD = 1.7$
NASA-TLX scores	$F(1.337, 20.049) = 10.135, p < 0.01,$ $power > 0.9, \eta^2 = 0.403$	HI: $M = 43.561$, $SD = 13.026$ MI: $M = 33.217$, $SD = 13.399$ RI: $M = 28.899$, $SD = 9.642$
percentage of time spent in autonomy	$F(2, 25) = 0.74, p > 0.05,$ $power < 0.8, \eta^2 = 0.50$	HI: $M = 57.7\%$, $SD = 25.9$ MI: $M = 53\%$, $SD = 23.25$ RI: $M = 49.1\%$, $SD = 12.22$
number of LOA switches	$F(2, 30) = 16.78, p < 0.01,$ $power > 0.9, \eta^2 = 0.528$	HI: $M = 10.4$, $SD = 8.64$ MI: $M = 13.6$, $SD = 11.78$ RI: $M = 3.13$, $SD = 1.15$

reported in Table 2. In all graphs throughout the article, the error bars indicate the standard error and the mean, the dots the data points, and the violin plots the smoothed probability density of the data. Also, in the graphs, the results of the pairwise comparisons are plotted when significant or relevant with a * for $p < 0.05$, ** for $p < 0.01$, *** for $p < 0.001$, and with *ns* when no significant result was found.

ANOVA for *primary task completion time* (see Figure 6(a)) showed overall significantly different means between HI variable-autonomy, RI and MI. Pairwise comparison revealed that HI performed significantly worse (i.e., slower completion time) than the other two control modes with $p < 0.01$. Also MI variable autonomy performed significantly better than RI ($p < 0.01$).

The effect of control mode on the number of *collisions* was not significant between RI, MI and HI variable autonomy mode.

We used a *primary task score* to be able to capture the speed-accuracy trade-off that different operators might have (i.e., how fast an operator is driving the robot versus how carefully). The primary task score (see Figure 6(b)) is calculated by adding a time penalty of 10 s for every collision, onto the primary task completion time for each participant. Hence, lower score is better. This is inspired by the performance scores used in the RoboCup competitions [25]. ANOVA analysis showed that control mode had a significant effect on primary task score. The LSD tests suggests

Fig. 6. **6(a)**: Primary task time-to-completion. **6(b)**: Primary task score combining time and collisions penalty. MI outperformed significantly both HI and RI in the primary task.

that HI variable autonomy is significantly ($p < 0.01$) worse than the MI mode and the RI mode ($p = 0.01$). Last, MI significantly ($p < 0.01$) outperformed RI regarding primary task score.

The average time per trial that the participants took to complete one series of the 3D object rotations, is denoted by the *Secondary task completion time* (see Figure 7(b)). ANOVA showed that there is a significant difference between the means of secondary task completion times. HI variable autonomy mode performed significantly worse (i.e., more time to complete) than the other modes with $p < 0.01$. Performance between the baseline trial, the MI mode and RI mode were without any significant difference ($p > 0.05$, i.e., same level of performance).

The *number of secondary task errors* (see Figure 7(a)) is the average number of mistake/errors per trial that the participants made during one series of the 3D object rotations. The number of secondary task errors did not have significant differences between the different control modes as shown by ANOVA.

Control mode had a significant effect on *NASA-TLX scores* (see Figure 8(b)) as shown by ANOVA. Pairwise comparisons showed that RI was perceived by participants as having the lowest difficulty, as compared to HI variable autonomy mode with $p < 0.01$ and MI mode with $p = 0.05$. HI variable autonomy is perceived as being more difficult than MI ($p < 0.01$).

Similar to our previous work [10], we analyzed the average *percentage of time spent in autonomy* for each controller. Due to corrupted data for one of the participants, data from 15 participants was used for every result reported on this particular metric. This is contrary to the rest of the analysis in which data from all 16 participants was used. ANOVA did not show any effect caused by control mode in the average percentage of time spent in autonomy.

The *number of LOA switches* (see Figure 8(a)) performed in each trial denotes the frequency in which operators switched LOA in HI; frequency that operators and the EMICS switched LOA in MI; and frequency in which the RI control switcher (i.e., the restricted EMICS) switched LOA in RI. ANOVA showed that control mode had a significant effect on the number of LOA switches. Pairwise comparisons showed that RI had significantly ($p < 0.01$) fewer LOA switches as compared to HI mode and MI mode. The number of LOA switches of HI and MI control modes were on the

Fig. 7. **7(a)**: Secondary task number of mistakes/errors. **7(b)**: Secondary task completion time.

same level (i.e., no statistical difference, $p > 0.05$). Of the 13.6 MI LOA switches, 1.3 were initiated from the EMICS.

We performed a correlation analysis using a two-tailed Pearson's r to investigate any relationships between different metrics and other variables. The *number of LOA switches* in HI and the *number of LOA switches* in MI are highly correlated ($r(14) = 0.849, p < 0.01$). Similar to HI (see Reference [10]) there was no correlation between the *number of LOA switches* and performance in the *primary task score* or *secondary task completion time* in RI and in MI.

No correlation was found between the *percentage of time spent in autonomy* in HI and MI ($r(13) = 0.202, p > 0.05$); and between the *percentage of time spent in autonomy* in HI and RI ($r(13) = 0.094, p > 0.05$). No correlation was found between MI and RI *percentage of time spent in autonomy* ($r(13) = 0.326, p > 0.05$).

The *percentage of time spent in autonomy* LOA is positively correlated with the *primary task score* in RI ($r(13) = 0.528, p < 0.05$). Also the same metrics are positively correlated in MI ($r(13) = 0.627, p < 0.05$). As expected given results in Reference [10], no correlation was found between *primary task score* in HI and *percentage of time spent in autonomy* in HI ($r(13) = 0.047, p > 0.05$).

Last, no correlation was found between *time spent in autonomy* and *secondary task completion time* for any of the three control modes, HI (as found in Reference [10]), MI, and RI.

4.4 Discussion

Mixed-Initiative (MI) and **Robot-Initiative (RI)** outperformed **Human-Initiative (HI)** in terms of primary task performance. What this shows primarily, in the context of this experiment, is that the fuzzy robot control switcher (i.e., both in MI and in RI) is capable of successfully measuring human-robot system performance, inferring if a switch in the LOA is needed, and initiating a LOA switch. This is particularly true given that RI performs better than HI, meaning that the fuzzy control switcher initiative is at least as good as operator's judgment in switching LOA (i.e., compared to the HI of Reference [11]) in the context of this experiment. Any potential learning effects regarding timely LOA switching in RI did not affect the results, since the operator was not allowed to switch LOA (i.e., in RI). One possible explanation for the fact that MI outperforms RI is learning effects in the primary task and the LOA switching. If that is the case, then it is

Fig. 8. **8(a)**: Number of LOA switches per control mode. **8(b)**: NASA-TLX score showing the overall trial difficulty/workload as perceived by the operators.

due to the fact that we used the same participants in an identical experimental setup. However, given the long time difference between the current and the previous experiment in Reference [11] (7 months difference), this is not the only possible explanation. Another explanation might be that the combination of human and EMICS LOA switching capabilities that MI offers is beneficial to performance.

