

Robot Health Indicator: A Visual Cue to Improve Level of Autonomy Switching Systems

Aniketh Ramesh
Axr1050@student.bham.ac.uk
Extreme Robotics Laboratory,
University of Birmingham
Birmingham, England

Madeleine Englund
Andreas Theodorou
Maen0191@student.umu.se
Andreas.Theodorou@umu.se
Umeå university
Umeå, Sweden

Rustam Stolkin
Manolis Chiou
R.Stolkin@bham.ac.uk
M.Chiou@bham.ac.uk
Extreme Robotics Laboratory,
University of Birmingham
Birmingham, England

ABSTRACT

Using different Levels of Autonomy (LoA), a human operator can vary the extent of control they have over a robot's actions. LoAs enable operators to mitigate a robot's performance degradation or limitations in the its autonomous capabilities. However, LoA regulation and other tasks may often overload an operator's cognitive abilities. Inspired by video game user interfaces, we study if adding a 'Robot Health Bar' to the robot control UI can reduce the cognitive demand and perceptual effort required for LoA regulation while promoting trust and transparency. This Health Bar uses the robot vitals and robot health framework to quantify and present runtime performance degradation in robots. Results from our pilot study indicate that when using a health bar, operators used to manual control more to minimise the risk of robot failure during high performance degradation. It also gave us insights and lessons to inform subsequent experiments on human-robot teaming.

1 INTRODUCTION AND RELATED WORK

During mobile robot navigation tasks, factors like uneven terrain, obstacles, sensor noise etc. can degrade a robot's performance. If left unattended, performance degradation can cause robots to fail or perform sub-optimally [1]. Timely operator intervention or triggering of recovery behaviours can mitigate this. Using Levels of Autonomy (LoA) [2], the extent of control that a human operator has over the robot's actions can be varied. For example during high performance degradation, an autonomous robot can be switched to manual control by an experienced human operator. The LoA can be switched back once the robot is capable of functioning autonomously again.

In Human Initiative (HI) LoA switching systems, the robot operator is in charge of LoA regulation during the task. Compared to fully autonomous robots, manually controlled robots, and other implementations of variable autonomy systems, HI-LoA shows better task performance during remote navigation tasks in unknown environments [3]. However, the additional perceptual and mental effort required to monitor robot operation data and determine if LoA switching is needed, imposes comparatively higher levels of cognitive workload on the operator. While overloading an operator's cognitive abilities for prolonged periods can reduce task performance [4] due to stress, fatigue and varying levels of trust in the system [5, 6], low workload can cause out-of-loop performance problems like complacency and over-trusting the system [7, 8]. Effective HI-LoA system design remains an open problem in

existing literature [3]. A well-designed system should keep operator cognitive workload within a 'sweet spot' [9] (i.e., acceptable levels). LoA switches by operators should be due to a clear understanding of the robot's capabilities and limitations, not due to trust issues. Transparent design, i.e., where the operator can accurately interpret the robot's capabilities, goals, and its progress, have reported effective calibration of trust [10], even with increasing levels of automation [11]. Systems where information about the robot is presented over the graphic/visual modalities are trusted more than the audio and textual modalities [12]. Independent of systems design considerations, significant differences still exist in people's perceived reliability of variable autonomy systems. However, they can be minimised through standardised training [11].

Video games serve as inspiration to understand how real-time information can be presented in a transparent manner, while keeping the operator cognitive workload in the sweet spot. Video game players report that they are receptive to information presented through visual cues as long as they are clear and consistent [13]. Game visual cues are usually colour coded UI elements overlaid on the interface during gameplay, e.g. the 'Health Bar'. In popular online games¹, remotely situated players use these cues to get their teammates attention. Similarly, robot operators can use colour coded visual cues to acquire quick situational awareness [14] while operating a robot with minimum perceptual effort. For robot interfaces, Murphy et al. [15] recommend that visual cues may be added to the robot control UI, without blocking useful information.

The existing literature shows that effective trust calibration, system transparency and standardised training are crucial for a well designed HI-LoA switching system. Information to assist the operator with making decisions about LoA regulation can be presented through properly designed visual cues. Therefore, in this study, we present our pilot experiments to explore how visual cues about a robot's performance degradation can help LoA regulation decisions during high cognitive workload situations. Based on the Robot Vitals and Robot Health framework [1] to quantify performance degradation, we design a colour coded 'Robot Health Bar' UI element. We hypothesise that using the Health Bar will reduce the perceptual effort required by operators to make LoA switching decisions, improve task performance, and result in a lower overall cognitive workload. We also took feedback from the participants about their overall trust in the system, its transparency and recommendation for improving the Health Bar design.

