Gorillas on a boat: Inattentional blindness during supervisory control of autonomous vessels

Erik Veitch[1*], Ole Andreas Alsos[1]

[1]Department of Design, NTNU, Trondheim, Norway

*erik.a.veitch@ntnu.no

Abstract. When focused on a task, unexpected visual stimuli can go by unnoticed. Inattentional blindness may be problematic for supervisory controllers of autonomous vessels, because this role specifically relies on identifying unexpected stimuli in case intervention is needed. In a simulation-based experiment (n=32), we show that 50% of participants did not perceive an unexpected visual stimulus (a gorilla passenger) when assigned the task of supervisory controller for an autonomous ferry. Additionally, eye-tracking showed that 12 of the 16 players who did not report the gorilla in fact gazed directly at it. Results also showed no correlation to skillset (gamer or navigator) nor to vigilance (5- or 30-minute scenarios); they did, however, show positive correlation to multitasking (1 or 3 supervised vessels). We discuss implications for designing work tasks for supervisory control of remotely operated and autonomous vessels.

1. Introduction

Inattentional blindness is a phenomenon whereupon we do not perceive objects in our visual field while focusing attention on a distractor task [1]. The most famous demonstration of this phenomenon was in an experiment where naïve participants were asked to count basketball passes in a video clip [2]. In the middle of the clip, a person in a gorilla costume walks through the frame, stops in the middle, thumps their chest, and walks off the frame. Amazingly, when asked if they saw anything unusual, only about half recounted seeing a gorilla. A variety of experimental studies have documented inattentional blindness [1,3–6]. Psychologists have long known that visual attention and awareness are not the same thing [7,8]. However, inattentional blindness may be problematic if one's task is specifically to look for objects or events that are out-of-the-ordinary and thereby potentially requiring intervention.

Unfortunately, this is precisely the case for supervisory controllers of autonomous vessels, a role that is emerging in tandem with autonomous vessels. Even the most advanced navigation algorithms rely on humans for supervision, and some even imply a back-up role [9]. Preventative action most often takes the form of manual intervention, often in the case of an unexpected event. Given the propensity for humans to be inattentively blind to unexpected objects while focused on a task, however, this raises an important question: to what extent can we rely on screen-based monitoring for safety-critical supervisory control? Does inattentional blindness play a significant role even for supervisory controllers whose role is defined by looking out for hazards, without specifying what those hazards might be?

In this study, we use an original experiment to investigate these questions. This comes at a time when defining the role of remote supervisory controllers is one of the highest priority issues in addressing regulatory gaps for autonomous ship operations [10]. At least two studies have shown that inattentional blindness can have dramatic results for airplane pilotage, demonstrating that ostensibly obvious objects

on a runway can go entirely unseen during simulated landing [11,12]. To date, however, no studies have examined how the phenomenon may affect remote supervisory control of autonomous vessels.

In our experiment, we also investigated whether certain factors affected inattentional blindness. Specifically, three factor effects were screened: (i) skillset, (ii) vigilance, and (iii) multitasking.

To test the skillset factor, we tested two groups of individuals: video game players and licensed maritime navigators. We hypothesized that gamers would demonstrate higher recall of the visual stimulus (fewer instances of inattentional blindness), because gamers are generally more accustomed to seeing out-of-the-ordinary visual objects during gameplay. This is line with theory explaining that even unseen stimuli are capable of priming and thus affecting subsequent recognition [4].

To test the vigilance factor, we separated participants into two groups: one with 5 minutes transpiring before appearance of the visual stimulus, and the other with 30 minutes before appearance. Previous work shows that a vigilance decrement sets in after approximately 20 minutes when engaged in a passive monitoring task [13–15]. Thus, we hypothesized that more individuals would recall the stimulus when vigilance decrement was present (in the 30-minute scenario) than not (in the 5-minute scenario), because attention on the distracting task was diminished in the former.

Finally, to test the multitasking factor, we divided participants once more into two groups: one supervising a single vessel and one with three vessels. We hypothesized that because visual attention was divided in the multitasking scenario, recall frequency would be diminished.

