Effects of Haptic Shared Control on the Operation of Marine Boom Cranes

MSc Thesis - Robotics

E.E. Jerez



TU Delft - Heerema Marine Contractors

TU Delft

Effects of Haptic Shared Control on the Operation of Marine Boom Cranes

by

E.E. Jerez

TU Delft Supervisor:	Dr. Luka Peternel - Associate Prof. Cognitive Robotics
Heerema Marine Contractors Supervisor:	Geert Meskers - Specialist Engineer Simulation
Department:	Human Robot Interaction, Cognitive Robotics [CoR]
Faculty:	Faculty of Mechanical Engineering, TU Delft
Internship Department:	Simulation, Heerema Marine Contractors Nederland SE
Publication Date:	20-10-2026

Cover: SSCV Sleipnir, Heerema Marine Contractors Nederland SE

TU Delft

MARINE CONTRACTORS

Effects of Haptic Shared Control on the Operation of Marine Boom Cranes

Eduardo Enrique Jerez
Cognitive Robotics / Simulation
TU Delft / Heerema Marine Contractors

Supervised by:
Dr. Ir. Luka Peternel, *Cognitive Robotics, TU Delft.*
Geert Meskers *Simulation, Heerema Marine Contractors.*

***Abstract*—This thesis investigated whether the addition of haptic guidance produced a measurable improvement in a shared control task using a simulated offshore boom crane. A counterbalanced within-subject study was conducted with 26 participants, comparing the same autonomous controller (the "autopilot") with and without haptic guidance during a payload moving task on the deck of a crane vessel (analogous to a pick-and-place task). With haptic guidance, the autopilot physically actuated the joysticks, while without haptic guidance, the autopilot actuated the crane directly while the joysticks remained centered. It was hypothesized that haptic guidance would improve task performance and safety, subjective operator experience and perception of the autopilot, and reduce perceived workload. No significant differences were observed in completion time, three-dimensional trajectory length, control-input jitter, collisions, or payload sway, corresponding to task performance and safety. Haptic guidance produced higher scores in subjective operator experience and autopilot perception, as well as a lower perceived workload. Participant comments most often mentioned higher perceived transparency on the actions of the autopilot when haptic guidance was present. The study's findings suggest that haptic guidance primarily improves the perceived quality and transparency of the autonomous controller, and reduces perceived workload, rather than improving the measured performance or safety in a shared-control task involving marine boom cranes.**

***Index Terms*—Haptic feedback, Haptic shared control, shared control, boom crane, vessel-mounted crane.**

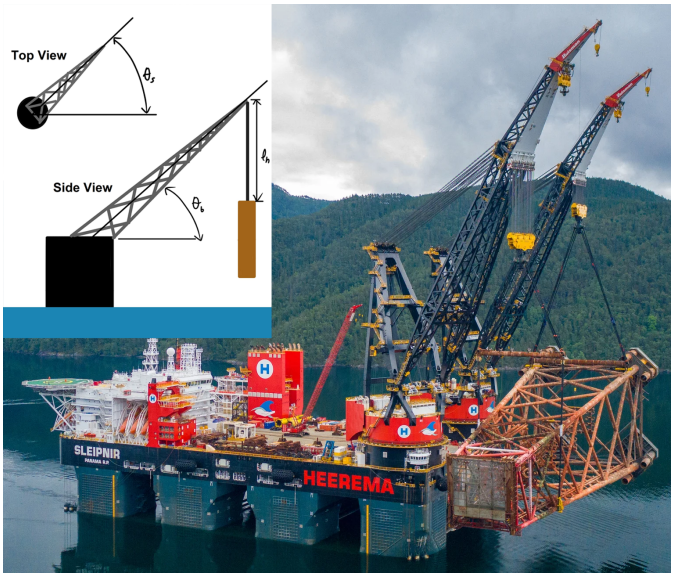


Fig. 1: SSCV Sleipnir lifting a jacket. Top left corner shows the actuation of a boom crane, where θ_s is the slew angle, θ_b is the boom angle and ℓ_h is the hoist length.

I. Introduction

Boom cranes are cranes which are built with a pivoting structural element called the "boom", and are the types of cranes usually employed to lift the heaviest loads. A boom crane has two joints and a hoist cable length which can be actuated as shown in Figure 1:

- Booming consists of pivoting the boom arm up and down.
- Slewing consists of rotating the entire base of the crane about the vertical axis.
- Hoisting consists of shortening or lengthening the hoist cable.

Vessel-mounted marine boom cranes like the Sleipnir, seen in Figure 1, are amongst the largest cranes in the world, with lifting capacities up to 10,000 metric tons [1]. These cranes are extremely important, as there is currently no other device which is capable of installing the large structures necessary for the construction of modern offshore infrastructure, including oil rigs, wind turbines, electrical substations, amongst others.

Controlling these cranes poses a large number of challenges related to their enormous sizes and lifting capacities, coupled with the fact that they are mounted on floating vessels. Higher order effects resulting from the finite stiffness of the crane structure and vessel deck, vessel dynamics, operational safety, and operational performance (that is, how quickly and precisely a task can be completed) all pose their own significant challenges.

Furthermore, most of the heavy lift offshore boom cranes currently in operation make very limited use of automation, in many cases only the ability to hold the payload at a fixed height. Efforts are ongoing to develop and implement diverse control systems for different types of cranes [2], although these efforts have largely been focused on other types of cranes like rotary cranes [3] [4] [5].

In the context of vessel-mounted boom cranes, the State of the Art includes the development of a custom novel haptic device by Chu et al. [6], in which a device based on the dynamics of the crane is used to allow the operator to feel the forces that the end effector of the crane experiences; which is not a viable solution for heavy lift cranes which are traditionally operated via joysticks and which often move much more slowly than the small crane used in the paper. Other work, like that by Sun et al. [7], Huang et al. [8] and Maleki et al. [9] studies the dynamics of boom cranes, but does not deal with vessel mounted cranes which are affected by vessel dynamics. Work also exists in the context of control and sway reduction in boom cranes, like the paper by Arnold et al. [10] in which the authors implement a model predictive control (MPC) approach to mitigate sway in the operation of on-shore boom cranes, but which does not look into the shared control aspect of operation.

The state of the art, thus, shows a significant gap in terms of the shared control of boom cranes - that is, the protocols by which the operator and autonomous controller share the control tasks - and especially in regards to those that are very large and vessel-mounted which must, due to the need for simple and space-efficient control interfaces, use joysticks. Within the shared control of these cranes, little attention aside from the study by Chu et al. [6] has been paid to haptic guidance, or haptic shared control: this approach involves having the controller exert physical force on the control interface in order to produce the desired control inputs - thereby physically sharing the control authority with the operator - and is well suited to cranes, as it is ideal to implement on the joystick-based control interfaces that many cranes use.

Haptic shared control strategies have gained much traction and become an increasingly researched topic, due to their potential to allow for the operator and autonomous controller to share control in an intuitive and seamless manner. This is particularly true in the automotive field, as discussed by Abbink et al. [11], which examines different contemporary haptic shared control approaches, focusing on whether the implementations they utilize use a fixed or a variable degree of haptic authority. However, much research is ongoing as well outside of the automotive field, as exemplified by the paper by Yu et al. [12], which investigates different control sharing protocols between an operator using a haptic device and an autonomous controller attempting to navigate a series of obstacles of which one was movable.

Some contemporary papers display the degrees of complexity current haptic shared control systems are capable of addressing. One such paper is that by Mulder et al. [13], in which the authors set up a driving simulator and tested participant's lateral error in a lane-following task, using haptic shared control on the steering wheel with several different levels of haptic authority. The authors showed that the implementation of shared control reduced the control activity required to increase safety performance in terms of lateral error, and while full automation further reduced the required human control activity to zero and further increased safety (again, in terms of lateral error), it put the driver in a supervisory position, which introduces risks associated with the possibility of disengagement or inattention, amongst other issues. Another such paper is that by Benloucif et al. [14], where the authors developed a cooperative control approach meant to assist the driver via haptic feedback on the steering wheel when changing lanes and avoiding obstacles during highway driving. The authors of this paper included several aspects of driving into their solution, including driver intent deduction in order for the autonomous controller to make driving decisions, trajectory planning, a fuzzy steering controller, as well as compensation for vehicle dynamics and human perception. In this paper, the authors set up a driving simulator experiment consisting of driving on a highway with instructions to stay in the right lane, but with a series of obstacles (stationary cars) placed at different points, partially obstructing the lane. Participants performed the experiment with no assistance, with traditional lane assist, and with the full shared control implementation; and the results showed both better objective performance and a better subjective driver experience with the shared control implementation. Unfortunately, neither of these studies (nor any of the automotive ones) is directly applicable to the field of boom crane operation, not least because the control interfaces are completely different (steering wheels vs. joysticks), but also because the control tasks themselves are very different (driving vs. the pick-and-place-like nature of crane operation). Studies like these nonetheless show the enormous potential that haptic shared control has to assist and enhance human operators in complex control tasks.

In this context, research on Haptic Shared Control within the context of human-robot collaboration provides a very useful guideline. Mol et al. [15] set up a robot-assisted force-position task in which 4 conditions were tested, wherein the robot assisted with either the force or position, both, or neither. Rodríguez-Sedano et al. [16] set up a teleoperated multi-stage pick-and-place experiment in which the participants had to perform the experimental task with and without haptic feedback. Both of these studies measured the effects of different conditions - within haptic shared control - on the subjective operator experience using NASA-TLX as well as their own custom subjective operator experience questionnaires, and found the implementation of different haptic shared control conditions to affect the subjective operator experience and perceived workload, most often positively. However, these types of collaborative and assisted tasks are still very different, both operationally and in terms of scale, from the operation of a heavy lift boom crane.

Thus, to address this gap in the research regarding the haptic shared control of (heavy lift) marine boom cranes, this study poses the following question: "How does haptic guidance affect the operation of marine boom cranes?"