¹See <https://overwatch.blizzard.com/en-us/>

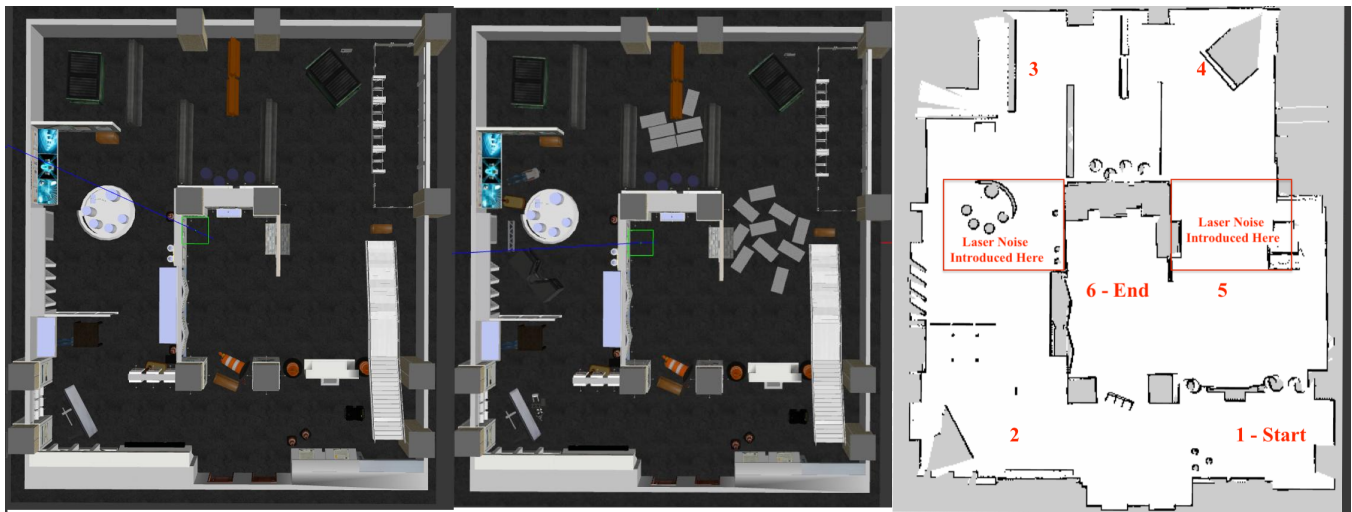


Figure 1: Condition A arena (L to R): 1) Empty, 2) With obstacles and uneven terrain, 3) 2D map with waypoints marked and locations where laser noise was introduced.

2 SYSTEM DESIGN

The experiment consisted of two tasks, similar to the experiments setup by Chiou et al. [3, 16]. The primary task was a mobile robot navigation task using the Clearpath Husky Robot, simulated on Gazebo. As shown in Figure 1, these arenas were designed to mimic urban search and rescue scenarios, and were populated with performance degrading factors commonly found in them like - obstacles, uneven terrain and laser noise. A 2D laser scan of an empty map (Figure 1 - Right) was generated before the performance degrading factors were added. The difference between the map used for robot navigation and the actual arena affects robot navigation planning, thereby adding another performance degrading factor. Two different arenas were created to compare operator performance with (condition A) and without the Health Bar (condition B).

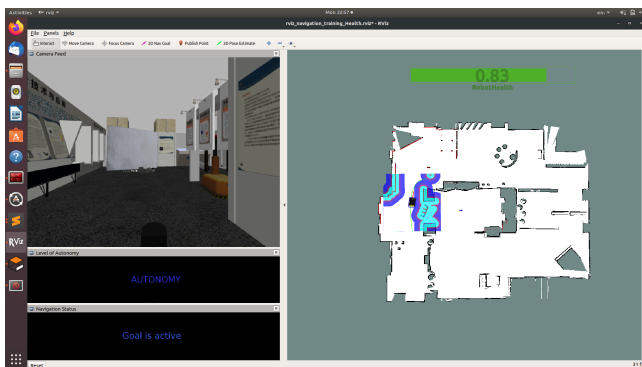


Figure 2: Condition A - Interface with Robot Health Bar. The green colour indicates the robot is 'healthy'