In terms of secondary task performance, MI and RI outperform HI. The secondary task completion time for MI and RI is faster than HI and on the same level of performance as the baseline condition (i.e., secondary task conducted in isolation from the primary task). This can be explained by two possible causes: (a) secondary task learning effects; (b) the fact that operators were feeling confident enough about the robot to neglect it and focus on the secondary task, given its capabilities to take initiative and alert the operator. The latter is reinforced by anecdotal evidence as an informal discussion with the participants revealed that most of them trusted the EMICS to take control and progress toward the navigation goal if needed. As one of the participants put it “even if I completely neglect the robot, at least it will do something meaningful.”

NASA-TLX showed that RI was perceived as marginally easier (i.e., less workload) than MI; and both MI and RI were perceived as significantly easier than HI. This can be an indication of the extra cognitive overhead that operators might have when they need to switch LOA based on judgment (e.g., HI). It is also an indication that having operators knowing that the EMICS will help them if needed, makes the task seem as easier. This is particularly true about the RI, as reflected in the NASA-TLX scores. Operators were restricted not to switch LOA themselves and thus, were not concerned about the LOA switching, just complied with EMICS’s initiative. This is in contrast to MI in which operators still had to do some thinking about switching LOA. This might provide another explanation for the RI outperforming MI in the primary task, that is MI potentially providing beneficial engagement and workload levels (i.e., not too low to degrade engagement and not too high to degrade performance).

A human-factors explanation for the improved performance in MI and RI, both in primary and secondary tasks, as well as for reduced workload, may come from the operator’s attentional model developed by Johnson et al. [28]. They model attention, workload, and performance across

automation mode transitions in aviation and aerospace systems. In the HI condition operators have the extra cognitive overhead of having to switch LOA based on their judgment. MI and RI conditions reduce this overhead. As a result, the operator has more cognitive resources free, which in turn can be re-allocated more efficiently via attention in the task at hand after a LOA switch. This improves performance as predicted by the model of Johnson et al. [28].

Regarding the number of LOA switches in MI and HI, not only they were on the same level but they were also highly and positively correlated. The positive correlation means that participants followed similar patterns in the *number of LOA switches* both in HI and in MI. For example, participants that switched LOA very frequently in HI, also switched LOA very frequently in MI and vice versa for low frequency. This suggests that operators follow the same style of LOA switching regardless of the fact that the EMICS can take initiative. Taking into account the findings of Reference [10], this might be an indication that operators switch LOA based on reasons beyond performance such as personality traits or preferences. Furthermore, the idea that some of the operator's LOA switching is redundant, is further reinforced by the fact LOA switches in RI were much less frequent than HI and MI.

Last, on average, of the 13.6 LOA switches in MI, 1.3 were initiated by the EMICS. This means that despite operators' proven LOA switching capabilities and learning effects, there are circumstances in which the EMICS is successfully contributing by taking initiative.

5 EXPERIMENT 2: EVALUATION USING REAL ROBOT AND TEST ENVIRONMENT

The experiment presented in this section was designed to investigate how well the EMICS could generalize. Generalization in this context means that the EMICS should be able to function well (i.e., deliver similar performance to previous evaluation) in a different setting and under different conditions than the ones that led to its design (e.g., in the previous simulated environment). This is achieved by conducting an experiment with a real robot in a less controlled setting. More specifically the experiment aims to: (a) provide evidence that the EMICS can generalize; (b) factor out any learning effects that might have affected the results of the previous experiment; (c) shift our experimental paradigm toward more complex, realistic environments and tasks.

5.1 Experimental Setup—Apparatus, Robot Test Arena, and Control Modes

In the experiment described here, we used identical software; variable autonomy control switchers; interface (see Figure 12(a)); and OCU (see Figure 9(b)) to our previous experiment in simulation (see Section 4). The robot used was a Pioneer-3DX equipped with a laser range finder and a camera (see Figure 9(a)). Operators controlled the robot remotely, from a separate location, via the control interface. Any **situation-awareness (SA)** was solely gained from the control interface. The communications link between the robot and the OCU was achieved via WiFi.

Our system offers the same two LOAs as in our previous experiment: **teleoperation** and **autonomy**. Three different control modes were tested in the experiment described here: (1) pure **teleoperation**, in which the operator was restricted to using only teleoperation LOA; (2) **Human-Initiative (HI)**, in which the operator could dynamically switch between the teleoperation and autonomy LOA using a button press; (3) **Mixed-Initiative (MI)**, in which both the operator and the EMICS had the ability and authority to dynamically switch between autonomy and teleoperation LOA.

Part of the first floor of School of Computer Science, University of Birmingham building, was used as the arena for this experiment (see Figure 12(b) and Figure 11(a)). More specifically, a long corridor (see Figure 10(a)), 2 offices (see Figure 10(b)) and an open space of approximately 101 square meters in total were used. The experiment took place on weekends and off hours to prevent human activity in the building being a confounding factor.

Fig. 9. 9(a): Pioneer 3DX robot used in the experiment. 9(b): Operator Control Unit (OCU), composed of a laptop, a joypad, a mouse, and a screen showing the control interface. The same OCU was used in all variable autonomy experiments. Note the floor plan in front of the screen, used for the secondary task.

Fig. 10. The first floor of School of Computer Science, University of Birmingham building was used as the arena for the USAR experiment. 10(a): Long corridor that connected the search areas (i.e., offices). 10(b): One of the offices used as a search area.

5.2 Tasks and Performance Degradation Factors

The overall theme of the experiment was an **Urban Search and Rescue (USAR)** scenario. In this scenario the robot operator had to remotely control the robot in the search zone (i.e., the building) and identify the positions and the statuses of victims and potential hazards. As is often the case in real operations, we assume that some prior knowledge of the building (e.g., a map and/or potential points of interest) is given to the robot operator from the authorities or mapped by a UAV. This is represented in the experiment by the SLAM map that appears on the interface (see Figure 12(b)) and by the floor plan used in the secondary task, as kept in the university's records (see Figure 11(a)).