To answer this question, a human factors experiment was set up wherein test subjects control one of the cranes of a heavy lift crane vessel, using Heerema's crane simulator. A custom haptic device was developed in-house, based on the 3 single-axis joystick setup of the SSCV Thialf (a vessel with

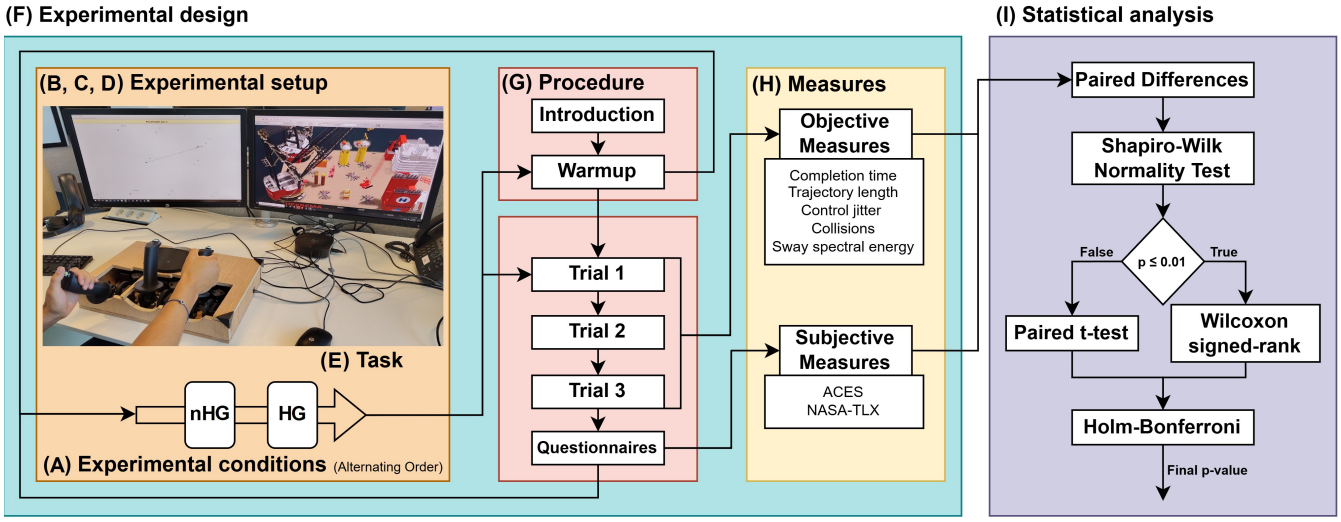


Fig. 2: Methodological pipeline based on the structure used by Mol et al. [15]. The experimental setup (seen on the left) consists of the haptic device; which contains 3 single-axis joysticks equipped with haptic feedback, a screen displaying the crane simulator in a controllable free floating view, a second screen displaying a simple top-view map which serves as the autopilot control interface, as well as a keyboard and a mouse. The author sat to the right with the experiment control panel on a separate screen.

similar configuration to the SSCV Sleipnir), which allows the implementation of force feedback on the controls of the crane.

The way that operators generally control these heavy lift vessel-mounted boom cranes is by simply setting each joystick axis to a fixed position, and waiting until the desired final position has been reached. This is due to the very slow speeds at which the crane is operated - moving from one point to another can take from several minutes to hours. Given the geometry of a boom crane, this makes achieving radial (or arc-based) trajectories elementary, but makes straight line trajectories fundamentally difficult to achieve, as they require continuous and progressive correction of the control inputs.

Thus, to address the lack of crane automation, an autonomous control system was developed, hereafter referred to as "autopilot", which allows the operator to move the payload along a straight line trajectory from an initial position to a target position. This autopilot system was then equipped with the capacity to perform haptic guidance, that is, to actuate the joysticks physically to achieve the desired control inputs - this is hereafter referred to as "haptic guidance".

Each participant was therefore exposed to two conditions: crane autopilot with haptic guidance and crane autopilot with no haptic guidance. Measures come from recorded simulation parameters and subjective opinion questionnaires. Through statistical analysis of the measured and subjective data, the author hopes to gain insight into the effects of haptic guidance on the operational performance and safety, and the operator's experience and workload. The author's hypotheses are:

- H1 The proposed haptic guidance improves the task performance and safety.
- H2 The proposed haptic guidance improves the operator's subjective experience.
- H3 The proposed haptic guidance improves (reduces) the operator's perceived workload.

II. Methodology

The experiment setup consists of a custom-made haptic device, as well as an autopilot system for the crane complete with a simple control panel, and the crane simulator itself, provided by Heerema. The code containing all the functionality (except any HMC proprietary code related to the vessels or simulator) is accessible on GitHub under ejerez/MScThesis [17].

A. Experimental conditions

The experiment included only two conditions:

Condition 1 - nHG: The autonomous controller (hereafter referred to as "autopilot") actuates the crane directly and the joysticks only produce haptic feedback to simulate the notches and springs present on the current SSCV Thialf crane joysticks. Enabling autopilot requires the joystick for that axis to be centered and the autopilot button pressed. Pulling the joystick out of the center disables the autopilot for that axis and returns manual control to the user.

Condition 2 - HG: The autopilot controls the crane by physically actuating the joysticks, thus providing haptic guidance by producing an "autopilot notch", which captures the joystick and tracks it towards the commanded autopilot input. Autopilot can be enabled from any joystick position by pressing the autopilot button. Pulling the joystick out of the autopilot notch disables autopilot for that axis and returns manual control to the user, eliminating the autopilot notch.

B. Experimental setup: Haptic device

The custom-made haptic device can be seen in Figure 3 and consists of 3 single-axis joysticks capable of haptic feedback, which works by simulating a number of mechanical springs and notches.



Fig. 3: Custom haptic device for crane control. The joysticks control boom, slew, and hoist from left to right. The green button on each joystick also has an LED light and serves to engage and disengage the autopilot, as well as an autopilot status indicator.

To produce a default feeling as similar as possible to that of the real crane joysticks, each axis of the device provides haptic feedback by means of simulating one center notch; which snaps the joystick to the very center, two half-notches which push the joystick away from the center notch; these three notches effectively simulate two bumps located at ± 0.05, around the center, and depending on the axis either a free region and an end spring, or a return spring. This default haptic feedback is based on the joysticks currently used for the SSCV Thialf crane.

Additionally, when the autopilot is enabled with haptic guidance, a larger "autopilot notch" of half-width 0.30 overrides all other notches and springs, and effectively captures the joystick until the user pulls or pushes it out of the notch to reclaim control.

TABLE I: Simulated Haptic Feedback Elements

Axis	Element	Attractor	Region
All	Center notch	$x_a = 0$	$\lvert x \rvert \leq 0.05$
All	Outer notch	$x_a = \pm 0.10$	$0.05 < \lvert x \rvert \leq 0.10$
Slew	Return spring	$x_a = \pm 0.10$	$\lvert x \rvert > 0.125$
Boom and Hoist	Free range		$0.125 < \lvert x \rvert \leq 0.675$
Boom and Hoist	End spring	$x_a = \pm 0.70$	$\lvert x \rvert > 0.675$
Autopilot with HG	Autopilot notch	$x_a = x_{ap}$	$x_{ap} - 0.3 < x < x_{ap} + 0.3$

The positions and ranges of the notches and springs are shown in Table I, where x is the joystick position, x_a is the attractor center and x_ap is the commanded autopilot target. There are a few extra quirks that are important to mention: the area at $0.10 < |x| \leq 0.125$ is a transition in which the voltage is smoothed between the outer center notches and the free area or return spring, which is necessary to prevent instability at the edge of the outer notches. There is also a deadband at $|x| \leq 0.18$, meaning that the control input commanded by the joystick is zero until it is moved beyond ± 0.18, then increases linearly until the maximum control input at 1.

$$T_{\text{notch}} = T_{\max} \tag{1}$$

$$T_{\text{spring}} = 0.80 T_{\max} \left(\frac{\min(|x - x_a|, d_{\max})}{d_{\max}} \right) \tag{2}$$

$$T_{\text{free}} = 0 \tag{3}$$

To produce a feeling of a stiff mechanical notch, simulated notches are constant torque, as shown in Equation 1, where $T_{\max}$ is the maximum torque available. Springs, on the other hand, scale torque linearly as defined in Equation 2, where $x \in [-1, 1]$ is the joystick position; x_a is the notch center and $d_{\max}$ is the maximum extension of the spring. Note that in the region that is neither within a notch nor a spring, the "free" torque is zero, as shown in Equation 3.

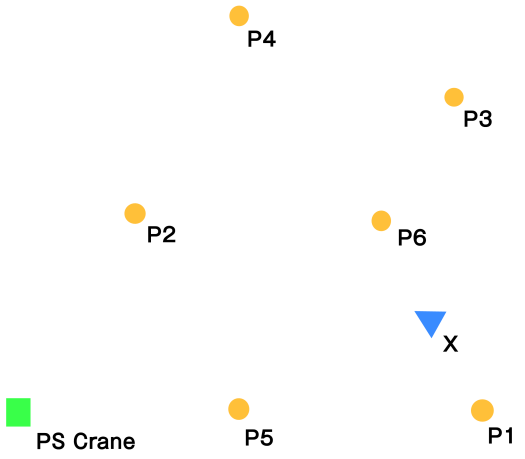


Fig. 4: Simplified representation of the autopilot control screen, showing a simplified top-view displaying the position of the crane **PS Crane**, estimated payload position **X**, and grillages **P** 1 through 6. In the real screen, actuation indicators on the bottom left corner allow the operator to monitor the commanded actuation levels when the autopilot is enabled with haptic guidance disabled.

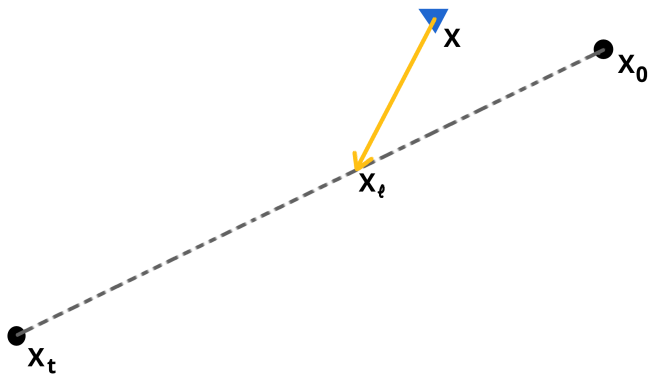


Fig. 5: Graphical representation of the autopilot in operation. The autopilot actuates the payload $\mathbf{X}$ towards the look-ahead target $\mathbf{X}_\ell$, which tracks along a straight line from the initial position $\mathbf{X}_0$ to the final target $\mathbf{X}_t$.

C. Experimental setup: Autopilot

An autopilot system was developed as an implementation of shared control. It consists of an autonomous controller that allows the operator to set a target point, and subsequently moves the payload along a straight line trajectory towards that target point. The decision was made to use straight line trajectories as the default controller behavior for simplicity, and due to the complexity of manually following a straight trajectory with the crane as discussed in section I. This autonomous controller uses a simple proportional controller to follow a target that tracks ahead of the payload along the straight line trajectory set by the user.

In order to share control with the autonomous controller, the operator can enable and disable the autopilot on each axis independently, and can also simply regain control by pulling or pushing the joystick out of its current position.