RViz was used to visualise the robot map, sensor data, and give commands to the robot. The standard RViz interface in Figure 3 was used for Condition B, and the interface with a 'Robot Health Bar'

shown in Figure 2 was used for Condition A. The Husky robot used two LoAs - 1) Waypoint-based navigation and 2) Manual Control by an operator using a Joystick. The Robot Health at each instant was calculated using the Robot Vitals and Robot Health framework [1], and the Health Bar was created and displayed using the JSK Visualisations² ROS Package. Here, Robot Health is defined as 'an overall scalar estimate of a robot's ability to carry out its tasks without its capabilities being impaired by any performance degrading factors'. Therefore, an operator can monitor the health to detect situations where a robot is likely to fail, and trigger LoA switches to assist the robot. To improve readability, the Robot Health was standardised to the range [0, 1]. The Red-Amber-Green colour coding convention was used for the Health Bar as shown in Figure 2. Based on preliminary experiments, health above 0.7 was decided as 'Healthy', and the colour of the Health Bar was green. Values between 0.5 - 0.7 were coloured in amber, and values below 0.5 were coloured red. Instead of sharp colour changes, the Health Bar gradually went from green to amber to red. This was done to minimise rapid colour changes due to small fluctuations in health near the threshold values.

The secondary task was a 3D object rotation task [17], used to induce additional cognitive workload for the operator. Here participants were successively presented with two 3D objects. They had to determine whether the objects were the same or different. Instead of having a fixed time limit like in Chiou et al.[16], here the participants did this task throughout the runtime. Both the primary and secondary tasks were displayed on adjacent screens. Participants were explicitly instructed 'to prioritise the robot navigation task and minimise the likelihood of robot failure'. They were also told to simultaneously do the secondary task to the best of their ability, and that they were not being evaluated on their performance in it.

²http://wiki.ros.org/jsk_visualization/

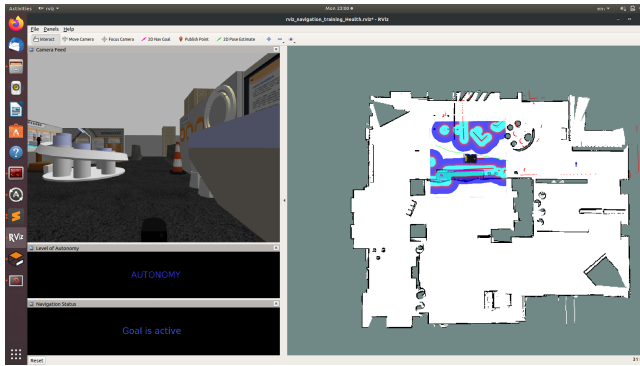


Figure 3: Condition B - Interface without Robot Health Bar

3 EXPERIMENT METHODOLOGY

Eight test subjects participated in the experiment and performed both conditions (i.e. within subjects design). The order of the conditions was counterbalanced to minimise learning and fatigue effects. First, participants had to fill out background information (see appendix A.1). Then they were introduced to basic robot navigation in a training arena similar to the one used in Chiou et al.[16]. All participants trained on this arena till they were able to demonstrate a minimum proficiency in robot navigation. This ensured that confounding factors due to a variation in skill levels were minimised. Next, they were shown the 2 different LoA, how to switch between them and were given some time to practice LoA switching on the training arena.

Participants were then introduced to the interface with the Health Bar and asked open-ended questions without prior explanation (see appendix A.2). After answering the questions, they were given the following instructions regarding the Health Bar: “The Robot Health Bar indicates how much a robot’s performance is degraded by environmental factors. These environmental factors can be anything ranging from bad terrain to laser noise. During low health, the Health Bar will become redder. During high health, the Health Bar will become greener. The lower the robot’s health, the more likely it will fail. You may use the Health Bar to help you determine when the robot requires a LoA switching”. The effect of performance degrading factors on the Health Bar was demonstrated by introducing obstacles and laser noise in the training arena. Participants were also given time to familiarise themselves with the interface before starting the navigation task. Next participants were familiarised only with the secondary task, and allowed multiple practice runs till they felt comfortable with it. Their baseline on the secondary task alone was measured before the experiment.