In Section 2 we outline the method used to investigate inattentional blindness and formally test our hypotheses. In Section 3 we present the results, and in Section 4 we discuss the implications of the results for relevant aspects including training and work design. In Section 5 we present concluding remarks about what inattentional blindness means for supervisory control of autonomous vessels.

2. Method

We used an open-source simulator based on the Gemini platform [16] to test whether volunteer participants perceived an unexpected stimulus during a simulation of remotely controlled, automated ferry operations. Tests were completed from March-June 2022 and are also reported in [17] in the context of factor screening for general supervisory control performance. This section briefly outlines the following methodological aspects: (i) the simulator, (ii) the experimental design, and (iii) participant recruitment. Additional details are found in [17].

2.1. Simulator

The simulator is hosted at the NTNU Shore Control Lab, which is a research platform for investigating human-machine collaboration in autonomous vessel operations [18]. The simulator has two parts: (i) a Scenario Builder [19] and (ii) a Scenario Player [20]. This division allows researchers to build custom scenarios (e.g., vessel, environment, traffic) and invite users to execute the scenarios. Observing the results can yield insights into aspects of human behaviour, emergency response, and human-computer interaction for the novel role of supervisory controller when completing monitoring and intervention tasks. It can also be valuable for prototyping interface designs during early design stages of remote control center infrastructure. For this study, it allowed us to define the two distinct modes of control (automatic and manual) as well as the control transfer between the two modes. As we will see later, it also allowed us to insert a gorilla as a passenger aboard the simulated passenger ferry. Most importantly for our research, it allowed us to investigate aspects that are difficult to study in real-life situations in a repeatable, controlled way.

The milliAmpere2 was featured as the simulated "own-boat" in the simulator. The milliAmpere2 is an autonomous urban passenger ferry designed and built at NTNU, which is based on its successful predecessor, milliAmpere [21] (Figure 1). In the simulator, we recreated the route the ferry traversed during public trials in 2022 (Figure 2). This entailed creating a "digital twin" of the milliAmpere2, which modelled important details like the vessel's inertial, hydrodynamic, and maneuvering properties. The digital twin also modelled physical appearances of the ferry and its environment, including docks and urban infrastructure. Even boat traffic in the simulator was based on real vessels in Trondheim's canals.

During simulation, the milliAmpere2 digital twin operated autonomously much like the real one did, stopping to allow crossing traffic to pass and allowing compliancy with operators' control takeovers using a joystick for manual maneuvering. In the simulator, animated passengers lined up, boarded the ferry, and deboarded, much like they did during the real-life trials. This thematic strategy was chosen to improve the scenario fidelity [22] and to allow the opportunity for full-scale validation in similar trials.

The open-source Gemini platform was built in Unity [23], which uses the PhysX engine for modeling dynamics. In 2021-22, the "Scenario Builder" and "Scenario Player" functionalities were added to the Gemini platform and was made available at GitHub (https://github.com/mikael-rh/ScenarioBuilder) with some licensing restrictions.

Figure 1. milliAmpere2 during field trials (Photo Mikael Særeteid, 30 Sep 2022).

Figure 2. milliAmpere2 digital twin in Gemini simulator (screenshot).

2.2. Test design

Participants executed the simulated scenarios individually (n=32; two scenarios each). The lengths of the scenarios varied (either 5 minutes or 30 minutes each) and the total experiment time took about 1 to 2 hours (depending on scenario length). Each trial followed the same protocol and script. Prior to testing, each participant underwent the same two training scenarios, familiarizing participants with the simulator and their role.

The participants were told that they were supervisory controllers whose job it was to ensure that passengers boarded and deboarded safely and to take over control, if necessary. At one point in the scenario, an event occurred requiring manual intervention. Participants completed two scenarios: in the first, an alarm appears warning of an autonomous control system failure ("handover"); in the second, the participant must avoid a collision with a small boat ("takeover"). These scenarios were always completed in the same order. A full two minutes (120 seconds) before the critical event occurs in each scenario, an animated gorilla boards the ferry, directly in sight of the supervisory controllers. Inspired by the famous "Gorillas in our midst" experiment [2], the gorilla was specifically used to check for inattentional blindness. The gorilla appeared in the same location and at the same time in both scenarios and was equally as visible for the same amount of time. The scenario design essentially forces supervisory controllers to bring the gorilla into their field of view; however, given that they control the cameras, total visibility time varied among participants from just a few seconds to over a minute.