The autopilot was operated via a screen represented in Figure 4, and the buttons on the joysticks shown in Figure 3: Clicking on the screen set the clicked point as the final target as shown in Figure 5, and pressing the autopilot buttons on the joysticks switched the autopilot on, prompting it to start moving the payload along a straight line to the final target. Each axis could also be controlled manually (separately) by disabling the autopilot with either a press of the button or moving it out of the center or autopilot notch. Reengaging the autopilot was then done by simply pressing on the button again, and would result in the straight-line trajectory being recalculated from the current position of the payload.

$$\mathbf{s} = \mathbf{X}_t - \mathbf{X}_0, \tag{4}$$

$$\lambda = \begin{cases} 0, & \mathbf{s}^\mathsf{T}\mathbf{s} \le \varepsilon_\ell, \\ \operatorname{clip}\left(\dfrac{(\mathbf{X} - \mathbf{X}_0)^\mathsf{T}\mathbf{s}}{\mathbf{s}^\mathsf{T}\mathbf{s}}; 0, 1\right), & \text{otherwise}, \end{cases} \tag{5}$$

$$\mathbf{X}_c = \mathbf{X}_0 + \lambda\mathbf{s}, \tag{6}$$

$$\mathbf{r} = \mathbf{X}_t - \mathbf{X}_c, \tag{7}$$

$$\mathbf{X}_\ell = \begin{cases} \mathbf{X}_t, & \|\mathbf{r}\|_2 \le L_\ell, \\ \mathbf{X}_c + L_\ell \dfrac{\mathbf{r}}{\|\mathbf{r}\|_2}, & \text{otherwise}. \end{cases} \tag{8}$$

Equations 4 to 8 define the controller's instantaneous target. $\mathbf{X}_0$, $\mathbf{X}_t$, and $\mathbf{X}$ are the initial position, final target, and estimated payload position, respectively.

λ is used to project the closest point $\mathbf{X}_c$ on the segment s in Equations 5 and 6. $L_\ell = 3.0\,\mathrm{m}$ is the look-ahead distance, and $\mathbf{X}_\ell$ is the look-ahead, or instantaneous, target of the controller.

This effectively traces a straight line between the current position and the final target (which is user defined), then finds the closest point of that line to the payload position, and calculates an instantaneous target for the controller $L_\ell = 3.0\,\mathrm{m}$ ahead along the line, or at the final target (whichever is closest). The intention is to constrain the payload to following a path as close as possible to a straight line between the initial position and final target.

Note that the controller does not control the payload position directly, which would often oscillate too quickly for the crane to react. Instead, the payload's position is estimated via forward kinematics, in order to control the approximate steady-state position of the payload (that is, where it would be without sway), thus preventing excessive control inputs by the autopilot.

$$\theta = q_s - q_{s,0}, \tag{9}$$

$$\beta = q_b + \beta_0, \tag{10}$$

$$\rho = h_x + L_b \cos\beta, \tag{11}$$

$$\mathbf{J}_s = \rho \begin{bmatrix} -\sin\theta & \cos\theta & 0 \end{bmatrix}^\mathsf{T}, \tag{12}$$

$$\mathbf{b} = \begin{bmatrix} -\sin\beta & 0 & \cos\beta \end{bmatrix}^\mathsf{T}, \tag{13}$$

$$\mathbf{J}_b = L_b \mathbf{R}_z(\theta)\mathbf{b}, \tag{14}$$

$$\mathbf{J}_h = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^\mathsf{T}, \tag{15}$$

$$\mathbf{J} = \begin{bmatrix} \mathbf{J}_s & \mathbf{J}_b & \mathbf{J}_h \end{bmatrix}. \tag{16}$$

The task Jacobian used for the forwards and inverse kinematics is defined in Equations 9 to 16, where $i \in \{s, b, h\}$ are slew, boom, and hoist, respectively. A right-handed cartesian coordinate system is used, with z the vertical direction. q_s and q_b are the slew and boom angles, $q_{s,0}$ is the initial slew, θ is relative slew angle, β_0 is the initial boom angle, and β is the corrected boom angle. h_x is the boom hinge radial offset, L_b is the boom length, and ρ is the current horizontal radius.

$$\mathbf{P}_i = \begin{cases} \operatorname{diag}(1, 1, 0), & i = b, \\ \mathbf{I}_3, & i \in \{s, h\}, \end{cases} \tag{17}$$

$$\mathbf{d}_i = \mathbf{P}_i \mathbf{J}_i, \tag{18}$$

$$\widehat{\mathbf{d}}_i = \begin{cases} \mathbf{d}_i / \|\mathbf{d}_i\|_2, & \|\mathbf{d}_i\|_2 > \varepsilon_J, \\ \mathbf{0}, & \text{otherwise}, \end{cases} \tag{19}$$

$$e_i = a_i \widehat{\mathbf{d}}_i^\mathsf{T} (\mathbf{X}_\ell - \mathbf{X}). \tag{20}$$

Independent per-axis errors are calculated as shown in Equations 17 to 20. To isolate control of each direction of actuation, the controller uses the boom angle exclusively to control the radius, and the hoist length to control the payload height. $\mathbf{P}_i$ removes the boom's vertical component to isolate the radius-error, and $\mathbf{d}_i$ and $\widehat{\mathbf{d}}_i$ are the projected and normalized directions. $\varepsilon_J = 10^{-12}$ prevents division by zero, $a_i \in \{0, 1\}$ equals 1 for axis i if the autopilot is enabled for that axis, ensuring that the error for a given axis is discarded if that axis is being manually controlled, and e_i is the signed projected error.

$$\mathcal{A} = \{i \mid a_i = 1\}, \tag{21}$$

$$\mathcal{A}_{-i} = \mathcal{A} \setminus \{i\}, \tag{22}$$

$$i^\star = \begin{cases} \arg\max_{j \in \mathcal{A}} |e_j|, & |\mathcal{A}| \geq 2 \text{ and } \min_{j \in \mathcal{A}} |e_j| \geq e_{\text{stop}}, \\ \varnothing, & \text{otherwise}, \end{cases} \tag{23}$$

$$m_i = \begin{cases} 0, & i = i^\star \text{ or } \mathcal{A}_{-i} = \varnothing, \\ \max_{j \in \mathcal{A}_{-i}} |e_j|, & \text{otherwise}, \end{cases} \tag{24}$$

$$c_i = \begin{cases} 1, & i = i^\star, \\ \text{clip}\left(1 - \dfrac{m_i - e_{\text{slow}}}{e_{\text{stop}} - e_{\text{slow}}}; 0, 1\right), & \text{otherwise}, \end{cases} \tag{25}$$

$$p_i = \begin{cases} 0, & a_i = 0 \text{ or } |e_i| \leq e_{\text{db}}, \\ \min\left(\dfrac{|e_i| - e_{\text{db}}}{e_{\text{sat}} - e_{\text{db}}}, 1\right), & \text{otherwise}, \end{cases} \tag{26}$$

$$v_i = \text{sgn}(e_i) p_i c_i. \tag{27}$$

The P term of the proportional controller is calculated as shown in Equations 21 to 27. $\mathcal{A}$ contains only the axes for which the autopilot is enabled, $i^\star$ is the active axis with largest error, m_i is the largest error amongst other active axes, c_i is the cross-axis scaling, p_i is own-axis scaling, and v_i is the signed product. The constants are set as $e_{\text{db}} = 0.2\,\text{m}$, $e_{\text{sat}} = 3.0\,\text{m}$, $e_{\text{slow}} = 4.0\,\text{m}$, and $e_{\text{stop}} = 6.0\,\text{m}$.

Therefore for an arbitrary axis $i = a$, $p_a = 0$ for an axis error $e_a \leq 0.2\,\text{m}$, then p_a grows linearly until it reaches saturation at 1 when $e_a \geq 3.0\,\text{m}$, while the cross axis term c_i equals 1 if $m_i \leq 4.0\,\text{m}$, and reduces linearly until it reaches desaturation at 0 for $m_i \geq 6.0\,\text{m}$. This ensures that the faster axes will slow down if the slower ones are falling behind, which is particularly important because the slew axis can move much faster than the boom axis. If all active errors reach e_{stop}, $c_{i\star} = 1$, preventing the cross-axis term from fully locking up the autopilot.

No I or D terms were used, as the performance of the autopilot was already satisfactory for the experimental task with just a P term.

$$\ell_c = \ell_0 + \frac{|q_h|}{r_d} \tag{28}$$

$$T = 2\pi \sqrt{\frac{\ell_c + h_{CG}}{g}} \tag{29}$$

$$D_{\text{dec}} = v_{\text{blk}} T \tag{30}$$

$$\Delta\gamma_k = \frac{t_k - t_{k-1}}{2T} \tag{31}$$

$$\gamma_k = \begin{cases} \max(\gamma_{\min}, \gamma_{k-1} - \Delta\gamma_k), & \|(\mathbf{X}_t - \mathbf{X})_{xy}\|_2 \leq D_{\text{dec}}, \\ \min(1, \gamma_{k-1} + \Delta\gamma_k), & \text{otherwise} \end{cases} \tag{32}$$

$$\widetilde{u}_{i,k} = \gamma_k v_i \tag{33}$$

Command smoothing was added in order to minimize control-induced oscillations, and is implemented as per Equations 28 to 33, where ℓ_c is the hoist cable length, q_h, r_d the hoist drum position and radius, ℓ_0 the initial hoist cable length, $g = 9.81\,\text{m}\,\text{s}^{-2}$, and T the estimated sway period, where h_{CG} is the offset between the hook and the total payload CG. v_{blk} is the measured horizontal payload speed, D_{dec} is the estimated current deceleration distance, $\Delta\gamma_k$ is the ramp increment, γ_k is the speed factor, t_k is the iteration time, and $(\cdot)_{xy}$ selects horizontal components. $\widetilde{u}_{i,k}$ is the raw control inputs.

This smoothly accelerates the payload from standstill to full speed in $2T$, thus dampening the natural sway frequency. In order to decelerate in $2T$, the current deceleration distance D_{dec} is calculated every iteration given the current measured payload speed, and constant deceleration is started when the distance to the final target becomes smaller than D_{dec}. During acceleration $\gamma_0 = 0$, but during deceleration a lower bound of $\gamma_{\min} = 0.2$ is enforced, to prevent the payload from stopping moving without reaching the final target.

$$\alpha_k = 1 - \exp\left(-\frac{t_k - t_{k-1}}{\tau_h}\right) \tag{34}$$

$$w_{i,k} = \begin{cases} w_{h,k-1} + \alpha_k(\widetilde{u}_{h,k} - w_{h,k-1}), & i = h, \\ \widetilde{u}_{i,k}, & i \in \{s, b\}, \end{cases} \tag{35}$$

$$u_{i,k} = \text{clip}(w_{i,k}; -1, 1) \tag{36}$$

The final actuation is calculated as per Equations 34 to 36, in which a low-pass filter is applied to the hoist axis to reduce oscillation in the commanded input, with weight α_k, $\tau_h = 0.5\,\text{s}$, and $w_{h,0} = 0$, and finally the control inputs are clipped to $[-1, 1]$. $w_{i,k}$, and $u_{i,k}$ are the filtered and normalized control inputs.

$$\left|\frac{\text{d}u_{i,k}}{\text{d}t}\right|_{max} = \frac{1}{T} \tag{37}$$

There is one final layer of sway prevention as shown in Equation 37. The autopilot has a rate limiter which bounds the absolute value of input changes to $1/T$, ensuring that changes of direction at full speed won't excessively excite the natural sway frequency. However, this addition was implemented near the end of the experiment phase, and was only present for the last 3 participants.