Before each condition, participants were shown the start and finish points in the navigation task and told that the next waypoint would be assigned automatically on reaching the current one. Then the participants carried out each experimental condition followed by a NASA-TLX form to evaluate the perceived cognitive workload during the task. Lastly, participants completed open-ended questions and a transparency/trust questionnaire after the experimental trials and the NASA-TLX forms (see A.2 and A.3).

Table 1: Summary of Statistical Analysis

Primary Task	Condition A - With Health Bar	Condition B -Without Health Bar	Pairwise T Test, Two Tailed
Completion Time	M 174.38, SD = 30.91	M 175.25, SD = 32.35	T-Score = 2.36, P Value = 0.96
Percentage of the Task Robot was Manually Controlled	M 49.13, SD = 19.97	M 39.77, SD = 17.77	T-Score = 2.36, P Value = 0.08
Percentage of the Task Robot was Unhealthy	M 61.91, SD = 21.80	M 78.03, SD = 21.36	T-Score = 2.36, P Value = 0.17
Percentage of the Task Robot was Manually Controlled when Unhealthy	M 64.59, SD = 21.16	M 45.17, SD = 18.67	T-Score = 2.36, P Value = 0.02*
Secondary Task	Condition A - With Health Bar	Condition B -Without Health Bar	Pairwise T Test, Two Tailed
Average Response Time	M 10.616.18, SD = 3.673.76	M 9.053.49, SD = 3.679.02	T-Score = 2.36, P Value = 0.22
Total Answered	M 20.63, SD = 5.80	M 23.63, SD = 8.35	T-Score = 2.36, P Value = 0.19
Accuracy %	M 90.13, SD = 7.00	M 94.59, SD = 4.98	T-Score = 2.36, P Value = 0.02*
NASA-TLX	Condition A - With Health Bar	Condition B -Without Health Bar	Pairwise T Test, Two Tailed
Mental Demand	M 80.63, SD = 16.78	M 80.08, SD = 17.73	T-Score = 2.36, P Value = 0.89
Physical Demand	M 28.75, SD = 13.30	M 36.88, SD = 25.35	T-Score = 2.36, P Value = 0.19
Temporal Demand	M 66.88, SD = 26.98	M 70.06, SD = 17.11	T-Score = 2.36, P Value = 0.60
Performance	M 43.75, SD = 29.73	M 43.13, SD = 24.63	T-Score = 2.36, P Value = 0.94
Effort	M 58.75, SD = 30.21	M 59.38, SD = 27.83	T-Score = 2.36, P Value = 0.91
Frustration	M 49.38, SD = 28.21	M 53.13, SD = 27.51	T-Score = 2.36, P Value = 0.36

* $p < 0.05$

4 RESULTS

Robot health values under 0.7 were classified as ‘unhealthy’. This criterion was heuristically determined based on previous studies on the robot vitals and robot health framework [1]. Results from the experiments conducted on 8 participants were computed and tested for statistical differences using two-tailed pairwise T-Tests. These results are summarised in table 1.

The total percentage of runtime that the robot was autonomous, manually controlled and ‘unhealthy’ was calculated for each experimental condition. The percentage of run time that the robot health was ‘unhealthy’ showed no statistical differences between the two conditions. Similarly, no significant differences were observed in the percentage of run time the robot was manually controlled. However, the percentage of run time that the robot was manually controlled when it was unhealthy, showed a significant ($p < 0.05$, $t = -2.36$) difference between both conditions. That is, operators manually controlled unhealthy robots for longer when the health of the robot was displayed (i.e., Condition A). As a result, the robot was ‘unhealthy’ for an average of 61.91% of the runtime in condition A, but 78.03% of the runtime in condition B.

The change in perceptual effort required to carry out LoA switching tasks was measured using the operator’s accuracy on the secondary task. When the health bar was not displayed (Condition B), operators showed significantly higher levels of accuracy ($p < 0.05$, $t = 2.36$) on the secondary task. This indicates the perceptual effort required to use the interface with the robot health bar, is higher. However, the NASA-TLX scores showed no significant differences in the overall cognitive workload imposed by both conditions.

Open-ended questions asked before the experiment showed that some participants were slightly confused about what the Robot Health Bar was. One participant said - “The health bar looks like a timer, because it has a number. I thought it meant seconds”. Others successfully grasped the idea behind the UI element. Another participant thought “When the number goes down the robot dies”, indicating a game-like perception of the Health Bar’s function. However, all participants understood it was there to assist or alert an operator performance degradation.