After completing the scenarios, the participants were asked in an exit interview whether they noticed the gorilla, and if so, in which scenario(s). Eye-tracking data (collected with PupilLabs "PupilCore" goggles) were used to confirm which participants gazed directly at the gorilla, irrespective of what was reported in the interviews.

Aside from simply checking for inattentional blindness, the test was designed to explore three independent factor effects. A factorial design was chosen [24] by defining test treatments based on three categorical factors at two levels each, for a total of 2^3 permutations (see [17] for details about the overall

experiment design.) For this study, each permutation was effectively repeated four times, for a total of 2^3 x 4 = 32 runs. This number of runs was judged to balance experimental accuracy (given expected factor effect variation) with the resource demands associated with recruiting volunteers and running trials. Runs were completed in randomized order within four blocks to ensure that no run-order-associated confounding factors went undetected. Note that the full experiment contained two additional variables (namely time pressure and availability of a decision support system), but since these were not of interest to inattentional blindness, their treatments were considered as repetitions here. A detailed description of the full experiment, which was designed primarily to screen factor effects on general remote supervisory performance, is available in [17].

We used analysis of variance (ANOVA) to analyze the results. This effectively allowed us to measure the signal-to-noise ratio determining whether the variance in measured outcomes was associated to changes in factor levels. If an association was detected, we inferred the relationship as causal if there was less that a 10% probability that it was due to chance alone (alpha level = 0.10). Because the results were coded as either 1 or 0 (for either "detection" or "no detection"), we applied an arcsine square root transformation to the results before completing the ANOVA.

2.3. Test participants

Two groups were recruited based on skillsets: (i) video game players (n=16; 3 female, 13 male; age 25 ± 5 sd) and (ii) licensed maritime navigators (n=16; 2 female, 14 male; age 43 years ± 9 sd). Participants aged 18-65 were recruited using sign-up posters and social media (gamers) and through internal networks and local interest groups (navigators). While aging has been shown to increase inattentional blindness [25], the average difference between the groups in this study (18 years) was not sufficient to produce this effect. In terms of skillsets, there was little overlap between groups: gamers had negligible maritime experience and navigators had negligible gaming experience relative to their groups' counterparts. Participants with eyeglasses wore contact lenses due to the eye tracking headwear. Because personal and health-related data were recorded in the experiment, written consent was obtained from all participants and a data management plan was established with the Norwegian Center for Research Data (NSD). Additional demographic details are available in [17]. All participants were rewarded with a gift card.

3. Results

Half the participants (n=16) reported that they did not see the gorilla in the first scenario. Of these, 75% (n=12) gazed directly at the gorilla according to eye-tracking data (Figure 3).

Interestingly, almost all reported seeing the gorilla in the second scenario. Using eye-tracking data, we found that out of the 12 participants that gazed directly at the gorilla without perceiving it, 10 of them reported seeing it in the second trial. Even among those who had the gorilla only in their peripheral vision, 75% (3 out of 4) reported seeing it in the second trial. This indicated that simply having the gorilla in one's field of view (whether in direct gaze or peripherally) was a strong predicter of whether the gorilla was perceived in the second trial.

Three participants reported seeing no stimulus either in the first or second trials. Examining these individual trials closer, we found that all of them had limited visual exposure to the stimulus. In the first scenario, one participant did not have the gorilla in the scene all at, one only peripherally, and one only peripherally with a brief saccade. For these trials, visual exposure was limited in the second scenario, too. In this scenario (and in respective order), one player had only a brief saccade, one never had it in the scene, and one only peripherally with a brief saccade. In other words, good visual contact with the unexpected stimulus (consisting of more than a cursory saccade and/or peripheral vision) was highly correlated with perceiving the same stimulus in a second trial, provided that good visual exposure was provided again in the second trial.