The primary task was to navigate the robot between the different areas in the map in a predefined order. From point A to points B, C, D and then back to point A (see Figure 12(b)). In every one of those areas/points (excluding A) one victim and one hazard sign was placed. The victims were represented by stuffed animals. A meerkat represented an alive victim and a teddy-bear a dead victim (see Figure 13). Hazards were represented by three commonly used hazard signs for flammable materials, radiation, and bio-hazard (see Figure 13). Inside the search areas the robot

Fig. 11. **11(a)**: Floor plan, as kept in the university records, of the area that the experiment took place. This floor plan was printed and given to participants for the secondary task. **11(b)**: Floor plan of the experiment area annotated by an operator during the secondary task. The handwritten annotations of the operator are not to be confused with the blue letters overlaid to denote the search areas for clarity in the figure.

had to stop in the center of the room. Then the operator would have to identify and memorize the position and status of the victim and the sign. Both the victim and the sign were visible from the center of the area and in 360 degrees around the robot (i.e., the victim and the sign could be anywhere around the central position, but visible); no further exploration was needed.

In addition to the primary task, the operator had to perform a secondary task every time the robot exited one of the search areas. This secondary task was designed to induce additional workload in the operator and degrade performance on the primary task. A pen and a paper displaying the floor plan was placed in-front of the operator (see Figure 9(b)). When the operator was asked to perform the secondary task by the experimenter, they had to sketch on the floor plan paper (see Figure 11(b)): (a) the position of the victim denoted by a small x; (b) the status of the victim denoted by a letter (A for alive, D for dead); (c) the position of the hazard denoted by a small o; (d) the status of the hazard denoted by a letter (R—radiation, F—flammable, and B—biohazard); (e) the path that the robot followed from the previous visited point/area. The pen and the floor plan were placed in front of the operator and in a central position (see Figure 9(b)). This way individual differences regarding handedness (i.e., left- versus right-handed participants) did not bias the results. Similar tasks to the one described here are typically required from robot operators in real world disaster response [34] as they are asked to sketch similar information for the SAR team. The speed and ability of the operators to annotate hazards, victims and paths, along with their correct status and respective positions, were the measures of secondary task performance.

Fig. 12. **12(a)**: Control interface as presented to the operator. **Left**: Video feed from the camera, the control mode in use and the status of the navigation goal. **Right**: Map showing the position of the robot (blue footprint and red arrow), the current goal (blue arrow), the AI planner path (green line), the obstacles' laser reflections (red) and the walls (black). **12(b)**: SLAM map of the arena as displayed to the human operator on the interface. Operators had to navigate in turn from point A to B to C to D and then back again to point A.

Fig. 13. The hazard signs used to denote flammable materials, bio-hazard and radiation risk. On the bottom right the stuffed animals representing the victims of the USAR scenario. A meerkat represents a victim that is alive, and a teddy-bear represents a victim that is dead.

The robot performance was degraded by two different degradation factors. The first factor was a box placed in one of the offices. This box was not part of the prior knowledge (i.e., the box was not in the SLAM map). The box also narrowed the entrance to the office. For these two reasons navigation performance would degrade (i.e., the robot is either moving very slowly or is stuck), as

the robot would try to plan a new path through this very tight passage and not through the obstacle. The second performance degradation factor was naturally occurring noise in the laser sensor in certain parts of the arena. This noise was due to shiny surfaces deflecting the laser's beams and it was not controlled by the experimenter. However, due to the high and semi-systematic frequency of its appearance in the specific area of the map, it was adopted as part of the experimental design. Both of these factors, naturally occurring laser noise and unknown obstacles, often occur in real scenarios as the environments are dynamic. This adds to the realism of the experiment.

Last, in one of the offices the WiFi signal was weak. As a result, delays in the control commands and in SA updates in the interface (e.g., location in map, video feedback, etc.) occurred. This was systematic throughout all of the participants and trials. Hence, it did not constitute a confounding factor, while contributing further to the realism of the experiment.

5.3 Participants and Experimental Design

A total of 12 volunteers participated with the majority been males in their early 20s as they were recruited from the School of Computer Science's population. A within-groups experimental design was followed as every participant performed one trial for each of the three control modes (i.e., teleoperation, MI, and HI). The order of the three trials was rotated between all the different permutations for different participants. This counterbalancing technique was used to prevent learning and fatigue effects from introducing confounding factors to the results, since every participant performed all three trials. Additionally, for the secondary task, the signs, the victims, and their positions were randomized in every trial. Again, this was to eliminate any learning effects. A prior experience questionnaire showed, as in our previous experiments, that the majority of the participants were experienced in playing video games. Moreover, 8 of the 12 participants participated in our previous experiment.

Participants underwent extensive standardized training, similar to our previous experiment. The training lasted approximately 20 to 30 min depending on the participant. Due to space constraints a sub-region of the experiment's area was used for the training. In order for the participants to proceed with the experiment they had first to demonstrate their abilities by completing three standardized test trials, one for every control mode. These trials mimicked the actual experimental trials (i.e., same primary and secondary tasks). The training and the test trials ensured that all participants had attained a common minimum skill level in operating the robot.

Participants were instructed to perform the primary task (controlling the robot to search the areas) as quickly and safely (i.e., avoiding collisions) as possible. Additionally they were told that when instructed to perform the secondary task (i.e., annotating information regarding victims, hazards, and paths), 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. Last, participants were told that the best performing individuals in each trial would be rewarded with an extra gift voucher. However, they were not informed about how the best performing participants would be decided, as it would have had them biased toward specific factors. The purpose of this extra gift voucher was to provide an incentive for participants to achieve the best performance possible on both primary and secondary task.

At the end of each trial, participants had to complete an online **NASA Task Load Index (NASA-TLX)** questionnaire. As in our previous experiments, NASA-TLX was used to rate the level of difficulty and workload the participants experienced during each trial.

5.4 Results

A repeated measures one-way ANOVA was used. The independent variable was the control mode with three levels: teleoperation, HI, and MI. In the cases that sphericity assumption was violated,

Table 3. Table Showing ANOVA Results and Descriptive Statistics for All the Metrics