Most participants preferred the interface of condition A (with the ‘Robot Health bar’), finding it intuitive for regulating LoA - “When the health was low, it was better to manually control the robot”. One participant said - “Without the health bar I had to actively use

my brain to detect when the robot needed help". Three participants, however, felt that the colours were sometimes inaccurate and did not match the robot's state. One of them pointed out that while red and green clearly helped indicate which LoA was better, amber was confusing because they did not know what to do. Although not significant, operators that had experience with operating robots (according to the background information) seemed to prefer more transparency in the UI, while the rest preferred less transparency. Novice operators (as reported by the background information), instead, stated that they responded to the colours of the UI to a greater extent than the participants with more experience in operating robots.

Pearson's correlation showed a significant negative trend between how often the participants operated remote-controlled vehicles and how easy it was to know when to change LoA for alternative 1 (Figure B.3a) in the questionnaire ($p < 0.01$). Significant positive trends were also observed for how often they used AI for work and how easy it was to know when to change LoA for alternative 3 in the questionnaire (Figure B.3c) ($p < 0.05$), and how easy it was to understand why to change LoA for alternative 2 (Figure B.3b) ($p < 0.01$) and alternative 3 in the questionnaire (Figure B.3c) ($p < 0.01$) (see Table B.1).

In terms of determining when to change LoA for the alternatives in the questionnaire, four participants rated either the highest or both the lowest and highest levels of transparency equally highly on the Likert scale (Figure B.2a). Conversely, the remaining four participants rated the lowest level of transparency highest in comparison to the other two levels. Overall, the medium level of transparency was rated the lowest. As for why to change LoA, six participants either increased or maintained their rating for each increase in the level of transparency, resulting in the highest level of transparency being rated the highest (Figure B.2b). However, two participants decreased their rating on the Likert scale for each increase in the level of transparency.

5 DISCUSSION AND INSIGHTS

Our experimental results illustrate how informing an operator about a robot's performance degradation through visual cues can affect the operator's driving and LoA regulation style. When shown the health bar, operators triggered LoA switches to mitigate situations where the robot was 'unhealthy'. This style of HI-LoA regulation reduces the aggregate risk of robot failure. The interface with the health bar did not impose a significant additional cognitive workload on the operator during the experiment, and did not significantly change the task completion time. Contrary to our hypothesis, LoA regulation using the Health Bar required significantly ($p < 0.05$) higher perceptual effort. In condition A, Participants carried out fewer rotations and gave less accurate answers. One likely explanation for this is that adding a Health Bar increased the number of points on the UI the operator had to focus on, which increased the perceptual effort for the primary task.

Table B.1 shows that novice robot prefer significantly less transparency than experienced robot operators. Participant feedback revealed that while novice operators responded more to the colours of the health bar, experienced operators preferred having more information displayed to inform their LoA switching decisions. This

indicates that participants value transparency only when they understand the necessity of transparency. People with different levels of experience have varied mental models about the Robot, leading them to perceive the system and the Health Bar differently. This makes it difficult to standardise explanations. Depending on their experience, each participant may want to clarify different aspects of the Human-Robot System, leading to differences in the way they are primed for the experiment. Therefore, rigorous training coupled with detailed explanations about the different components of the system are required to minimise the differences in perception. Finally, the major limitation of this study is the small sample size. Therefore all the results and insights presented in this study can require validation with a larger set of participants.

Participant Recommendations to improve the Robot Health Bar and User Interface Design:

- Make the Health Bar more Salient, so that it attract attention when necessary
- Do not include a drop down-menu in the Health Bar as shown in the Questionnaire
- Use a percentage to the health value instead of a value between $[0, 1]$
- To attract attention to 'low health' use multi-modal awareness cues like sound alerts or UI elements e.g. By Shaking or Blinking the health bar after a threshold value.

6 CONCLUSION AND FUTURE WORK

In this study, we explored if visual cues about a robot's performance degradation can reduce the perceptual effort required to make LoA switching decisions for remote mobile robot navigation tasks. Inspired by video games, we designed a 'Robot Health Bar'. This health bar displayed the total runtime performance degradation the robot is facing. A total of 8 participants carried out a mobile robot navigation task with and without the health bar UI element under high cognitive workload, and their performance was measured. Adding a Robot Health Bar to the robot control UI significantly changed how the operator makes Level of Autonomy Switching Decisions. When the Health Bar was displayed, operators were more attentive and took control of the robot more to minimise the risk of the robots failing. Visual cues that indicate a robot's health can serve as an effective way of ensuring safer control of robots, especially in extreme environments where robot missions have high levels of risk and environmental adversities. In the future, we aim to explore how the insights gained from this study can scale to multi-robot systems.