D. Experimental setup: Simulator

In order for the task to be realistic and challenging, Heerema's crane simulator was used. This is a high-fidelity simulator used to train crews and practice crane operations.

Within the simulator, the participant was given control of the port-side crane of the SSCV Sleipnir, with the hook block attached on the main hoist, itself attached to an arbitrarily selected wind turbine TP (Top Piece), as seen in Figure 6, selected because it is a payload type that is commonly lifted by these types of cranes. To increase the realism, the measured payload position was taken to be the position of the block, which would be easily measurable, as the main hoist block always remains attached to the crane.



(a) Deck layout of the SSCV Sleipnir.



(b) Top Piece (TP) attached to the hook (Block).

Fig. 6: Deck layout (left) and payload (right) used for the experiment. Note the obstacles placed on the deck. The TP Grillages are the radial structures placed at several locations. The reference payload location used by the automatic controller was actually the position of the Block rather than the TP, since in reality it would be much easier to measure the position of the Block given that, unlike the payload, the Block is rarely swapped or changed.

Due to a limited field of view when displaying the simulator on a conventional monitor, a free-floating view was instead provided, and the participants were allowed to request perspective adjustments to be able to better see what the payload was doing.

E. Task

The experimental task was a close analogue to a pick-and-place task, a frequently utilized task-style in research involving teleoperation, although the payload was permanently attached to the crane block rather than having to be physically picked up.

To provide a variety of targets and possible paths, six (6) different TP Grillages were placed on the Sleipnir's deck as shown in Figure 6, with the "pick" analogue being to lift the TP from the current grillage, and the "place" analogue being to lower the TP onto the target grillage. Each task trial simply consisted of moving the TP from one grillage to another, using the autopilot whenever desired.

In order to introduce more complexity to the task, and to force the participants to follow more complex trajectories with the autopilot or use some degree of manual control, a number of obstacles were also placed on the deck, obstructing the path between each grillage pair and necessitating that the participant take targeted actions to go either above or around the obstacles. Since the autopilot only moves the payload along a straight path, multiple successive targets or manual intervention were necessary to move around obstacles; although the exact strategy used to avoid each obstacle was left to the discretion of the participant.

Thus, the task can be divided roughly into three phases:

- **Lifting**: Lifting the payload from the initial grillage.
- **Moving**: Moving the payload to the target grillage, while avoiding any obstacles along the way.
- **Lowering**: Placing the payload down at the target grillage.

Only the moving phase, consisting of moving the payload to the target location, was of significant interest to the present study, as it was the only phase for which the autopilot was useful.

$$z_m \geq 3.8\,\mathrm{m} \tag{38}$$

$$\|\mathbf{X} - \mathbf{X}_t\|_2 \leq 0.5\,\mathrm{m} \tag{39}$$

$$d_{xy}^{\mathrm{goal}} = \sqrt{(x_m - x_g)^2 + (y_m - y_g)^2} \tag{40}$$

$$d_{xy}^{\mathrm{goal}} \leq 4.0\,\mathrm{m}, \qquad z_m \leq 0\,\mathrm{m} \tag{41}$$

A number of different criteria were used to define successful completion of a task trial. Equation 38 describes the condition for the autopilot to become available, that is, the payload has to be at least $3.8\,\mathrm{m}$ off the deck. z_m is calibrated to zero when the simulator exercise is first loaded.

The autopilot convergence threshold is as defined in Equation 39. When this criterion is satisfied, the autopilot automatically switches off.

As per Equation 40, the euclidean distance between the payload and the target d_{xy}^{goal} is calculated in the XY plane. The threshold for trial completion is described in Equation 41, and when these two conditions are met, the trial is automatically ended.

F. Experimental design

The experiment was a within-subject design in which 26 participants completed multiple task trials under each of the two conditions. The participants were aged between 20 and 49 years old (with 15 between 20 and 29, 8 between 30 and 39, and 3 between 40 and 49), 22 of whom were male and 4 female, and only 2 reported previous crane experience. Each participant filled and signed an Informed Consent Form prior to participating, and the experiment was approved by the Human Research Ethics Committee of the TU Delft.

The order of the conditions was alternated between participants in order to mitigate order effects. Each trial consisted of moving from one grillage to another, along a pre-determined route that requires navigating a combination of moderately challenging and challenging obstacles.

The complete experiment consisted initially of 10 trials in the order:

P1→P5→P4→P1→P3→P4→P2→P5→P3→P1→P6

However, after the first 4 participants, the number of trials was reduced to just 6 in an effort to reduce the amount of time it took to perform the full experiment, and thus be able to find a larger number of people willing to participate. The experiment then consisted of 6 trials in the order:

P1→P6→P4→P3→P1→P2→P5

In both 10 and 6 trial experiments, the first half of the trials (1-5 and 1-3, respectively) were performed under the first condition, and the second half (6-10 and 4-6, respectively) were performed under the second condition.

To isolate the effects of haptic guidance, the only difference between the conditions was whether the autopilot physically actuated the joysticks (HG) or not (nHG) when enabled. The functioning of the autopilot and crane, and task procedure, were exactly the same aside from that, requiring the human to manually lift and place down the load, as well as to decide what trajectory to follow between grillages and set the autopilot targets to achieve their desired trajectory.

G. Experimental Procedure

Participants initially were given an explanation about how the device, simulator and task worked. They were simply instructed to lift the payload up manually, then move it to the target grillage, using the autopilot as much as desired, and finally lower the payload manually onto the target grillage. Landing precisely was not required, and the successful landing thresholds were lax to prevent the landing phase from taking up a disproportionate amount of time.

Before beginning the measured trials, participants performed a "warmup" trial to familiarize themselves with the operation of the crane and, in particular, the autopilot. The first half of the warmup trial was performed under the first condition and consisted of:

- Lift up the payload from P1.
- Test the three control axes: boom, slew and hoist.
- Utilize the autopilot to move towards P5.
- Modify the autopilot trajectory by taking manual control of the slew axis and then returning it to autopilot.
- Modify the autopilot trajectory by setting a new final target at a different location.

The second half of the warmup trial was performed under the second condition and consisted of:

- Use the autopilot to return to P1.
- Lower the payload manually onto P1.

Once the warmup had been completed, the measured trials began immediately. The first half of the trials were performed under the first condition, and upon finishing them, the participant was prompted to fill in 2 questionnaires. The second half of the trials were then performed under the second condition, and upon finishing them, the participants were again asked to fill in the 2 questionnaires.

H. Quantitative measures

The metrics defined per proposed hypothesis comprised the following measures:

- **H1. Task performance and safety**: Of which performance is defined in terms of task completion time, which quantifies how quickly the task was finished; payload trajectory length, which quantifies how long the path taken was; and control jitter, which quantifies how smooth the control inputs were. Safety was defined in terms of collision count, which quantifies how many times the payload collided with one of the obstacles on the deck; and sway spectral energy, which quantifies how much the payload swayed (or oscillated) during the task.
- **H2. Subjective experience**: Of which the operator's self-perception was measured by the Human Perspective criteria of the ACES questionnaire developed by Mol et al. [15], and the operator's perception of the autopilot was measured by the Robot Perspective criteria of the ACES questionnaire [15].
- **H3. Perceived workload**: Measured by means of the unweighted NASA-TLX questionnaire developed by NASA [18].

With i, c, and r denoting the participant, condition, and trial, and s and e denoting the start and end of the moving phase of the task, and k denotes an iteration or sample, the measures were calculated as follows:

$$T_{icr} = t_{icr,e} - t_{icr,s} \tag{42}$$

$$L_{icr}^{\text{hook}} = \sum_{k=s+1}^{e} \|\mathbf{h}_{icr,k} - \mathbf{h}_{icr,k-1}\|_2 \tag{43}$$

$$C_{icr} = n_{icr,e}^{\text{col}} - n_{icr,s}^{\text{col}} \tag{44}$$

$$J_{icr} = T_{icr} \sum_{a\in\{s,b,h\}} \text{trapz}_{f_0,\ldots,f_M} \left((2\pi f)^2 \widehat{P}_{u_a u_a}(f)\right) \tag{45}$$

Where Equation 42 shows the **task completion time** in simulator time, Equation 43 shows the **trajectory length** of the payload with $\mathbf{h}_{icr,k}$ the measured block position, Equation 44 shows the number of collisions with $n_{icr,k}^{\text{col}}$ the **cumulative collision count**, and Equation 45 showing the **control input jitter** energy with $\widehat{P}_{u_a u_a}(f)$ the one-sided Welch density of the control inputs. Note that the control inputs were sampled at $5\,\text{Hz}$, with derivatives larger than $0.5/\text{s}$ discarded to account for impulse artifacts introduced by the autopilot in nHG mode when taking manual control before the rate limiter was implemented. Welch analysis used a Hann window, constant detrending, $n_{\text{perseg}} = \min(256, N_g)$, 50% overlap, and density scaling.

$$\mathbf{r}_{icr,k} = \mathbf{g}_{icr,k} - \mathbf{h}_{icr,k} \tag{46}$$

$$r_{xy,icr,k} = \sqrt{r_{x,icr,k}^2 + r_{y,icr,k}^2} \tag{47}$$

$$\theta_{icr,k} = \text{atan2}\left(r_{xy,icr,k}, -r_{z,icr,k}\right) \tag{48}$$

$$\begin{bmatrix}\theta_{x,icr,k}\\ \theta_{y,icr,k}\end{bmatrix} = \frac{\theta_{icr,k}}{r_{xy,icr,k}} \begin{bmatrix}r_{x,icr,k}\\ r_{y,icr,k}\end{bmatrix} \tag{49}$$

$$E_{icr}^{\text{sway}} = T_{icr} \sum_{a\in\{x,y\}} \text{trapz}_{f_1,\ldots,f_M} \left(\widehat{P}_{\theta_a\theta_a}(f)\right) \tag{50}$$

The **sway spectral energy** was calculated as per Equations 46 to50, where $\mathbf{g}_{icr,k}$ is the measured TP center of gravity. The angle of sway is considered zero when the TP is aligned with the Z axis. In order to exclude the steady state energy, the spectral energy was integrated from f_1, which is the first positive Welch frequency bin.

ACES is a subjective opinion questionnaire developed by researchers at TU Delft and contains two perspectives, a **human perspective** and an **autopilot or robot perspective**. The ACES questionnaire is comprised of 4 questions for each perspective in an effort to quantify perceived **autonomy**, **competence**, **engagement**, and **usefulness**. [15]

To assess reliability, Cronbach's alpha was calculated separately for each experimental condition. For the human perspective internal consistency was $\alpha = 0.59,\ 95\%$ CI $[0.26, 0.80]$ under condition nHG and $\alpha = 0.60,\ 95\%$ CI $[0.27, 0.80]$ under condition HG. For the autopilot perspective internal consistency was $\alpha = 0.63,\ 95\%$ CI $[0.33, 0.82]$ under condition nHG and $\alpha = 0.65,\ 95\%$ CI $[0.37, 0.83]$ under condition HG.