metric	ANOVA control mode effect	descriptive statistics
primary task completion time	$F(2, 22) = 7.382, p < 0.01,$ $power > 0.9, \eta^2 = 0.402$	HI: $M = 327.9$ s, $SD = 23.38$ MI: $M = 349.5$ s, $SD = 35.95$ teleop: $M = 366.3$ s, $SD = 36.26$
collisions	$F(2, 22) = 0.786, p > 0.05,$ $power < 0.8, \eta^2 = 0.067,$	HI: $M = 0.08, SD = .289$ MI: $M = 0.08, SD = .289$ teleop: $M = 0.25, SD = 4.52$
primary task score	$F(2, 22) = 8.724, p < 0.01,$ $power > 0.9, \eta^2 = 0.442$	HI: $M = 328.8, SD = 22.42$ MI: $M = 350.3, SD = 34.9$ teleop: $M = 368.8, SD = 35.57$
secondary task completion time	$F(2, 22) = 0.68, p > 0.05,$ $power > 0.85, \eta^2 = 0.058$	HI: $M = 54$ s, $SD = 12.86$ MI: $M = 54.8$ s, $SD = 20.16$ teleop: $M = 60.3$ s, $SD = 16.37$
secondary task errors	$F(2, 22) = 0.789, p > 0.05,$ $power < 0.8, \eta^2 = 0.067$	HI: $M = .83, SD = 1.03$ MI: $M = 1.5, SD = 1.314$ teleop: $M = 1.42, SD = 1.62$
NASA-TLX scores	$F(1.263, 13.896) = 5.001, p < 0.05,$ $power = 0.603, \eta^2 = 0.313$	HI: $M = 25.9, SD = 12.39$ MI: $M = 29.9, SD = 14.94$ teleop: $M = 38.9, SD = 17.4$
number of LOA switches	$t(11) = -4.076, p < 0.01$	HI: $M = 9.08, SD = 3.42$ MI: $M = 15.08, SD = 5.73$ EMICS in MI: $M = 3.42, SD = 1.83$

For the number of LOA switches, the t-test result is reported.

a Greenhouse-Geisser correction was used with the ANOVA. For pairwise comparisons after a significant ANOVA result, Fisher's **least-significant difference (LSD)** test was used to determine the conditions that differed. Similar to the rest of the article, we consider a result to be significant when it yields a p value less than 0.05. We also report on the statistical power of the results and on the effect size using η^2 . The detailed statistical calculations are reported in Table 3.

ANOVA for *primary task completion time* (see Figure 14(a)) showed overall significantly different means between HI variable-autonomy, MI, and teleoperation. Pairwise comparison reveals that HI performed significantly better (i.e., lower mean completion time) than teleoperation ($p < 0.01$). MI primary task completion time was not significantly different from teleoperation and HI ($p > 0.05$).

Similar to our previous experiments the *primary task score* (see Figure 14(b)) was calculated to capture any speed-accuracy trade-offs in the primary task. The primary task score was calculated by adding a time penalty of 10 sec for every collision onto the primary task completion time for each participant. ANOVA analysis showed that control mode had a significant effect on the primary task score. As expected due to the small number of collisions, LSD pairwise tests showed very similar results with the primary task completion time. The HI control switcher significantly ($p < 0.01$) outperformed pure teleoperation. The primary task score for MI control switcher was in the same level (i.e., no statistical difference found) as pure teleoperation and HI ($p > 0.05$).

Secondary task completion time refers to the total time per trial that the participants took to complete the full annotation in the floor footprint. A full annotation is defined as a sketch that has annotated positions, statuses and paths for all the three search areas. ANOVA did not suggested a significant difference between the mean secondary task completion times for the different control modes.

Fig. 14. **14(a)**: Primary task mean time-to-completion. **14(b)**: Primary task score combining time and collisions penalty. HI significantly outperformed teleoperation in primary task.

The *secondary task number of errors* was measured. A position error was defined an annotation not representing the true position of a victim or hazard with relative accuracy. A status error was defined an annotation not representing the correct status, e.g., victim annotated alive (i.e., meerkat) when is it in-fact dead (i.e., teddy-bear). A path error was defined: (a) a path that was not annotated; (b) a path that was heavily colliding with the walls in the floor plan; (c) a not accurately depicted path. In a real situation, the rescue team should be able to find the victims that are alive, be aware of any hazards and their nature. This should happen by following the path and the position annotations provided and by using common sense regarding space orientation. This is represented in our experiment by the secondary task errors. No significant differences were observed between the different control modes with respect to the number of secondary task errors according to ANOVA.

Control mode had a significant effect on NASA-TLX scores (see Figure 15(b)) as suggested by ANOVA. Pairwise comparisons showed that teleoperation was perceived as harder (i.e., more workload) compared to HI with $p < 0.05$ and marginally harder than MI with $p = 0.05$. HI variable autonomy is perceived as having the same difficulty as MI ($p > 0.05$).

The mean number of LOA switches in HI and in MI were compared using a paired samples t-test (see Figure 15(a)). The reason for using a paired samples t-test is that the experiment had a within-groups design, i.e., we expect some correlation in the results given that the same participants performed both conditions. The mean number of LOA switches in MI was significantly higher than the number of LOA switches in HI, as shown by the t-test. Using Pearson's correlation, no correlation was found in mean LOA switches between HI and MI ($r(10) = 0.472$, $p > 0.05$). Again, the means of HI and MI are significantly different. No correlation was found between the *primary task completion time*, *secondary task completion time*, NASA-TLX score, and the *number of LOA switches* in MI.

The mean number of LOA switches in MI due to EMICS's initiative is $M = 3.42$, ($SD = 1.83$). This number of LOA switches due to EMICS's initiative is positively correlated ($r(10) = 0.681$, $p < 0.05$) with the *primary task completion time*. No correlation was found between the *secondary task*

Fig. 15. **15(a)**: Number of LOA switches during HI, MI, and EMICS initiated LOA switches during MI. **15(b)**: NASA-TLX score showing the overall trial difficulty-workload as perceived by the operators.

completion time, *NASA-TLX score* and the mean number of LOA switches in MI due to EMICS's initiative.

5.5 Discussion

Regarding performance in the primary task, the HI control switcher outperformed teleoperation both in terms of score and time-to-completion. This evidence further reinforces our prior findings [11] that HI variable autonomy outperforms individual LOAs such as teleoperation. Of particular importance is that this evidence came from conducting a scenario in which we aimed for realism and did not restrict potential degradation factors (e.g., naturally occurring noise or communication issues), while at the same time having a controlled experiment.

In contrast to our previous experiment (see Section 4), results from the use of MI control switcher are more difficult to interpret. A trend can be seen as MI performed better than teleoperation in the primary task, however the result was not statistically significant. A comparison between HI and MI does not lead to reliable conclusions, as no statistical difference was found. We believe this is due to conflicts between the operator's and the EMICS's initiative regarding LOA switching. This conflict arose from a restriction that the MI control switcher has. It is designed on the assumption that the agent who is in control (e.g., the MI system or the operator) follows the path yielded from the expert planner relatively closely. Our experiment was designed to control for exploration strategies by restricting operators in only visiting the center point of the offices/areas, as victims and hazards were visible from that point. However, there were occasions that operators decided to engage in some exploration or follow a less restricted path (i.e., compared to the one the expert planner yielded). For example, operators that decided to move closer to a hazard sign to see the letters more clearly or to improve lighting conditions on the sign. Another example is provided by the operators that struggled to pass through the narrow passage created by the unseen object. In such cases the EMICS inferred a performance drop or a deviation from the navigation goal (via the goal directed motion error); as a result the EMICS switched to autonomy. At the same time, if operators had not yet finished their action, they would switch back to having control (i.e., teleoperation). This is further reinforced by participants' feedback. Many of them noted that they

generally trusted the EMICS as there were cases in which it switched LOA in a meaningful way. However, they felt restricted from driving freely and they also felt the EMICS was intrusive at times. Others noted that they felt the MI capabilities were redundant, given that HI control switcher was easy enough to use.