REFERENCES

- [1] Aniketh Ramesh, Rustam Stolkin, and Manolis Chiou. Robot vitals and robot health: Towards systematically quantifying runtime performance degradation in robots under adverse conditions. *IEEE Robotics and Automation Letters*, 7(4): 10729–10736, 2022.
- [2] Thomas B Sheridan and William L Verplank. Human and computer control of undersea teleoperators. Technical report, Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab, 1978.
- [3] Manolis Chiou, Rustam Stolkin, Goda Bieksaite, Nick Hawes, Kimron L Shapiro, and Timothy S Harrison. Experimental analysis of a variable autonomy framework for controlling a remotely operating mobile robot. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3581–3588. IEEE, 2016.

- [4] Matthew S. Prewett, Ryan C. Johnson, Kristin N. Saboe, Linda R. Elliott, and Michael D. Covert. Managing workload in human–robot interaction: A review of empirical studies. *Computers in Human Behavior*, 26(5):840–856, 2010. doi: 10.1016/j.chb.2010.03.010.
- [5] Kei Mizuno, Masaaki Tanaka, Kouzi Yamaguti, Osami Kajimoto, Hirohiko Kuratsune, and Yasuyoshi Watanabe. Mental fatigue caused by prolonged cognitive load associated with sympathetic hyperactivity. *Behavioral and Brain Functions*, 7(1):17, 2011. doi: 10.1186/1744-9081-7-17.
- [6] Siddharth Agrawal and Holly Yanco. Feedback methods in hri: Studying their effect on real-time trust and operator workload. *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018. doi: 10.1145/3173386.3177031.
- [7] Dakota C. Evans and Mary Fendley. A multi-measure approach for connecting cognitive workload and automation. *International Journal of Human-Computer Studies*, 97:182–189, Jan 2017. doi: 10.1016/j.ijhcs.2016.05.008.
- [8] Aya Hussein and Hussein Abbass. Mixed initiative systems for human-swarm interaction: Opportunities and challenges. In *2018 2nd Annual Systems Modelling Conference (SMC)*, pages 1–8, 2018. doi: 10.1109/SYSMC.2018.8509744.
- [9] C Chace Ashcraft, Michael A Goodrich, and Jacob W Crandall. Moderating operator influence in human-swarm systems. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pages 4275–4282. IEEE, 2019.
- [10] Jonathan Vitale, Meg Tonkin, Sarita Herse, Suman Ojha, Jesse Clark, Mary-Anne Williams, Xun Wang, and William Judge. Be more transparent and users will like you: A robot privacy and user experience design experiment. In *Proceedings of the 2018 ACM/IEEE international conference on human-robot interaction*, pages 379–387, 2018.
- [11] Jessie YC Chen and Michael J Barnes. Human–agent teaming for multirobot control: A review of human factors issues. *IEEE Transactions on Human-Machine Systems*, 44(1):13–29, 2014.
- [12] Tracy L. Sanders, Tarita Wixon, K. Elizabeth Schafer, Jessie Y. C. Chen, and P. A. Hancock. The influence of modality and transparency on trust in human-robot interaction. In *2014 IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, pages 156–159, 2014. doi: 10.1109/CogSIMA.2014.6816556.
- [13] Stein C. Llanos and Kristine Jørgensen. Do players prefer integrated user interfaces? a qualitative study of game ui design issues. In *DiGRA & #3911 - Proceedings of the 2011 DiGRA International Conference: Think Design Play*. DiGRA/Utrecht School of the Arts, 2011. URL <http://www.digra.org/wp-content/uploads/digital-library/11313.34398.pdf>.
- [14] Aaron Steinfeld. Interface lessons for fully and semi-autonomous mobile robots. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 3, pages 2752–2757. IEEE, 2004.
- [15] Robin R Murphy and Satoshi Tadokoro. User interfaces for human-robot interaction in field robotics. *Disaster Robotics: Results from the ImPACT Tough Robotics Challenge*, pages 507–528, 2019.
- [16] Manolis Chiou, Nick Hawes, and Rustam Stolkin. Mixed-initiative variable autonomy for remotely operated mobile robots. *ACM Transactions on Human-Robot Interaction (THRI)*, 10(4):1–34, 2021.
- [17] Giorgio Ganis and Rogier Kievit. A new set of three-dimensional shapes for investigating mental rotation processes: validation data and stimulus set. 2 2016. doi: 10.6084/m9.figshare.1045385.v13. URL https://figshare.com/articles/dataset/A_new_set_of_three_dimensional_stimuli_for_investigating_mental_rotation_processes/1045385.