The questions are scored on a Likert scale and are:

- I felt that X had a high level of autonomy while completing the task. [15]
- I felt that X had the necessary skills and knowledge to complete my/its part of the task. [15]
- I felt X was useful in performing my/its part of the task. [15]
- I felt X was engaged in the role. [15]

Where X is "I" for the human perspective and "the autopilot" for the autopilot perspective.

The ACES composite scores were calculated by simply obtaining the mean of the individual criteria.

NASA-TLX is a subjective opinion questionnaire which measures perceived workload, and consists of 6 questions which measure **Mental Demand**, **Physical Demand**, **Temporal Demand**, **Performance**, **Effort**, and **Frustration**. [18]

The questions are scored on a 100-point scale from "Very Low" to "Very High" (or "Perfect" to "Failure" for the fourth question), with 5-point steps, and are:

- How mentally demanding was the task? [18]
- How physically demanding was the task? [18]
- How hurried or rushed was the pace of the task? [18]
- How successful were you in accomplishing what you were tasked to do? [18]
- How hard did you have to work to accomplish your level of performance? [18]
- How insecure, discouraged, irritated, stressed or annoyed were you? [18]

As with ACES, the composite score for NASA-TLX was simply calculated as the mean of all the criteria.

I. Statistical analysis

To provide simple, robust analysis of the results, the condition effects were evaluated using paired tests. Normality was assessed for the paired HG minus nHG differences using the Shapiro-Wilk test. If normality was met with $p > 0.01$, a paired t-test was used, with effect size reported as Cohen's d_z. If normality was rejected with $p \leq 0.01$, a Wilcoxon signed-rank test was applied, with effect size reported as the matched-pairs rank-biserial correlation.

Collision outcomes always used the Wilcoxon test because of their sparse, discrete distribution. All tests were two-sided with $\alpha = 0.05$, and Holm-Bonferroni correction was applied within hypothesis metric families (that is, within the measures assigned to each metric of each hypothesis) to account for multiplicity. Mean paired differences and 95% confidence intervals were estimated using 10,000 participant-level percentile-bootstrap resamples.

A post-hoc analysis was conducted to measure the robustness of the statistical analysis. It performed alternative paired t-tests, Wilcoxon signed-rank tests, sign tests, and leave-one-participant-out analyses. Hypothesis-wise and metric-wise aggregate multiplicity corrections were additionally calculated. Potential effect modification was examined by condition order, six vs. ten trial experiments, and the rate-limiter implementation that affected the last 3 participants.

III. Results

The results can be seen in Table II, and have a few caveats that are important to mention. Firstly, data loss caused by various bugs means that not all the data is available for all participants: there are 18 participants with recorded task performance and task safety metrics (except for sway, for which only 14 are available). While the ACES data is available for all 26 participants, NASA-TLX is only available for 23.

Secondly, there were two ways in which the participants can be divided into two separate cohorts:

- 4 participants performed a 10-trial experiment, while the other 22 performed a shorter 6-trial experiment.
- 23 participants used the autopilot without an input change rate limiter, while the last 3 used it with input rate limiting.

While half the participants experienced Order 1 (nHG $\rightarrow$ HG) and the other half experienced Order 2 (HG $\rightarrow$ nHG), this half-and-half counterbalancing did not necessarily apply across cohorts, necessitating sensitivity testing to verify results.

A. Task Performance and Safety

Regarding task performance and safety, related to H1, no significant effects were observed for performance measured in completion time, 3D trajectory length and control input jitter; nor for safety measured in collision count and sway spectral energy. See Figure 7.

B. Subjective Experience

Subjective operator experience was evaluated using the ACES questionnaire and is displayed in Figure 8.

The operator's self-perception used the human perspective. A significant difference was found in the composite human ACES scores ($t(25) = 2.57,\ p = 0.0165$), suggesting an overall preference for HG over nHG. Looking into the individual criteria shows a significant difference only for the uncorrected perceived human competence, which does not remain significant after within-family Holm-Bonferroni correction.

Subjective autopilot perception was evaluated using the autopilot (robot) perspective. A strongly significant difference was found in the composite autopilot ACES scores ($W_{SR} = 22.5,\ p = 0.00044$) as shown in Figure 8, suggesting, again, an overall preference for HG over nHG. Looking into the individual criteria, shown also in Figure 8, a significant difference is also present for the uncorrected perceived autopilot autonomy and engagement criteria, of which only autopilot engagement remains significant after within-family Holm-Bonferroni correction ($t(25) = 3.67,\ p = 0.00458$), suggesting a perceived increase of the autopilot's engagement for HG compared to nHG.

It is worth mentioning that a significant order effect was found for the autopilot engagement criterion, with $p = 0.045$, suggesting - on average - a reduction of the perceived autopilot engagement on the second condition compared to the first.

C. Perceived Workload

Perceived workload was evaluated with the unweighted version of NASA-TLX, and can be seen in Figure 9.

The results show a significant decrease in overall scores ($t(22) = -2.55,\ p = 0.0183$), suggesting a decreased perceived workload for HG compared to nHG. Looking into the individual criteria shows only a significant difference for the uncorrected mental demand, which does not remain significant after within-family Holm-Bonferroni correction.

D. Participant comments

The comments were varied and made it clear that different participants held very different views on the advantages or disadvantages of HG compared with nHG. It is important to note that only 14 of the 26 participants decided to leave a comment, so the comments left may not reflect the global opinion. An overview of the statements asserted, as well as how many participants commented some version of a given statement, can be seen in Table V.

Most of the commenting participants agreed that HG increased the transparency of the autopilot actuation. Almost a third stated that they explicitly preferred to have HG over nHG, about a third stated they felt HG helped with more precise corrections or a better sense of control, while a minority stated a perceived reduction in workload with HG compared to nHG, or the potential of HG to aid in learning to operate the crane or autopilot. A few participants stated feeling discomfort with the lack of haptic guidance under nHG, while a few stated they felt little to no difference between the conditions and one thought that the haptic guidance was not necessarily useful because it did not convey whether what the autopilot was doing was correct or not.

E. Sensitivity analysis

All of the reported outcomes were tested with paired-t, Wilcoxon, sign test and leave-one-out. NASA-TLX and human ACES were no longer significant under the sign test, but no other changes were observed.

Additionally, all outcomes were tested based on two different cohort categorizations: 10 vs 6 trial participants and autopilot rate limited vs. no rate limiting participants. No significant interaction effects were found for either cohort split for any of the reported results.

F. Multiplicity correction

To check the potential effects of multiplicity, metric-wise and hypothesis-wise multiplicity correction was performed. Aggregation was based on subsection II-H, as shown in Table III and Table IV.

The paired p-values were aggregated using Fisher's combined probability test as shown in Equation 51, evaluated against a χ^2 distribution with $2k$ degrees of freedom, which produced a single p-value per aggregated composite. However, the components of a composite measure the same construct (metric), thus are correlated, which violates Fisher's independence assumption; so Tippett's minimum-p combination, shown in Equation 52, was also calculated.

TABLE II: Quantitative measure statistics

Outcome	n	Mean difference (HG-nHG) [95% CI]	Test	Effect size	Raw p	Holm p
H1 - Task performance						
Completion time	18	8.4 [-8.18, 25.8]	$t(17) = 0.945$	$d_z = 0.223$	0.358	0.886
3-D trajectory length	18	-13.7 [-38.1, 11.2]	$t(17) = -1.08$	$d_z = -0.255$	0.295	0.886
Control-input jitter energy	18	0.0698 [-0.374, 0.496]	$t(17) = 0.304$	$d_z = 0.0717$	0.765	0.886
H1 - Task safety						
Collisions	18	0.0833 [-0.0148, 0.183]	$W_{\mathrm{SR}} = 12$	$r_{\mathrm{rb}} = 0.564$	0.107	0.214
Sway spectral energy	14	-0.137 [-0.337, 0.0871]	$t(13) = -1.23$	$d_z = -0.328$	0.241	0.241
H2 - Subjective operator experience						
Human ACES (overall)	26	0.125 [0.0327, 0.22]	$t(25) = 2.57$	$d_z = 0.504$	**0.0165**	**0.0165**
Autonomy	26	0.127 [-0.0231, 0.273]	$t(25) = 1.63$	$d_z = 0.319$	0.116	0.232
Competence	26	0.177 [0.0192, 0.335]	$t(25) = 2.17$	$d_z = 0.425$	**0.0398**	0.159
Usefulness	26	0.00769 [-0.138, 0.162]	$t(25) = 0.0991$	$d_z = 0.0194$	0.922	0.922
Engagement	26	0.188 [0.00385, 0.373]	$t(25) = 1.96$	$d_z = 0.385$	0.0608	0.182
H2 - Subjective autopilot perception						
Autopilot ACES (overall)	26	0.159 [0.0798, 0.248]	$W_{\mathrm{SR}} = 22.5$	$r_{\mathrm{rb}} = 0.837$	**0.00044**$^{\circ}$	**0.00044**$^{\circ}$
Autonomy	26	0.196 [0.05, 0.358]	$t(25) = 2.44$	$d_z = 0.479$	**0.022**	0.0659
Competence	26	0.0231 [-0.112, 0.131]	$W_{\mathrm{SR}} = 47.5$	$r_{\mathrm{rb}} = 0.379$	0.169°	0.246°
Usefulness	26	0.0808 [-0.0115, 0.188]	$W_{\mathrm{SR}} = 44$	$r_{\mathrm{rb}} = 0.425$	0.123°	0.246°
Engagement	26	0.335 [0.162, 0.515]	$t(25) = 3.67$	$d_z = 0.72$	**0.00115**	**0.00458**
H3 - Perceived workload						
NASA-TLX (overall)	23	-5.04 [-9.09, -1.49]	$t(22) = -2.55$	$d_z = -0.532$	**0.0183**	**0.0183**
Mental Demand	23	-9.78 [-18.3, -2.39]	$t(22) = -2.42$	$d_z = -0.504$	**0.0243**	0.146
Physical Demand	23	-1.96 [-6.52, 2.39]	$W_{\mathrm{SR}} = 25$	$r_{\mathrm{rb}} = -0.242$	0.472°	0.488°
Temporal Demand	23	-4.78 [-11.1, 1.3]	$t(22) = -1.5$	$d_z = -0.312$	0.149	0.488
Performance	23	-5 [-10.2, 0.652]	$t(22) = -1.77$	$d_z = -0.369$	0.091	0.455
Effort	23	-3.48 [-11.1, 4.57]	$W_{\mathrm{SR}} = 34$	$r_{\mathrm{rb}} = -0.433$	0.137°	0.488°
Frustration	23	-5.22 [-11.5, 0.87]	$t(22) = -1.61$	$d_z = -0.335$	0.122	0.488

$\circ$ Indicates violation of normality (Shapiro-Wilk).