Every time a LOA switch takes place, the EMICS re-initializes the error exponential moving average for a time period of a couple of seconds. During this period the control switcher cannot initiate a LOA switch. In essence, this period of time acts as a minimum time between EMICS initiated switches. Despite this period, the conflict for control was not avoided, suggesting that it is happening on a longer time scale (e.g., several seconds). Instead, this conflict for control between the EMICS and the operator can be avoided to a large extent by MI control switchers that are context-aware. Imagine the following two situations and a MI control switcher that is not context-aware. In the first situation, the robot is idle (i.e., no progress toward the goal), because the operator is neglecting it to perform a secondary task. The performance error measured by the control switcher is the maximum. In the second case the robot is stuck in a corner and the operator is reversing to escape the enclosure. Similar to the previous situation the performance error is the maximum. An MI control switcher that does not take context into account, would initiate a LOA switch in both cases. In the first case, the switch into autonomy would be beneficial as the robot was idle. In the second case, it can potentially lead to a collision as the operator is in the middle of a maneuver. It can also lead to a control conflict as the operator can try to take control back to complete the maneuver. A context-aware control switcher would have been able to distinguish that although in both cases the performance error was large, the situation was different. Hence, in the second case a LOA switch would not have happened. Our EMICS in this specific example, is aware that the large error while the robot is reversing potentially means maneuvering to escape an entrapment (see fuzzy rule base in Section 3.2). However, with cases such the ones discussed in the previous paragraph (e.g., operator performing exploration), our control switcher is not able to cope. We identify context awareness as a major challenge for robotic MI systems. We believe that fuzzy control switchers can help toward this as they can be expanded with expert knowledge regarding context. Last, system transparency is another factor that might positively contribute toward tackling LOA switching conflicts.

The above evidence suggesting a control conflict is in accordance with the fact that MI had significantly more LOA switches compared to HI. In contrast to the results of Section 4 in which LOA switches of MI and HI were highly correlated, in the experiment reported here no correlation was found. This further reinforces the notion that possibly these extra LOA switches stem from the robot-operator conflict for control as they would switch LOA back and forth. Similarly, a number of LOA switches initiated from the EMICS might be due to this control conflict. Possible evidence might come from the positive correlation found between the mean number of EMICS's initiated LOA switches in MI and the primary task completion time. The more LOA switches the EMICS initiated, the more time it took participants to complete the primary task. This is either an indication of a conflict for control degrading performance, or it means that the EMICS is more active in switching LOA for the under-performing participants (i.e., they need more help from the EMICS).

Further HRI studies are required to investigate the phenomenon better as this conflict is a major challenge to overcome for MI systems. Last, the fact that the number of LOA switches is relatively high for both HI and MI, is in accordance with the findings of our previous two experiments (see Reference [10] and Section 4). These findings suggested that the high number of LOA switches can be due to reasons beyond performance (e.g., personality traits and/or preferences).

Regarding the difficulty/workload of the trials, NASA-TLX showed that teleoperation was perceived as the most difficult control mode compared to HI and MI. This suggests, similar to the findings of Reference [11], that the use of variable autonomy can alleviate operators from the

burden of control. Perceived difficulty for MI and HI was on the same level according to the pairwise comparisons. This is in contrast to the findings of Section 4 in which MI found to be easier than HI. However, due to the low statistical power, the possibility of MI and HI differing is not excluded.

Last, a possible explanation for performance and workload not been fully in agreement with the first experiment can be given by the model of Johnson et al. [28] (similar to the discussion in Section 4.4). At times, operators were allocating much of their cognitive resources to deal with the conflict for control. Hence, part of the performance degradation and workload arises from not been able to re-allocate their attention efficiently to the tasks.

6 LIMITATIONS, INSIGHTS, AND FUTURE WORK

The main limitation of the expert-guided MI control switcher presented here is that is not able to cope with the conflict for control. Two of the assumptions on which the EMICS was based are: (a) the human operator is willing to be handed or to hand over control based on the EMICS initiative; and (b) the agent to which the control will be handed is capable of correcting the task effectiveness degradation. A conflict for control arises when one or both of these assumptions are violated. The main reason for those assumptions being violated was the miss-match at times between what the two agents considered ideal performance. For example, the human was acting with the aim of improving exploration performance or avoiding an unforeseen obstacle, while the MI control system was trying to optimize for a different trajectory between known waypoints and without any context knowledge. Given the evidence presented in detail in the previous sections, we consider this conflict for control the main factor affecting the performance of MI in the second experiment compared to that of the first simulated experiment.

In our first simulated experiment both the primary and the secondary tasks were more restricted compared to the tasks in the second real world experiment. Hence, the operators' actions remained within the experiment's EMICS's working envelope (i.e., relatively simple navigation). This resulted in the EMICS being able to generate a timely LOA switch when needed; and most importantly it did not try to switch LOA when that would be unnecessary or intrusive. This in turn allowed the operators to trust the EMICS. Hence, the two above assumptions held true. In our second physical robot experiment, operators in many cases deviated from the navigational goal that they had given to the MI control system to explore further or move freely. The EMICS lacked the knowledge of the new goal (e.g., next to the hazard sign that the operator wanted to inspect instead of the center of the office), or the context of operator's commands. Due to this lack of knowledge the EMICS would initiate a LOA switch believing that the operator was not moving toward the goal, while instead the navigation goal had been changed. The operators found this intrusive and hence switched LOA, leading to the violation of both assumptions and to performance degradation.

In essence, the conflict for control arises because the EMICS and the operator have the authority to aggressively override each other's actions even in cases when they are not necessarily correct. When the EMICS incorrectly overrides the operator's actions it is due to: (a) an incorrect inference about the operator's intention (e.g., lack of knowledge of the current goal) and/or (b) a lack of contextual knowledge regarding the task and the operator's action (i.e., in which context a specific action is taking place and what this action is trying to optimize). When the operator overrides EMICS actions in the problematic cases, it is mostly because they perceive the EMICS initiative as intrusive and aggressive and hence they do not want to hand over control. This is similar to the findings of Dragan and Srinivasa [15] in shared control. They found that when the AI controller is aggressive (i.e., intrusive) and wrong (i.e., takes a wrong action) it negatively affects performance

and also user preference. However, in shared control any potential conflict for control is of a different nature as the operator cannot directly switch LOA and take over.