Figure 3. Eye-tracking viewer showing gaze (red/green dot). This participant did not "see" the gorilla despite gazing directly at it (dotted circle added).

3.1. Factor screening

Out of the three factors we tested (skillset, vigilance, and multitasking), only multitasking had a significant effect on inattentional blindness (Table 1). Skillset and vigilance had no effect, contrary to our hypotheses. In other words, the only factor affecting recall was whether participants' attention was divided across multiple vessels. About a third of participants (11 out of 32) recalled the gorilla when there was a single ferry, while only about a sixth did (5 out of 32) when there were three ferries. This confirms our hypothesis that divided attention lowers the frequency of recall; if there is more to which to attend, there is also more to miss. We will examine this result more closely in the Discussion.

Table 1. Results of factor screening on gorilla recall (* indicates significant effect).

	Recall frequencies (/16)		ANOVA		
Factor	*Low level*	*High level*	Error df	F-value	p-value
A-Skillset	*Gamers:* 8	*Navigators:* 8	28	0.00	1.00
B-Vigilance	*No decrement:* 7	*Decrement:* 9	28	0.52	0.48
C-Multitasking	*1 ferry:* 11	*3 ferries:* 5	28	4.67	0.04*

3.2. Transparency

The interview data in this experiment is openly available in a data archive [26].

4. Discussion

The results replicated the findings of the classic "Gorillas in our midst" study [2] on inattentional blindness. What is special about our study, however, is that it did so in the context of supervisory control of autonomous vessels. This highlights an important implication: even when one's assigned task is to visually monitor a process for deviations, one is not able to fully perceive the rich context of visual cues being monitored. Looking is not the same as seeing for supervisory control. Whether a visual stimulus is perceived while attention is focused on a distractor task appears as predictable as the toss of a coin.

Beyond this, the results were striking considering the rejection of two hypotheses. Firstly, we presumed that skillset would reduce instances of inattentional blindness because video game players were believed to be more accustomed to seeing out-of-the-ordinary visual objects during gameplay and were thus primed to perceive them. However, half the gamers and half the navigators recalled the gorilla,

showing no skillset effect whatsoever. Upon checking the literature, this result was found to have a historical precedent. Specifically, one study showed that even experienced radiologists were susceptible to inattentional blindness when scanning lung X-rays that had large gorilla-shaped tumors added to them [27]. The implication is that inattentional blindness may affect all supervisory controllers, regardless of experience or skillsets.

The second null result was that vigilance decrement had no effect on inattentional blindness. We had presumed that when vigilance decrement set in after about 20 minutes in passive monitoring, instances of inattentional blindness would decrease as attention originally focused on the distractor task was relaxed. However, the results indicated only a minor difference in gorilla recall frequency (only two more for the group in the attention-relaxing 30-minute scenario). This suggests that the distracting effect of the task at hand is present despite the drop in vigilance that occurs after extended periods in passive monitoring.

Multitasking was found to be the only factor that affected inattentional blindness/recall of the stimulus in our study. At first glance, this confirmed the logical premise of our hypothesis: if there is more to attend, there is also more to miss. However, upon closer inspection, this result also appeared to go against the grain of our original expectations. This is because the scenarios were designed in such a way as to force participants to attend to the vessel on which the gorilla was in plain sight, regardless of how many ferries were under supervision. Thus, one may justifiably presume that any effect of multitasking would be nullified. Examination of the multitasking scenario results shed some light on this conundrum. Specifically, those who had three ferries were significantly slower to react to the critical situation and thus had the stimulus in their field of view for a shorter period. (A detailed analysis of the reaction time results is found in [17].) Therefore, we may accept our hypothesis with one important caveat: multitasking does appear to affect inattentional blindness, but only because multitasking affects reaction time and thereby the duration of visual exposure to the stimulus.