TABLE III: Metric-wise multiplicity correction

Metric	Hypothesis	Components	Fisher p	Tippett p	Holm p
Performance	H1	Completion time; 3-D trajectory length; Control-input jitter energy	0.54	0.65	0.54
Safety	H1	Collisions; Sway spectral energy	0.12	0.203	0.241
Subjective operator experience	H2	Human ACES	**0.016**	**0.016**	0.066
Subjective autopilot perception	H2	Autopilot ACES	**0.00044**	**0.00044**	**0.002**
Perceived workload	H3	NASA-TLX	**0.018**	**0.018**	0.066

TABLE IV: Hypothesis-wise multiplicity correction

Metrics	Hypothesis	Components	Fisher p	Tippett p	Holm p
Performance; Safety	H1	Completion time; 3-D trajectory length; Control-input jitter energy; Collisions; Sway spectral energy	0.263	0.432	0.263
Subjective operator experience; Subjective autopilot perception	H2	Human ACES; Autopilot ACES	**0.000093**	**0.000879**	**0.000279**
Perceived workload	H3	NASA-TLX	**0.018**	**0.018**	**0.037**

TABLE V: Participant comments

Trend	Participants	Percentage of commenters
Haptics made crane/autopilot activity more observable or transparent	9	64.3%
Haptics explicitly described as advantageous, useful, or preferred	4	28.6%
Haptics aided small corrections, intervention, or sense of control	3	21.4%
Haptics reduced perceived workload	2	14.3%
Having no haptic guidance caused confusion or discomfort	2	14.3%
Haptics could support learning or understanding crane/autopilot behaviour	2	14.3%
Little or no difference was perceived between conditions	2	14.3%
Haptic cues did not clearly show whether an input was appropriate	1	7.1%
Autopilot was preferred for boom/slew but manual control for hoist	1	7.1%

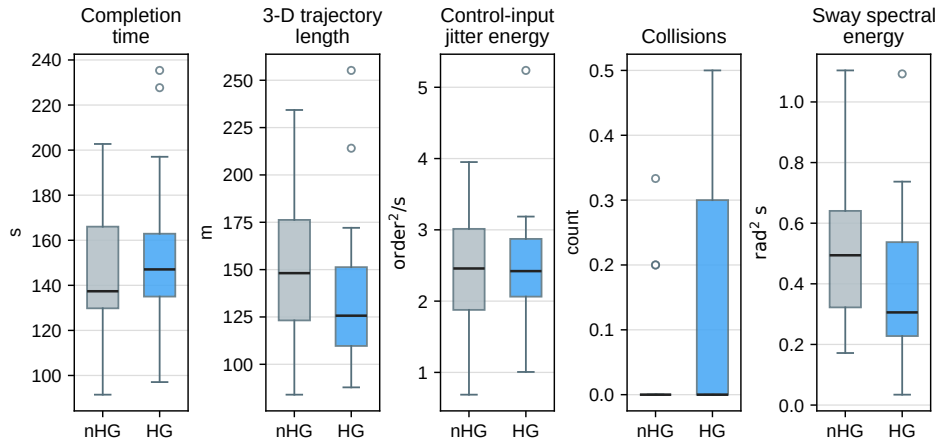


Fig. 7: H1 - Task performance and safety paired results, with performance measured in completion time, 3D trajectory length and control input jitter; and safety measured in collisions and sway spectral energy.

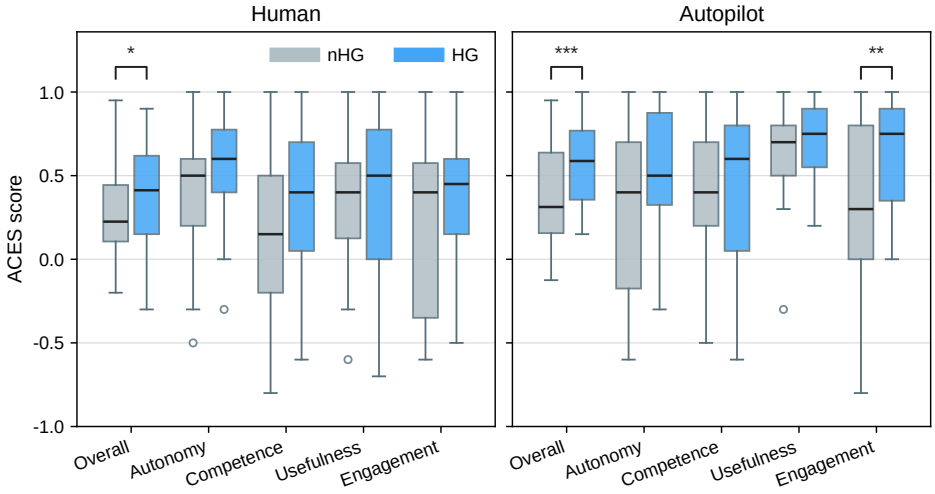


Fig. 8: H2 - Subjective experience paired results, with the operator's self-perception measured by the human ACES perspective, and the autopilot perception measured by the autopilot ACES perspective. * denotes significance between the two conditions.

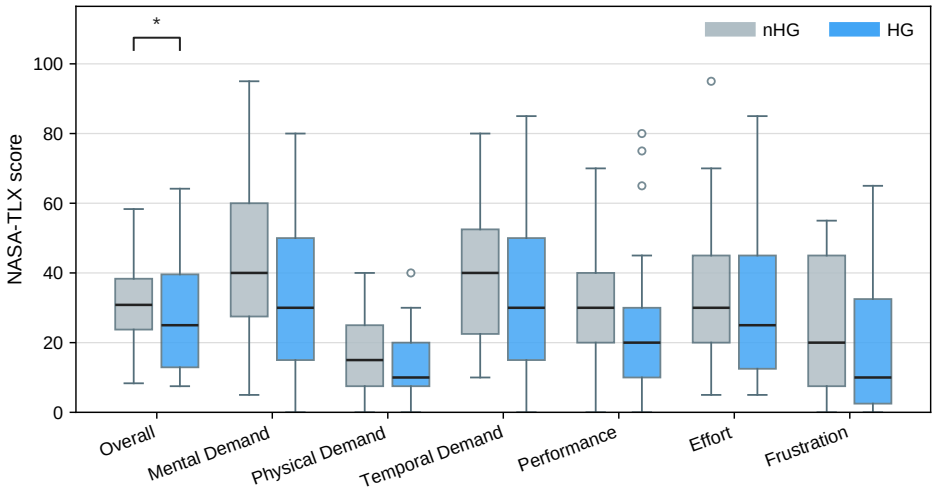


Fig. 9: H3 - Perceived workload paired results, measured with NASA-TLX. * denotes significance between the two conditions.

$$X^2 = -2\sum_{i=1}^{k} \ln(p_i) \tag{51}$$

$$p_T = 1 - (1 - p_{\min})^k \tag{52}$$

Only the subjective autopilot perception (autopilot ACES) remains significant under the stricter metric-wise multiplicity correction, as shown in Table III. This does not mean the results for other metrics should be automatically discarded, but does encourage interpreting them with caution. With hypothesis-wise multiplicity correction, both H2 and H3 remain significant, as shown in Table IV, with H2 more strongly so.

IV. Discussion

The present study investigated whether the addition of haptic guidance to a shared control approach is advantageous in the context of the operation of marine boom cranes. The results provide moderate but promising evidence that the addition of haptic guidance is advantageous, at least in terms of the subjective experience of the operators.

Haptic guidance did not improve task performance and safety: No significant differences were observed in completion time, trajectory length or control jitter between using haptic guidance or not. This is a fairly logical outcome, as the autopilot used by the participants to perform the tasks functioned exactly the same under both conditions except for the addition of haptic guidance, thus users experienced very similar degrees of automation and used very similar strategies to complete the tasks. This contrasts with Chu et al. [6], which found an improvement in sway mitigation, but which used a completely different type of haptic device and haptic guidance approach (letting the user feel the scaled-down end effector forces, an approach more typical of teleoperation).

Safety-wise, no significant improvement was observed in number of collisions or payload sway. As with the task performance, this is reasonable and relates to the fact that the autopilot provided the same degree of available automation under both conditions.

These outcomes also contrast with the findings in automotive studies like that by Mulder et al. [13] and Benloucif et al. [14], which found evidence of increased safety and performance, respectively, with the addition of haptic guidance. These differences, however, are likely to stem from the fact that driving a vehicle is not comparable with operating a crane, both in terms of the control mechanism and in terms of the task that is performed by the operator.

Haptic guidance improved operator's subjective experience: The human ACES results provide modest but still significant evidence that the experience of the operator improves on average with the addition of haptic guidance. Mixed opinions and a low statistical power due to a limited number of participants, however, mean that no conclusions can be strongly drawn about whether the operator's self-perceived autonomy, competence, usefulness or engagement are improved or not. Based on the comments, some of the participants felt it enabled them to perform small corrections more precisely or gave them a greater sense of control, which could be correlated to the observed trend towards a higher self-perception of competence, although this trend did not show significance after multiplicity correction.

Autopilot ACES scores suggest a moderate but significant improvement in the perception of the autopilot with haptic guidance, compared to without. Many of the participants (more than a third) mentioned a greater perceived transparency in autopilot operation with haptic guidance, and trends hint towards increased perceived autopilot autonomy and engagement, although only the engagement remains significant after multiplicity correction. Improved transparency is likely to be related to a feeling that the autopilot is more engaged with the task, thus explaining the increase in perceived engagement.

The results are consistent with much of the literature across different fields, including automotive studies like that by Benloucif et al. [14], as well as haptic shared control studies like those by Mol et al. [15] and Rodríguez-Sedano et al. [16], which found the operator's subjective experience to be affected (in varying degrees) by the addition of different types of haptic guidance. However, although the results are promising, they do not quite match the significance and scale of the results obtained by Mol et al. [15] - who developed the ACES questionnaire - despite the similarities between the two studies (same number of participants, similar statistical approaches and comparable experimental design). Among the potential reasons are the possibility that the differences between the conditions would be relatively small, or that ACES would not be well suited for this condition pair.

Haptic guidance decreased perceived workload: Although different participants had different opinions, the overall unweighted NASA-TLX score suggests a modest but significant decrease in perceived workload. The only individual criterion which shows any sort of defined trend by itself, however, is the perceived mental workload, which does not remain significant after multiplicity correction. Some participants also explicitly mentioned feeling like having haptic guidance decreased the mental workload.

A reduction in perceived workload is also consistent with state-of-the-art implementations of haptic guidance, as suggested by Mol et al. [15], although partially contradicting Rodríguez-Sedano et al. [16], in which the authors report higher mental workload but greater ease of use.