6.1 Advancing MI Control

Based on the evidence presented in this work, we propose that the key for advancing MI human-robot systems is advancing their human-robot teaming toward deploying efficiently their complementing capabilities (e.g., LOAs in this case). Toward this end, we propose two main directions for future research: (a) improving context awareness in MI control systems; (b) improving initiative hand-off coordination (i.e., the way control and/or initiative is shifting between agents).

Regarding context awareness in MI control, the evidence presented suggests that this would positively impact the conflict for control. A way to provide context is a MI control system capable of predicting operator's intent. For example, predicting operator's intent in the context of navigation (i.e., the navigational goal that the operator wants to move to) can inform the MI control switcher for an unforeseen change in task context, e.g., a new navigational goal can mean that the operator started exploration. In essence, context in the form of operator's intent, can provide the MI control switcher with the ability to understand that its expert policy for ideal performance does not apply in a specific situation and hence should rely on the human operator's decisions. Intent prediction is also known to improve performance in shared control manipulation tasks [26]. Additionally, semantic understanding of an environment (e.g., semantic mapping) will further improve context awareness and understanding of the operator's intention. Another way to provide context is a MI control system capable of inferring or taking into account the operator's status. For example, understanding if performance degradation is due to a tired operator [9] or if the operator is simply not available due to a secondary task [19] can improve the MI control switcher's assumptions about the capability and availability of the agent to which the control will be given.

Improving initiative hand-off coordination is crucial even if a capable context aware MI control system is assumed. There are always going to be situations in which the agents will have some level of disagreement regarding the LOA choice or actions to be taken. Hence, this disagreement should be resolved without a detrimental effect on performance. Using a negotiation interface, e.g., similar to negotiation theory used to negotiate evasive maneuvers in autonomous car research [44], can potentially enable smooth LOA switches as both agents can negotiate their choice contrary to the current aggressive override. For example, an operator might refuse to hand over control in a middle of a maneuver and similarly the MI control switcher might insist on handing over control in a critical situation. This will improve the assumption that the agent is always willing to hand over or be handed control. Additionally, transparency about the MI control switcher's choice (i.e., communicating to the operator the reasons behind a LOA switch) can further facilitate smoother hand-offs.

Last, a mathematical formalism for MI control in the context of LOA switching, similar to the one of Dragan and Srinivasa for shared control [15] is needed. This, in conjunction with the MI control taxonomy of Reference [27], would give a solid foundation and a common understanding for future research.

7 CONCLUSION AND IMPACT

This article presented an expert-guided approach to designing robotic MI control systems. It also presented an expert-guided Mixed-initiative control switcher and its experimental evaluation using both a high fidelity simulator and in a realistic scenario.

The proposed control switcher uses an online performance metric that represents the expected effectiveness of goal directed motion, with parameters determined by comparing the control switcher to human performance on a prior experiment. The proposed control switcher aims to measure human-robot team performance, infer if a LOA switch is needed, and switch LOA.

Evidence from our initial evaluation showed some of the potential advantages of MI control. The EMICS was found in simulation to outperform HI control in both primary and secondary task performance and in the workload as perceived by the operator. This in turn means that MI can potentially outperform control modes that lack any LOA switching capabilities in situations where using the complementing capabilities of humans and autonomous agents can be beneficial.

The second evaluation extended our experimental framework toward a less controlled and more realistic setting. The EMICS was used in a real robot performing in a USAR scenario. Results regarding performance advantages of MI over HI were less conclusive compared to the initial evaluation. However, the experiment yielded significant new insight into challenges and problems that need to be overcome in the design of MI systems. The difficulty in interpreting the results was partially due to variance and poor statistical power in some of the statistical calculations. However, the major factor that affected the results was the conflict for control situations that arose between the operator and the EMICS. This led operators and the EMICS to confusion on what they should do as the LOA was switching back and forth between autonomy (i.e. AI autonomous navigation in control) and teleoperation (i.e., human in control). We believe this control conflict is one of the major challenges for future MI control research and that the addition of further context-awareness (of task, environment, workload, operator's intent, etc.) is one way to address this problem. However, more generally, MI control has its merits, since it provides redundancy in the cases where an operator might not be able to switch LOA if needed, e.g., during loss of communication with the robot, or a sudden event impairing the operator.

Last, the USAR experiment provided important real world evidence that variable autonomy control in the form of HI has the potential to outperform teleoperation in a navigation task. This is in accordance to our previous findings in Reference [11] and denotes that human operators successfully use HI capabilities to overcome various performance-degrading factors and situations. Additionally, further evidence was provided to support the hypothesis that operators switch LOA for reasons other than performance.

Overall, we believe that this article has made a number of significant contributions to the MI research on mobile robots: (a) proposed a framework for designing MI control systems; (b) proposed an MI control switcher (i.e., the EMICS); (c) highlighted fields of interest and situations that would benefit from MI control while identifying the shortcomings that constitute open challenges for the research field.

APPENDICES

A FURTHER IMPLEMENTATION DETAILS

The following standard (i.e., commonly used) operators are used in the fuzzy inference process: for conjunction (i.e., “and”) the minimum operator is used; for disjunction (i.e., “or”) the maximum operator is used; for rule activation the minimum operator is used; for fuzzy rules aggregation the maximum operator is used. The fuzzification process is achieved using the fuzzy membership functions described in Table 4.

Table 4. The Fuzzy Membership Functions for Linguistic Values of Input Variables “Error” and “Speed” (See Figure 2)

Linguistic value	Membership functions
error small	$\mu_{A_1^1}(x_1) = \begin{cases} 0, & 0.06 \leq x \\ 1 & 0 \leq x \leq 0.035 \\ \frac{0.06-x}{0.06-0.035} & 0.035 \leq x \leq 0.06 \end{cases}$
error medium	$\mu_{A_1^2}(x_1) = \begin{cases} 0, & x \leq 0.045, 0.08 \leq x \\ \frac{x-0.045}{0.055-0.045} & 0.045 \leq x \leq 0.055 \\ 1 & 0.055 \leq x \leq 0.065 \\ \frac{0.08-x}{0.08-0.065} & 0.065 \leq x \leq 0.08 \end{cases}$
error large	$\mu_{A_1^3}(x_1) = \begin{cases} 0, & x \leq 0.065 \\ \frac{x-0.065}{0.085-0.065} & 0.065 \leq x \leq 0.085 \\ 1 & 0.085 \leq x \leq 0.1 \end{cases}$
speed reverse	$\mu_{A_2^1}(x_2) = \begin{cases} 0, & -0.02 \leq x \\ 1 & x \leq -0.03 \\ \frac{-0.02-x}{-0.02-(-0.03)} & -0.03 \leq x \leq -0.02 \end{cases}$
speed zero	$\mu_{A_2^2}(x_2) = \begin{cases} 0, & x \leq -0.03, 0.03 \leq x \\ \frac{x-(-0.03)}{0-(-0.03)} & -0.03 \leq x \leq 0 \\ \frac{0.03-x}{0.03} & 0 \leq x \leq 0.03 \end{cases}$
speed forward	$\mu_{A_2^3}(x_2) = \begin{cases} 0, & 0.02 \leq x \\ \frac{x-0.02}{0.03-0.02} & 0.02 \leq x \leq 0.03 \\ 1 & 0.03 \leq x \end{cases}$

For “error,” trapezoid membership functions have been used. For “speed” membership functions, two trapezoid for “reverse” and “forward” and one triangular for “zero,” were used. The membership functions were heuristically chosen to smoothly overlap throughout the universe of discourse. This is a common practice when designing fuzzy controllers. The membership function for “error large” was chosen based on the error threshold calculated in Section 3.1.