A OPEN-ENDED QUESTIONS AND QUESTIONNAIRES

A.1 Background information questionnaire

The background questionnaire consisted of the following questions:

- Q1: How often do you operate or use to operate remote controlled vehicles (e.g. Robots, Drones, Heavy machinery)
- Q2: How often do you play or used to play video games involving driving, flight simulation and third person shooters, RPG and sports?
- Q3: Do you use AI (e.g. Autonomous Robots, Machine Learning Algorithms, AI Tools) for work?
- Q4: Do you use AI (e.g. Personal Assistant) in your personal life?

The possible answers ranged between 1-5; i.e., least to most often.

A.2 Open-ended questions

Questions asked before participants had performed the experiment under both conditions are listed below:

- Q1: What are you thinking as you look at this?
- Q2: What is your first impression of this UI element?
- Q3: What do you think this UI element does or will do?

All questions asked after participants had performed the experiment under both conditions are listed below:

- Q1: Was anything surprising or did not perform as expected in either of the interfaces?
- Q2: Was the interface without the 'Robot Health Bar ' easy to understand?
- Q3: Was the interface with the health bar easy to understand?
- Q4: What did you think about the colours used in the health bar when the health changed?
- Q5: Was the LoA switching behaviour of the interface without the 'Robot Health Bar ' transparent?
- Q6: Was the LoA switching behaviour of the interface with the 'Robot Health Bar ' transparent?
- Q7: Is there anything else that you can think of, specific to the UI and UX to improve the use of LoA switching Robots?

A.3 Transparency/trust questionnaire

The transparency/trust questionnaire consisted of images of three different alternatives of a health bar ranging from low to high transparency and a 7-point Likert scale for each image (Figure B.3). The questions were the following pair, for each alternative:

- Q1: It is easy to understand WHEN to change Level of Autonomy (LoA) based ONLY on alternative 1?
- Q2: It is easy to understand WHY to change or not change Level of Autonomy (LoA) based ONLY on alternative 1?