Overall, no improvement in telemetry-recorded metrics but an improvement in perceived experience and workload: The measured performance of the participants did not differ significantly with and without haptic guidance. The subjective aspect of the task, however, seemed to be improved by the addition of haptic guidance. This introduces the possibility that performance could eventually be affected in cases in which prolonged or sustained crane operation is performed, as the subjective aspects of operation (self-perception, autopilot perception, workload) play a large role in the effects of fatigue.

An improvement in certain aspects of the operator experience with the addition of haptic guidance, without necessarily improving all aspects of the subjective experience significantly,

is consistent with the results obtained by Rodríguez-Sedano et al. [16], suggesting the interaction between operators and the autonomous controller, in the context of haptic shared control, is a rather complex and dynamic process, which is not yet fully understood.

Recommendations: Further research is potentially promising. Haptic guidance is well-suited to implementation or retrofit onto crane control interfaces, as they most commonly use a combination of joysticks, and the implementation of haptic guidance in addition to task automation has some potential to reduce workload and enhance the experience of the operator.

Furthermore, using haptic guidance as an aid for training operators could also demonstrate to be a useful resource, hinted by the comments of some of the participants in the present study, and backed by the propositions of other studies such as the one by Huegel et al. [19].

Limitations: The participants were mostly male (22 of 26), between 20 and 29 years old (15 of 26) and had no previous crane experience (24 of 26), thus the results of this study do not necessarily generalize to the general human population, nor to the population of crane operators. The differences across cohorts (10 vs. 6 trials, input rate limiting vs. no rate limiting) and data loss reduce the certainty of the results. Additionally, the Cronbach's alpha values for ACES were lower than the conventionally expected 0.7, indicating low internal consistency.

Holm-Bonferroni is valid under arbitrary dependence, therefore the results of the questionnaires' criteria are rather conservative. However, Fisher and Tippett aggregations assume independent p-values, meaning the hypothesis and metric aggregate multiplicity corrections may instead be rather optimistic.

Due to time constraints, it was not possible to have a crane operator test the finalized autopilot system or the haptic guidance, which constitutes an extremely valuable perspective on their functioning but is unfortunately missing from this study. While the task was reasonably realistic, the crane was set to move faster than it is allowed in reality, and the task did not necessarily represent the majority of the tasks crane operators perform, which can be very varied.

Additionally, prolonged use may yield different results than observed in the short experiment implemented in this study, due to fatigue and other effects. It is also important to note that the ACES questionnaire is not psychometrically validated, and thus the ACES results should be interpreted with caution.

REFERENCES

[1] "Sleipnir | Heerema." [Online]. Available: https://www.heerema.com/heerema-marine-contractors/fleet/sleipnir

[2] L. Ramli, Z. Mohamed, A. M. Abdullahi, H. I. Jaafar, and I. M. Lazim, "Control strategies for crane systems: A comprehensive review," *Mechanical Systems and Signal Processing*, vol. 95, pp. 1–23, Oct. 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0888327017301425

[3] M. Ahmad, R. R. Ismail, M. Ramli, N. Hambali, and N. Noordin, "Active sway control of a lab-scale rotary crane system," in *2010 The 2nd International Conference on Computer and Automation Engineering (ICCAE)*, vol. 1, Feb. 2010, pp. 408–412. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/5451923

[4] A. Takemoto, K. Yano, and K. Terashima, "Obstacle Avoidance Control System of Rotary Crane Using Proposed Haptic Joystick." IEEE Computer Society, Mar. 2005, pp. 662–663. [Online]. Available: https://www.computer.org/csdl/proceedings-article/whc/2005/23100662/12OmNx5Yv73

[5] K. Nakazono, K. Ohnishi, H. Kinjo, and T. Yamamoto, "Vibration control of load for rotary crane system using neural network with GA-based training," *Artificial Life and Robotics*, vol. 13, no. 1, pp. 98–101, Dec. 2008. [Online]. Available: https://doi.org/10.1007/s10015-008-0586-5

[6] Y. Chu, H. Zhang, and W. Wang, "Enhancement of Virtual Simulator for Marine Crane Operations via Haptic Device with Force Feedback," in *Haptics: Perception, Devices, Control, and Applications*, F. Bello, H. Kajimoto, and Y. Visell, Eds. Cham: Springer International Publishing, 2016, pp. 327–337.

[7] N. Sun, T. Yang, Y. Fang, B. Lu, and Y. Qian, "Nonlinear Motion Control of Underactuated Three-Dimensional Boom Cranes With Hardware Experiments," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 3, pp. 887–897, Mar. 2018. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8047336

[8] J. Huang, E. Maleki, and W. Singhose, "Dynamics and swing control of mobile boom cranes subject to wind disturbances," *IET Control Theory & Applications*, vol. 7, no. 9, pp. 1187–1195, 2013. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1049/iet-cta.2012.0957

[9] E. Maleki and W. Singhose, "Dynamics and Control of a Small-Scale Boom Crane," *Journal of Computational and Nonlinear Dynamics*, vol. 6, no. 031015, Mar. 2011. [Online]. Available: https://doi.org/10.1115/1.4003251

[10] E. Arnold, O. Sawodny, J. Neupert, and K. Schneider, "Anti-sway system for boom cranes based on a model predictive control approach," in *IEEE International Conference Mechatronics and Automation, 2005*, vol. 3, Jul. 2005, pp. 1533–1538. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/1626783

[11] D. A. Abbink, M. Mulder, and E. R. Boer, "Haptic shared control: smoothly shifting control authority?" *Cognition, Technology & Work*, vol. 14, no. 1, pp. 19–28, Mar. 2012. [Online]. Available: https://doi.org/10.1007/s10111-011-0192-5

[12] N. Yu, K. Wang, Y. Li, C. Xu, and J. Liu, "A haptic shared control algorithm for flexible human assistance to semi-autonomous robots," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2015, pp. 5241–5246. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7354116

[13] M. Mulder, D. A. Abbink, and E. R. Boer, "Sharing Control With Haptics: Seamless Driver Support From Manual to Automatic Control," *Human Factors*, vol. 54, no. 5, pp. 786–798, Oct. 2012. [Online]. Available: https://doi.org/10.1177/0018720812443984

[14] A. Benloucif, A.-T. Nguyen, C. Sentouh, and J.-C. Popieul, "Cooperative Trajectory Planning for Haptic Shared Control Between Driver and Automation in Highway Driving," *IEEE Transactions on Industrial Electronics*, vol. 66, no. 12, pp. 9846–9857, Dec. 2019. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8624569

[15] N. Mol, J. M. Prendergast, D. A. Abbink, and L. Peternel, "Fitts' List Revisited: An Empirical Study on Function Allocation in a Two-Agent Physical Human-Robot Collaborative Position/Force Task," *IEEE Robotics and Automation Letters*, vol. 11, no. 1, pp. 202–209, Jan. 2026. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/11248827

[16] F. J. Rodríguez-Sedano, M. Conde, F. Rodríguez-Lera, and J. Chaparro-Peláez, "Measuring the impact of haptic feedback in collaborative robotic scenarios," in *Universal Access in the Information Society, vol. 23, no. 3*, Aug. 2024, pp. 1031–1049. [Online]. Available: https://link.springer.com/article/10.1007/s10209-023-01040-8

[17] "Github ejerez/mscthesis repository." [Online]. Available: https://github.com/ejerez/MScThesis

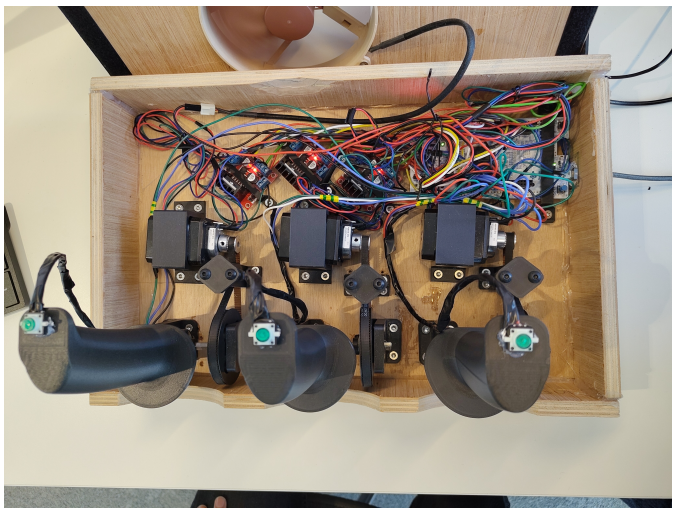
[18] "NASA Task Load Index (TLX) - NASA." [Online]. Available: https://www.nasa.gov/human-systems-integration-division/nasa-task-load-index-tlx/

[19] J. C. Huegel and M. K. O'Malley, "Progressive haptic and visual guidance for training in a virtual dynamic task," in *2010 IEEE Haptics Symposium*, Mar. 2010, pp. 343–350. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/5444632

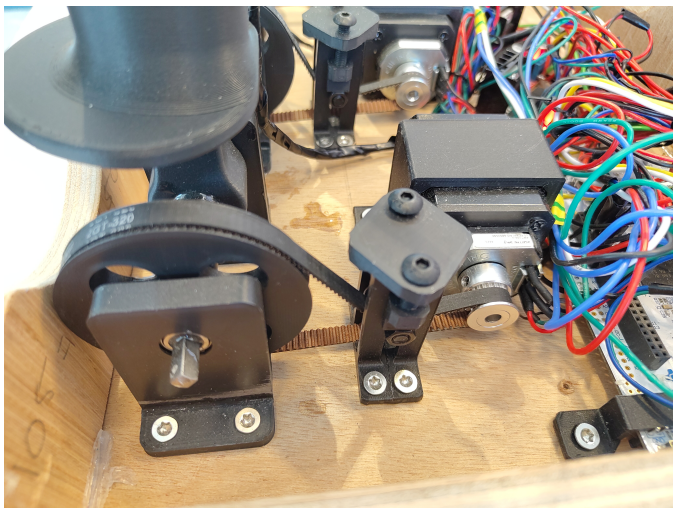
APPENDIX A
HAPTIC DEVICE

TABLE VI: Haptic-device and host communication rates.

Process	Nominal rate/interval	Location	Qualification
Electrical commutation	As fast as possible	Firmware	Variable; measured by debug counter, not fixed by a timer.
Motion and feedback control	1 kHz	Firmware	Scheduled every 1000 μs; no catch-up after an overrun.
Position telemetry	50 Hz	Firmware to host	Raw encoder angles plus applied normalized targets and voltage limits.
Shared-target command	30 Hz	Host to firmware	Nominal running worker cadence.
Keepalive	2 Hz	Host to firmware	Firmware feedback watchdog is 1.5 s.
Serial baud rate	115200 bit s^{-1}	USB CDC	Command and telemetry link.



(a) Electronics exposed. All wiring had to be custom made by hand, using silicone wires and DuPont-style connectors. The device is fed by a $12V$ power source which also powers the cooling fan.



(b) The joysticks are driven by NEMA13 stepper motors connected via a belt with a tensioner. All mechanical parts, including the joysticks themselves, the joystick gear, belt tensioner, etc. were designed and made in-house. The gear ratio is 1 to 4.