Perhaps the most striking result of the experiment was that while only half saw the gorilla in the first scenario, almost everyone saw it in the second scenario. From this, we can conclude that visual exposure to the stimulus was causally related to perceiving it upon re-exposure, even when that stimulus was not consciously perceived upon the first exposure. Findings in the literature support this result in the context of priming. In [4], Mack writes:

> "…unseen stimuli are capable of priming, that is, of affecting some subsequent act. (For example, if a subject is shown some object too quickly to identify it and is then shown it again so that it is clearly visible, the subject is likely to identify it more quickly than if it had not been previously flashed… Priming can occur if there is some memory of the stimulus, even if that memory is inaccessible.) (pg. 181)"

At least one experiment with a similar set-up demonstrated similar results to ours. In it, participants were given an attention-demanding task while exposed to visual pattern stimuli and were subsequently asked to repeat the task [28]. More than half the participants were inattentionally blind to the pattern upon first attempt; however, upon repeating the same task and exposure, all subjects reported seeing the pattern. The unexpected nature of the visual stimulus also appears to play a distinct role. Others have shown, for example, that inattentional blindness does not occur for iconic images like smiley faces or for familiar text like the name of a participant [1].

Brain scanning methods have also shed light on inattentional blindness. For example, one experiment used optical illusions and fMRI brain scans to show that stimuli that are not cognitively accessed are in fact processed in the brain as perceptual interpretations [29].

The paradigm of priming suggests that learning may be enforced without necessitating attention. Indeed, psychologists have long known that so-called subliminal perception can influence affect, learning, and cognition [30]. For example, in one experiment, subjects' preferences shifted for images of facial expressions presented for only 4 ms — a period so short subjects did not even recall consciously seeing them [31]. In another experiment, participants' task performance improved after repeated exposures to a visual stimulus they did not perceive [32]. The notion of subliminal learning begs the question: can we train remote controllers' perception through simple visual exposure? This may have

implications for training using simulators, in which relevant visual stimuli can be easily recreated in a safe, immersive environment.

Despite the provocative power of visual priming, the finding that only half of our supervisory controllers detected the gorilla upon first exposure implies that we may need to rethink how we design supervisory work tasks. Given what we know now about inattentional blindness, it is unreasonable to blame supervisory controllers in the event its occurrence leads to an accident—even if the hazard in question is judged a-posteriori to be "in plain sight." This begs the question: how can we support perception of a scene during supervisory control work? This design challenge is accentuated by the fact that widely accepted human-AI interaction principles do not even mention inattentional blindness [33,34]. At the very least, our experiment suggests that passive back-up or "in-the-loop" situational awareness models may not be well-aligned with how human visual attention really works. Fortunately, though, the situation is not lost. On the contrary, we are currently at a stage where we can have large impact on the design of the supervisory control role. Better alignment of this role to human abilities may involve defining a more active supervisory role: instead of just passively *viewing* the scene, supervisor might actively *interact* with a scene. Enhancing interaction can involve simple tasks like periodic manual maneuvering, logging events, or reporting system status and weather conditions at regular intervals. Another mitigating measure may involve having a partner during monitoring tasks, whose presence alone would theoretically halve the rate of inattentional blindness. In any case, designing interactions will rely on an accurate definition of the shared control dimensions in which a human and AI partner collaboratively control a system [35]. While many such collaborative systems exist (e.g., digital games, telerobotics, surgery, lane-keeping assisted driving), remotely operated and autonomous vessels are distinct in that they place humans in a safety-critical, remotely supervisory role. In this role, the AI is presumed to be deficient, and the supervisor takes over to save the day—not the other way round as in many other collaborative systems. A more precise definition of the shared control dimensions of remote supervisory control work will likely lead to designs of work tasks more appropriately aligned the abilities and limitations of human visual attention.

The work presented also suggests several directions for future investigations. One is to test the efficacy of visual priming during safety-critical operations relying on visual perception. This may have implications for training, for which programs are currently under development. Another research direction is to verify the results of inattentional blindness in full-scale, thus confirming that the phenomenon is of equal relevance to real-world conditions as it is to simulated conditions.

Conclusion

Gorillas are not likely to board autonomous passenger ferries anytime soon. However, if there is one thing we can expect during real-world operations, it is the unexpected—and it will be supervisory controllers' job to deal with it. Our experimental results adopted a gorilla as a colorful metaphor for unexpected visual stimuli to be detected during remote supervisory control tasks. It also provided a way to replicate studies that used a similar gorilla stimulus in their explorations of inattentional blindness [2,27].