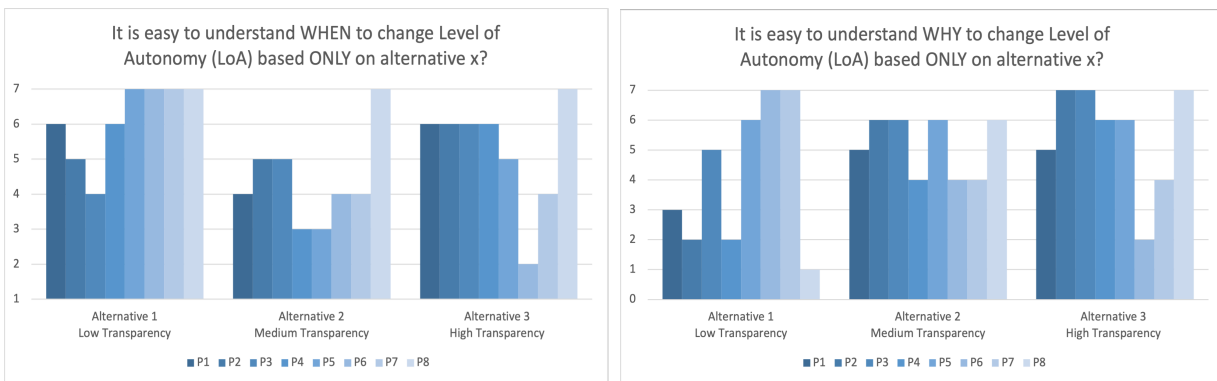
B RESULTS

Table B.1: Correlation Between Background and Online Questionnaire

		How often do you operate or use to operate remote-controlled vehicles (e.g. Robots, Drones, Heavy machinery)	How often do you play or used to play video games involving driving, flight simulation and third person shooters, RPG and sports?	Do you use AI (e.g. Autonomous Robots, Machine Learning Algorithms, AI Tools) for work?	Do you use AI (e.g. Personal Assistant) in your personal life?
It is easy to understand when to change Level of Autonomy (LoA) based ONLY on alternative 1?	Pearson's r	-0.874	0.045	-0.338	-0.336
	p-value	0.005**	0.916	0.413	0.416
It is easy to understand when to change Level of Autonomy (LoA) based ONLY on alternative 3?	Pearson's r	0.455	-0.277	0.802	0.120
	p-value	0.257	0.507	0.017*	0.778
It is easy to understand WHY to change or not change Level of Autonomy (LoA) based ONLY on alternative 2?	Pearson's r	0.363	0.051	0.895	0.191
	p-value	0.376	0.905	0.003**	0.651
It is easy to understand WHY to change or not change Level of Autonomy (LoA) based ONLY on alternative 3?	Pearson's r	0.433	-0.266	0.858	0.000
	p-value	0.284	0.524	0.006**	1.000

*p < 0.05, **p < 0.01

Note. Only a few selected results from the Pearson's correlation are displayed



(a) Results for the questions regarding when to change LoA

(b) Results for the questions regarding why to change LoA

Figure B.2: Results from questionnaire

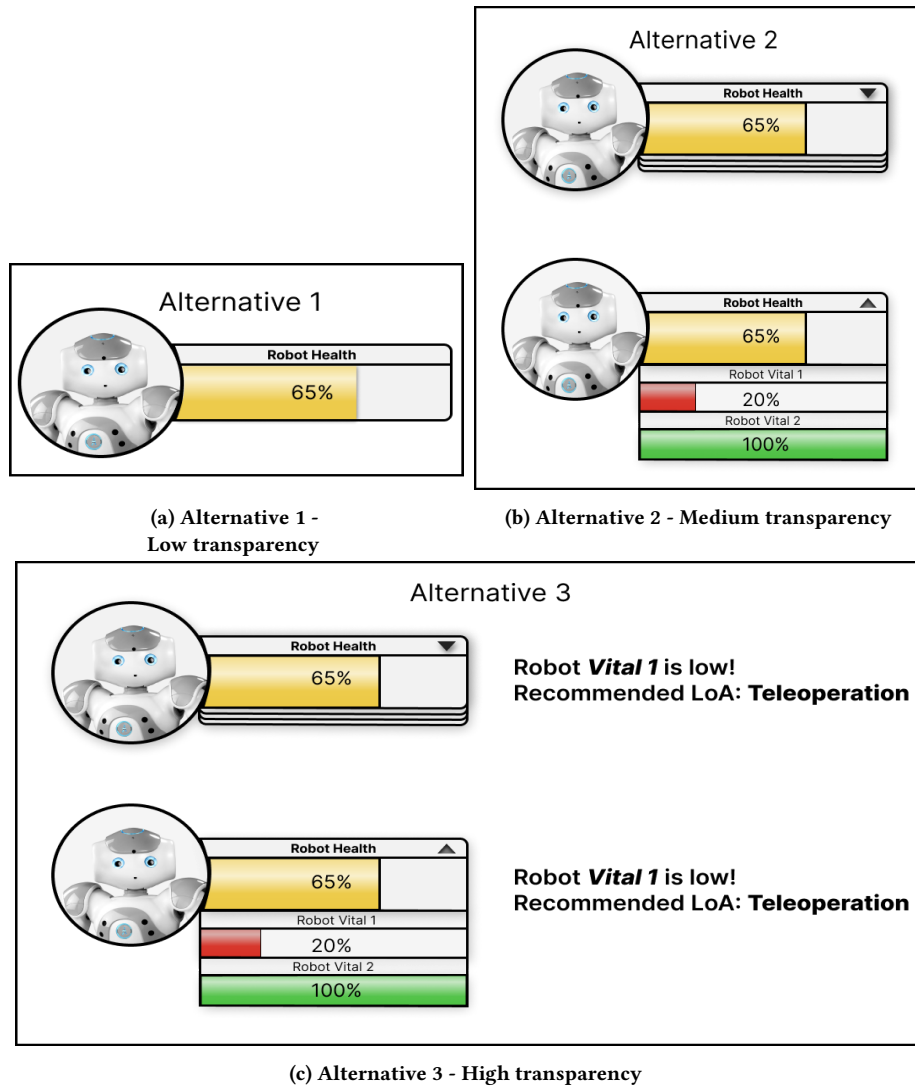


Figure B.3: Alternatives ranging from low to high transparency