Fig. 10: A look into the Haptic Device

TABLE VII: Default haptic attractor regions in normalized joystick coordinates.

Axis	Region	Attractor target $x^\star$	Voltage ceiling $V_{\max}(x)$
All	$a \le 0.05$	0	6.0 V
All	$0.05 < a \le 0.125$	$0.10\,\mathrm{sgn}(x)$	$6.0E(0.125 - a; 0.025)$ V
Slew	$a > 0.125$	$0.10\,\mathrm{sgn}(x)$	$E(a - 0.125; 0.025)\left[3.6 + 1.9\,\mathrm{clip}\left(\frac{a-0.150}{0.850}; 0, 1\right)\right]$ V
Boom, hoist	$0.125 < a \le 0.675$	None before velocity drag	0 V before velocity drag
Boom, hoist	$a > 0.675$	$0.70\,\mathrm{sgn}(x)$	$E(a - 0.675; 0.025)\left[3.6 + 1.9\,\mathrm{clip}\left(\frac{a-0.700}{0.300}; 0, 1\right)\right]$ V

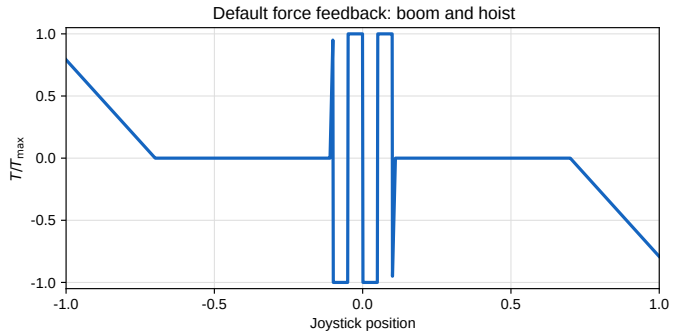


Fig. 11: Boom and hoist default haptic feedback showing torque as a function of joystick position.

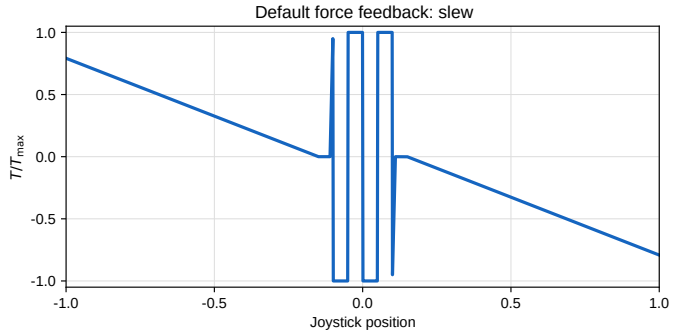


Fig. 12: Slew default haptic feedback showing torque as a function of joystick position.

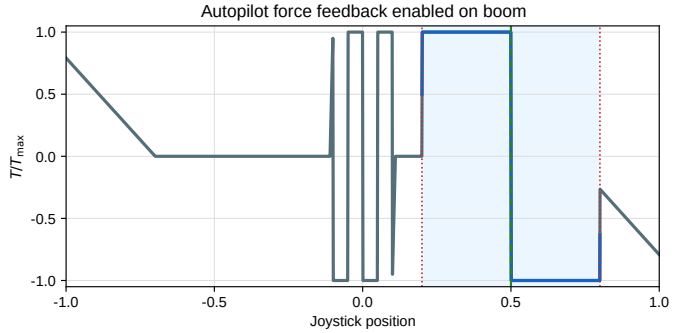
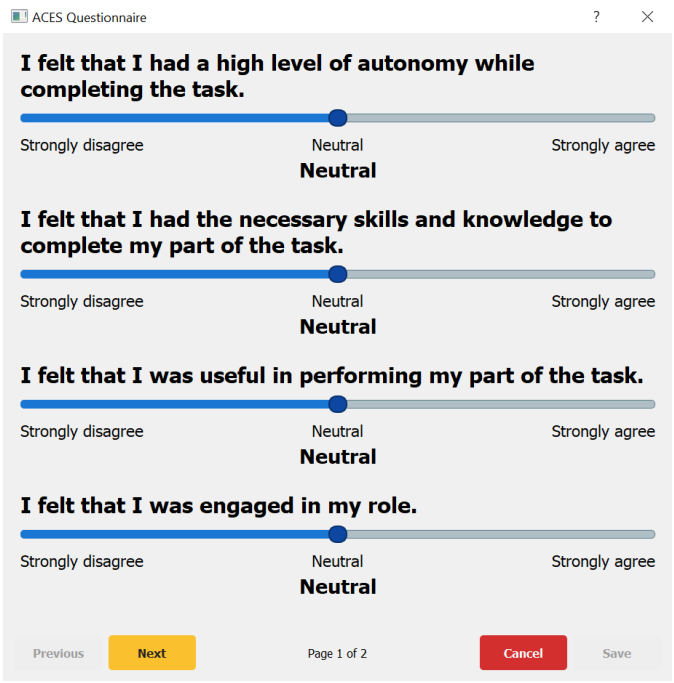


Fig. 13: Autopilot haptic feedback showing torque as a function of joystick position. The blue colored area is the range within the autopilot notch with the autopilot target at 0.5. Bringing the joystick out of the autopilot notch results in autopilot disconnect and the autopilot notch disappearing.

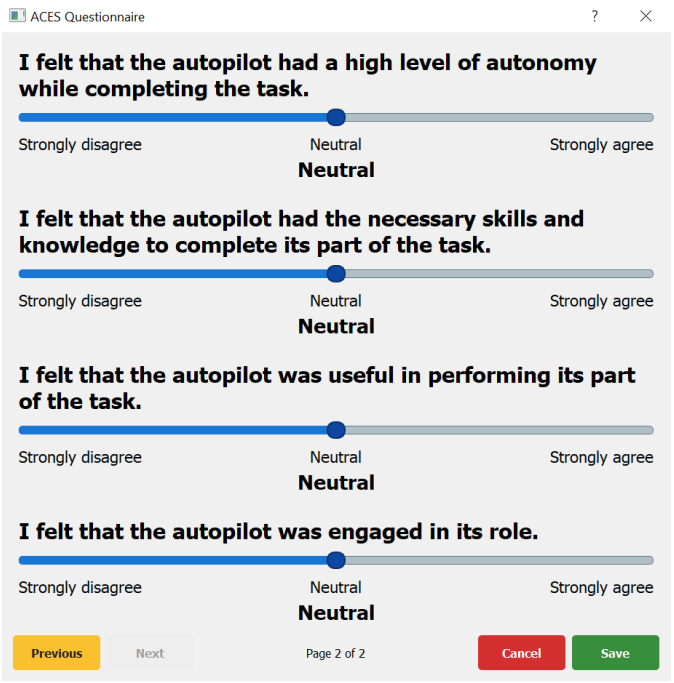
APPENDIX B
MEASURED METRICS

TABLE VIII: Measured Signals and Sampling Rates

Measure	Signal or calculation	Source rate	Stored rate
Completion time	Difference between trial stop and start timestamps	End of trial	Once per trial
Collision count	Number of distinct new collision messages since trial start	30 Hz polling	5 Hz
Measured payload position	Horizontal and vertical simulated payload position	Up to 8 Hz	5 Hz
Crane top-piece position	Simulated top-piece centre-of-gravity position	Up to 8 Hz	5 Hz
Corrected payload position	Forward-kinematic position + I term towards mean measured payload position in last $2T$	30 Hz	5 Hz
Selected drive target	Autopilot final target	Once per target set	5 Hz
Crane joint positions	Slew, boom, and hoist signals	30 Hz	5 Hz
Autopilot state and crane commands	Three active-axis flags and three normalized applied commands	30 Hz	5 Hz
ACES	Eight human/autopilot autonomy, competence, usefulness, and engagement items	Once per condition	Once per condition
NASA–TLX	Six unweighted workload items	Once per condition	Once per condition

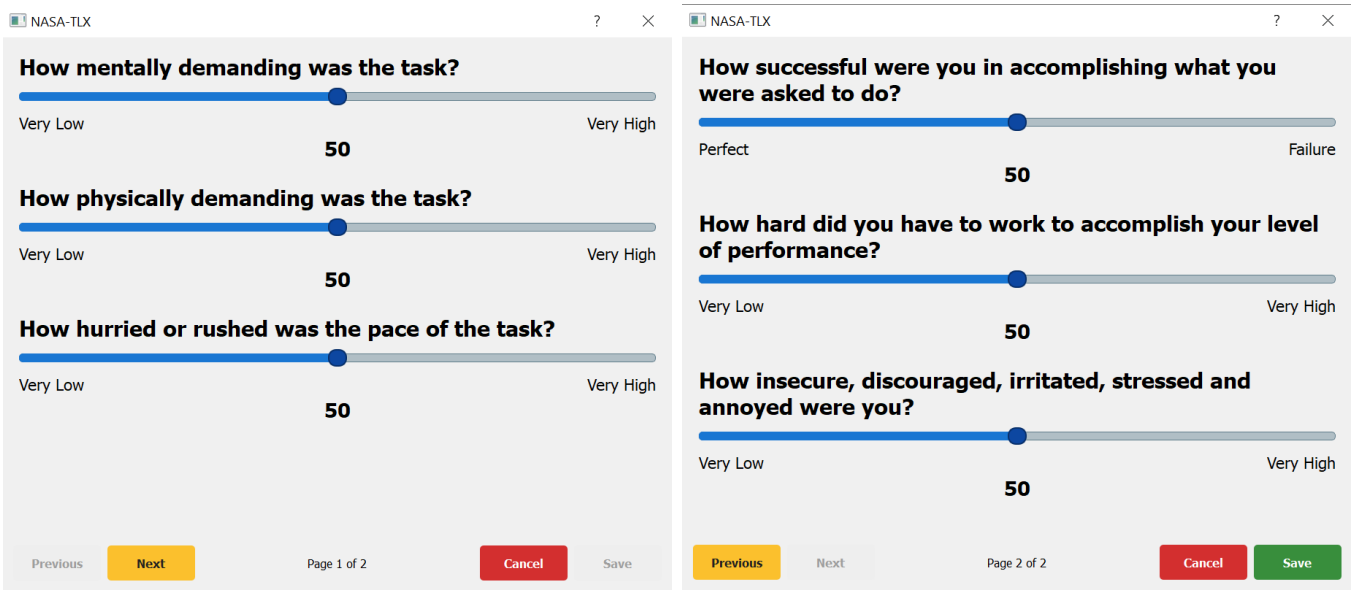


(a) Human perspective ACES questions.



(b) Autopilot perspective ACES questions.

Fig. 14: ACES subjective opinion questionnaire as answered by the participants.



(a) First page of NASA-TLX.

(b) Second page of NASA-TLX

Fig. 15: NASA-TLX perceived workload questionnaire as answered by the participants.