Our results suggest that remote supervisory controllers are not immune to inattentional blindness just because their assigned task is to look out for potentially hazardous objects or events. Furthermore, our results showed that skillset (whether gamer or navigator) and vigilance (whether monitoring for 5 or 30 minutes) had no effect on inattentional blindness. Multitasking (whether supervising 1 or 3 ferries) appeared to affect inattentional blindness, but only to the extent that multitasking was also linked to reaction time and thus to visual exposure duration.

Our results also suggest that simply exposing supervisory controllers to a possible hazard in a simulator environment will greatly increase chances of recognition upon re-exposure. Training programs for supervisory controllers may be able to leverage the powerful mechanism of visual priming, which was by far the most effective counter to inattentional blindness.

This work was motivated by regulatory gaps concerning the role of remote supervisory controllers. To this end, our experiment imagined a remote supervisory role based on a simulated autonomous ferry

and collected empirical observations that shed light on this novel role. In this context, our empirical results about inattentional blindness demonstrated the vulnerability of the remote controller's role as it is currently imagined. Designers addressing inattentional blindness may by extension address better overall alignment of work tasks with human capabilities. One approach may be to adopt a design paradigm of shared control, including designing more interactive work tasks and enabling teamwork during operations. Another approach may be to adopt simulation training that leverages visual priming.

References

[1]   Mack A and Rock I 1998 *Inattentional blindness* (Cambridge, MA, US: The MIT Press)

[2]   Simons D J and Chabris C F 1999 Gorillas in our midst: Sustained inattentional blindness for dynamic events *Perception* **28** 1059–74

[3]   Koivisto M and Revonsuo A 2008 The role of unattended distractors in sustained inattentional blindness *Psychol. Res.* **72** 39–48

[4]   Mack A 2003 Inattentional blindness: Looking without seeing *Curr. Dir. Psychol. Sci.* **12** 180–4

[5]   Most S B, Scholl B J, Clifford E R and Simons D J 2005 What you see is what you set: Sustained inattentional blindness and the capture of awareness *Psychol. Rev.* **112** 217–42

[6]   Richards A, Hannon E M and Derakshan N 2010 Predicting and manipulating the incidence of inattentional blindness *Psychol. Res.* **74** 513–23

[7]   Lamme V A F 2003 Why visual attention and awareness are different *Trends Cogn. Sci.* **7** 12–8

[8]   Koch C and Tsuchiya N 2007 Attention and consciousness: two distinct brain processes *Trends Cogn. Sci.* **11** 16–22

[9]   Veitch E, Christensen K A, Log M, Valestrand E T, Lundheim S H, Nesse M, Alsos O A and Steinert M 2022 From captain to button-presser: Operators' perspectives on navigating highly automated ferries *J. Phys.: Conf. Ser.* **2311** 012028

[10] IMO 2021 *Outcome of the regulatory scoping exercise for the use of Maritime Autonomous Surface Ships (MASS)* (London, UK: International Maritime Organization)

[11] Kennedy K D, Stephens C L, Williams R A and Schutte P C 2014 Automation and inattentional blindness in a simulated flight task *Proc. of the Human Factors and Ergonomics Society Annual Meeting* vol 58 (Thousand Oaks, CA, USA: SAGE Publications) pp 2058–62

[12] Haines R F 1991 A breakdown in simultaneous information processing *Presbyopia Research: From Molecular Biology to Visual Adaptation* ed G Obrecht and L W Stark (Boston, MA: Springer US) pp 171–5

[13] Cummings M L, Gao F and Thornburg K M 2016 Boredom in the workplace: A new look at an old problem *Hum. Factors* **58** 279–300

[14] Grier R A, Warm J S, Dember W N, Matthews G, Galinsky T L, Szalma J L and Parasuraman R 2003 The vigilance decrement reflects limitations in effortful attention, not mindlessness *Hum. Factors* **45** 349–59

[15] Molloy R and Parasuraman R 1996 Monitoring an automated system for a single failure: Vigilance and task complexity effects *Hum. Factors* **38** 311–22