B COMPLEMENTARY FIGURES

Fig. 16. **16(a)**: Secondary task completion time in the real robot experiment (i.e., experiment 2). **16(b)**: Secondary task total number of errors for each trial in the real robot experiment (i.e., experiment 2).

REFERENCES

- [1] Julie A. Adams, Pramila Rani, and Nilanjan Sarkar. 2004. Mixed initiative interaction and robotic systems. In *Proceedings of the AAAI Workshop on Supervisory Control of Learning and Adaptive Systems*. 6–13.
- [2] Michael Baker and Holly A. Yanco. 2004. Autonomy mode suggestions for improving human-robot interaction. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC'04)*, Vol. 3. 2948–2953.
- [3] Robert Goodell Brown. 1963. *Smoothing, Forecasting and Prediction of Discrete Time Series*. Courier Corporation. 1–468.
- [4] David J. Bruemmer, Douglas A. Few, Ronald L. Boring, Julie L. Marble, Miles C. Walton, and Curtis W. Nielsen. 2005. Shared understanding for collaborative control. *IEEE Trans. Syst. Man Cybernet.—Part A: Syst. Hum.* 35, 4 (July 2005), 494–504.
- [5] David J. Bruemmer, Julie L. Marble, Donald D. Dudenheffer, Matthew O. Anderson, and Mark D. McKay. 2003. Mixed-initiative control for remote characterization of hazardous environments. In *Proceedings of the 36th Annual Hawaii International Conference on System Sciences*. 9.
- [6] Jennifer L. Burke, Robin R. Murphy, Michael D. Covert, and Dawn L. Riddle. 2004. Moonlight in Miami: Field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise. *Hum.-Comput. Interact.* 19, 1–2 (2004), 85–116.
- [7] Tom Carlson and Yiannis Demiris. 2008. Collaborative control in human–wheelchair interaction reduces the need for dexterity in precise manoeuvres. In *Proceedings of Robotic Helpers: User Interaction, Interfaces and Companions in Assistive and Therapy Robotics, a Workshop at ACM/IEEE HRI*. 59–66.
- [8] Jennifer L. Casper and Robin R. Murphy. 2003. Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. *IEEE Trans. Syst. Man Cybernet.—Part B: Cybernet.* 33, 3 (2003), 367–385.
- [9] Caroline P. Carvalho Chanel, Raphaëlle N. Roy, Frédéric Dehais, and Nicolas Drougard. 2020. Towards mixed-initiative human–robot interaction: Assessment of discriminative physiological and behavioral features for performance prediction. *Sensors* 20, 1 (2020), 296. <https://doi.org/10.3390/s20010296>
- [10] Manolis Chiou, Goda Bieksaite, Nick Hawes, and Rustam Stolkin. 2016. 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. In *Proceedings of the AAAI Fall Symposium Series: Shared Autonomy in Research and Practice*. 304–310.
- [11] Manolis Chiou, Goda Bieksaite, Rustam Stolkin, Nick Hawes, Kimron L. Shapiro, and Timothy S. Harrison. 2016. Experimental analysis of a variable autonomy framework for controlling a remotely operating mobile robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'16)*. 3581–3588.

- [12] Manolis Chiou, Nick Hawes, Rustam Stolkin, Kimron L. Shapiro, Jess R. Kerlin, and Andrew Clouter. 2015. Towards the principled study of variable autonomy in mobile robots. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC'15)*. 1053–1059.
- [13] Manolis Chiou, Mohammed Talha, and Rustam Stolkin. 2019. Learning effects in variable autonomy human-robot systems: How much training is enough? In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC'19)*. 720–727. <https://doi.org/10.1109/SMC.2019.8914558>
- [14] Lorin D. Dole, David M. Sirkin, Robin R. Murphy, and Clifford I. Nass. 2015. Robots need humans in the loop to improve the hopefulness of disaster survivors. In *Proceedings of the IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN'15)*. 707–714.
- [15] Anca D. Dragan and Siddhartha S. Srinivasa. 2013. A policy-blending formalism for shared control. *Int. J. Robot. Res.* 32, 7 (2013), 790–805. <https://doi.org/10.1177/0278364913490324>
- [16] Douglas A. Few, David J. Bruemmer, and Miles Walton. 2006. Improved human-robot teaming through facilitated initiative. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'06)*. 171–176.
- [17] Alberto Finzi and Andrea Orlandini. 2005. Human-robot interaction through mixed-initiative planning for rescue and search rovers. In *Proceedings of the Congress of the Italian Association for Artificial Intelligence*. 483–494.
- [18] Giorgio Ganis and Rogier Kievit. 2015. A new set of three-dimensional shapes for investigating mental rotation processes: Validation data and stimulus set. *J. Open Psychol. Data* 3, 1 (2015).
- [19] Thibault Gateau, Caroline P. Carvalho Chanel, Mai-huy Le, and Frederic Dehais. 2016. Considering human's non-deterministic behavior and his availability state when designing a collaborative human-robots system. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'16)*. 4391–4397. <https://doi.org/10.1109/IROS.2016.7759646>
- [20] Matthew Gombolay, Anna Bair, Cindy Huang, and Julie Shah. 2017. Computational design of mixed-initiative human-robot teaming that considers human factors: Situational awareness, workload, and workflow preferences. *Int. J. Robot. Res.* 36, 5–7 (2017), 598–617.
- [21] Matthew C. Gombolay, Reymundo A. Gutierrez, Shanell G. Clarke, Giancarlo F. Sturla, and Julie A. Shah. 2015. Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Auton. Robots* 39, 3 (2015), 293–312. <https://doi.org/10.1007/s10514-015-9457-9>
- [22] Michael A. Goodrich, Dan R. Olsen Jr., Jacob W. Crandall, and Thomas J. Palmer. 2001. Experiments in adjustable autonomy. In *Proceedings of the IJCAI Workshop on Autonomy, Delegation, and Control: Interacting with Autonomous Agents*. 1624–1629.
- [23] Benjamin Hardin and Michael A. Goodrich. 2009. On using mixed-initiative control: A perspective for managing large-scale robotic teams. In *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI'09)*. 165–172.
- [24] A. Hong, O. Igharoro, Y. Liu, F. Niroui, G. Nejat, and B. Benhabib. 2019. Investigating human-robot teams for learning-based semi-autonomous control in urban search and rescue environments. *J. Intell. Robot. Syst.* 94, 3–4 (2019), 669–686. <https://doi.org/10.1007/s10846-018-0899-0>
- [25] A. Jacoff, E. Messina, B. A. Weiss, S. Tadokoro, and Y. Nakagawa. 2003. Test arenas and performance metrics for urban search and rescue robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'03)*, Vol. 3. 3396–3403.
- [26] Siddarth Jain and Brenna Argall. 2019. Probabilistic human intent recognition for shared autonomy in assistive robotics. *ACM Trans. Hum.-Robot Interact.* 9, 1 (2019), 1–23. <https://doi.org/10.1145/3359614>
- [27] Shu Jiang and Ronald C. Arkin. 2015. Mixed-initiative human-robot interaction: Definition, taxonomy, and Survey. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC'15)*. 954–961.
- [28] Aaron W. Johnson, Kevin R. Duda, Thomas B. Sheridan, and Charles M. Oman. 2017. A closed-loop model of operator visual attention, situation awareness, and performance across automation mode transitions. *Hum. Factors* 59, 2 (2017), 229–241.
- [29] Yugang Liu and Goldie Nejat. 2016. Multirobot cooperative learning for semiautonomous control in urban search and rescue applications. *J. Field Robot.* 33, 4 (2016), 512–536. <https://doi.org/10.1002/rob.21597>
- [30] E. H. Mamdani and S. Assilian. 1975. An experiment in linguistic synthesis with a fuzzy logic controller. *Int. J. Man-Mach. Studies* 7, 1 (1975), 1–13.
- [31] Vikram Manikonda, Priya Ranjan, Zachary Kulis, and Calhoun Drive. 2007. A mixed initiative controller and testbed for human-robot teams in tactical operations. In *Proceedings of the AAAI Fall Symposium*. 92–99.
- [32] Julie L. Marble, David J. Bruemmer, Douglas A. Few, and Donald D. Dudenhoeffer. 2004. Evaluation of supervisory vs. peer-peer interaction with human-robot teams. In *Proceedings of the 37th Hawaii International Conference on System Sciences*. 1–9.