[16] Vasstein K, Brekke E F, Mester R and Eide E 2020 Autoferry Gemini: a real-time simulation platform for electromagnetic radiation sensors on autonomous ships *IOP Conf. Ser.: Mater. Sci. Eng* **929** 012032

[17] Veitch E, Alsos O A, Cheng T, Senderud K and Utne I B 2023 Human factor influences on supervisory control of remotely operated and autonomous vessels (Preprint)

[18] Alsos O A, Veitch E, Pantelatos L, Vasstein K, Eide E, Petermann F-M and Breivik M 2022 NTNU Shore Control Lab: Designing shore control centres in the age of autonomous ships *J. Phys.: Conf. Ser.* **2311** 012030

[19] Ek S 2022 *Design of simulator for researching autonomous marine vessels* Master thesis (Trondheim, Norway: NTNU)

[20] Hansen M 2022 *Developing a video game for research and prototyping of unmanned maritime vessels* Master thesis (Trondheim, Norway: NTNU)

[21] Brekke E F, Eide E, Eriksen B-O H, Wilthil E F, Breivik M, Skjellaug E, Helgesen Ø K, Lekkas A M, Martinsen A B, Thyri E H, Torben T, Veitch E, Alsos O A and Johansen T A 2022 milliAmpere: An autonomous ferry prototype *J. Phys.: Conf. Ser.* **2311** 012029

[22] Liu D, Macchiarella N D and Vincenzi D A 2008 Simulation fidelity *Human Factors in Simulation and Training* ed P A Hancock, D A Vincenzi, J A Wise and M Mouloua (CRC Press: Boca Raton, FL, USA) pp 61–73

[23] Unity Technologies 2022 Unity

[24] Montgomery D C 2017 *Design and analysis of experiments* (John Wiley & Sons, Inc.: Hoboken, NJ, USA)

[25] Graham E R and Burke D M 2011 Aging increases inattentional blindness to the gorilla in our midst. *Psychol. Aging* **26** 162–6

[26] Veitch E 2022 Dataset for: Human factor influences on supervisory control of remotely operated and autonomous vessels (DataverseNO)

[27] Drew T, Võ M L-H and Wolfe J M 2013 The invisible gorilla strikes again: Sustained inattentional blindness in expert observers *Psychol. Sci.* **24** 1848–53

[28] Pitts M A, Martínez A and Hillyard S A 2012 Visual processing of contour patterns under conditions of inattentional blindness *J. Cogn. Neurosci.* **24** 287–303

[29] Vandenbroucke A R E, Fahrenfort J J, Sligte I G and Lamme V A F 2014 Seeing without knowing: Neural signatures of perceptual inference in the absence of report *J. Cogn. Neurosci.* **26** 955–69

[30] Zajonc R B 1980 Feeling and thinking: Preferences need no inferences *Am Psychol* **35** 151–75

[31] Murphy S T and Zajonc R B 1993 Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures *J Pers Soc Psychol* **64** 723–39

[32] Watanabe T, Náñez J E and Sasaki Y 2001 Perceptual learning without perception *Nature* **413** 844–8

[33] Shneiderman B 2020 Human-centered artificial intelligence: Reliable, safe & trustworthy *Int. J. Hum.-Comput. Interact.* **36** 495–504

[34] Amershi S, Weld D, Vorvoreanu M, Fourney A, Nushi B, Collisson P, Suh J, Iqbal S, Bennett P N, Inkpen K, Teevan J, Kikin-Gil R and Horvitz E 2019 Guidelines for human-AI interaction *Proc. of the 2019 CHI Conf. on Human Factors in Computing Systems (Glasgow)* CHI '19 (New York, NY, USA: Association for Computing Machinery) pp 1–13

[35] Cimolino G and Graham T C N 2022 Two heads are better than one: A dimension space for unifying human and artificial intelligence in shared control *Proc. of the 2022 CHI Conf. on Human Factors in Computing Systems (New Orleans)* CHI '22 (New York, NY, USA: Association for Computing Machinery) pp 1–21