- [33] Eitan Marder-Eppstein, Eric Berger, Tully Foote, Brian Gerkey, and Kurt Konolige. 2010. The office marathon: Robust navigation in an indoor office environment. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'10)*. 300–307.
- [34] Robin R. Murphy. 2004. Human-robot interaction in rescue robotics. *IEEE Trans. Syst. Man Cybernet. Part C (Appl. Rev.)* 34, 2 (May 2004), 138–153.
- [35] Robin R. Murphy and Jennifer L. Burke. 2005. Up from the rubble: Lessons learned about HRI from Search and Rescue. In *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*. 437–441.
- [36] Selma Musić and Sandra Hirche. 2017. Control sharing in human-robot team interaction. *Annu. Rev. Control* 44 (2017), 342–354. <https://doi.org/10.1016/j.arcontrol.2017.09.017>
- [37] Keiji Nagatani, Seiga Kiribayashi, Yoshito Okada, Kazuki Otake, Kazuya Yoshida, Takeshi Nishimura, Tomoaki Yoshida, Eiji Koyanagi, Satoshi Tadokoro, Mineo Fukushima, and Shinji Kawatsuma. 2013. Emergency response to the nuclear accident at the Fukushima Daiichi nuclear power plants using mobile rescue robots. *J. Field Robot.* 30, 1 (2013), 44–63.
- [38] Farrukh Nagi, L. Perumal, and J. Nagi. 2009. A new integrated fuzzy bang-bang relay control system. *Mechatronics* 19, 5 (2009), 748–760.
- [39] Curtis W. Nielsen, Douglas A. Few, and Devin S. Athey. 2008. Using mixed-initiative human-robot interaction to bound performance in a search task. In *Proceedings of the IEEE International Conference on Intelligent Sensors, Sensor Networks and Information Processing*. 195–200.
- [40] Adam Norton, Willard Ober, Lisa Baraniecki, Eric McCann, Jean Scholtz, David Shane, Anna Skinner, Robert Watson, and Holly Yanco. 2017. Analysis of human–robot interaction at the DARPA Robotics Challenge Finals. *Int. J. Robot. Res.* 36, 5–7 (2017), 483–513. <https://doi.org/10.1177/0278364916688254>
- [41] Thomas V. Perneger. 1998. What’s wrong with Bonferroni adjustments. *BMJ Clin. Res. Ed.* 316, 7139 (1998), 1236–8.
- [42] Juan Rada-Vilela. 2013. fuzzylite: A fuzzy logic control library in C++. In *Proceedings of the Open Source Developers Conference*.
- [43] Marc Rigter, Bruno Lacerda, and Nick Hawes. 2020. A framework for learning from demonstration with minimal human effort. *IEEE Robot. Autom. Lett.* 5, 2 (2020), 2023–2030. <https://doi.org/10.1109/LRA.2020.2970619>
- [44] Simon Rothfuss, Robert Schmidt, Michael Flad, and Soren Hohmann. 2019. A concept for human-machine negotiation in advanced driving assistance systems. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC'19)*. 3116–3123. <https://doi.org/10.1109/SMC.2019.8914282>
- [45] Kenneth J. Rothman. 1990. No adjustments are needed for multiple comparisons. *Epidemiology* 1, 1 (1990), 43–46.
- [46] David Sharek. 2011. A useable, online NASA-TLX tool. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. 1375–1379.
- [47] Jian Shen, Javier Ibanez-Guzman, Teck Chew Ng, and Boon Seng Chew. 2004. A collaborative-shared control system with safe obstacle avoidance capability. In *Proceedings of the IEEE Conference on Robotics, Automation and Mechatronics*, Vol. 1. 119–123.
- [48] Thomas B. Sheridan and William L. Verplank. 1978. Human and computer control of undersea teleoperators. Technical report, MIT Man-Machine Systems Laboratory.
- [49] David L. Streiner and Geoffrey R. Norman. 2011. Correction for multiple testing: Is there a resolution? *CHEST J.* 140, 1 (2011), 16.
- [50] Holly A. Yanco, Adam Norton, Willard Ober, David Shane, Anna Skinner, and Jack Vice. 2015. Analysis of human-robot interaction at the DARPA robotics challenge trials. *J. Field Robot.* 32, 3 (2015), 420–444.

Received February 2020; revised January 2021; accepted May 2021