

Active Telepresence Assistance for Supervisory Control: A User Study with a Multi-Camera Tele-Nursing Robot

Alexandra Valiton, Hannah Baez, Naomi Harrison, Justine Roy, and Zhi Li¹

Abstract—Supervisory control of a humanoid robot in a manipulation task requires coordination of remote perception with robot action, which becomes more demanding with multiple moving cameras available for task supervision. We explore the use of autonomous camera control and selection to reduce operator workload and improve task performance in a supervisory control task. We design a novel approach to autonomous camera selection and control, and evaluate the approach in a user study which revealed that autonomous camera control does improve task performance and operator experience, but autonomous camera selection requires further investigation to benefit the operator’s confidence and maintain trust in the robot autonomy.

I. INTRODUCTION

Nursing is a critical role in society, including assisting the elderly and supporting patient recovery after disease or surgery. Globally, demand is soon expected to exceed the supply of healthcare professionals. Nursing robots present an opportunity to supplement the nurse workforce for both quarantine and routine patient care [1]. The teaming of human and (semi-) autonomous robots is a practical solution until autonomous robots can reliably perform a wide variety of tasks in cluttered human environments [2]. While direct teleoperation of high-DOF robot systems is complex and challenging [3], supervisory control of autonomous robots [4], [5] improves task performance [6], reduces workload [7], and allows a single human operator to supervise tasks among multiple robots [1], [8]. In order to supervise robot autonomy that is unreliable due to perception limitations or task uncertainty, appropriate active telepresence control is critical to maintaining operator situational awareness. For instance, an operator may need to select and control appropriate telepresence cameras to observe objects being manipulated from different viewpoints. Our prior study found that, with multiple cameras available, camera selection and manual camera control imposed a significant cognitive workload, particularly for novice operators [9].

Remote perception is a significant and challenging component of robot teleoperation [10]. Specifically, the limited field of view, unfamiliar frame of reference, multiple camera perspectives, and lack of depth perception make it difficult to perceive and understand the remote environment. Another challenging aspect of remote perception is dynamic coordination of perception and action [11]. Contemporary designs that support remote perception have mostly focused on autonomous control of a single telepresence camera, which

may avoid visual occlusion by the robot itself [12], the environment [13], or a third party [14].

We propose a novel approach to autonomous camera selection and control to provide remote perception support to robots with multiple moving cameras. We conducted a user study ($N = 14$) in which a humanoid nursing robot autonomously completed a sorting task under supervisory control. The two telepresence cameras available for the participants were a fixed camera on the robot head and an actuated camera attached to the robot hand not performing manipulation. Different levels of robot autonomy were provided to the participants, including manual camera selection and control (Mode 1), manual camera selection with autonomous camera control (Mode 2), and autonomous camera selection and control (Mode 3). With errors randomly introduced to the robot action sequence and motions, we evaluated the participants’ task performance in terms of accuracy and confidence in error detection. We also compared the camera usage, workload (through NASA-TLX and secondary task performance), and trust in the robot autonomy for remote perception support.



Fig. 1: Robot performs workspace organizing task under supervisory control. Participants are supposed to detect the errors in robot motion and actions.

Our study found that participants prefer to use the combination of a “head camera” with a large field of view and a “hand camera” with a large and dexterous motion range. For supervisory control, participants prefer to use the hand camera to detect errors in task precision and use the head camera to detect errors in large robot motions and action sequences. We found that autonomous camera control has significant impacts on task performance and workload, while autonomous camera selection significantly decreased the participants’ confidence in the accuracy of error detection, and participants’ trust in the autonomy for active telepresence. We saw that both the decision time for robot control and confidence in the accuracy of error detection decreased as the level of autonomy for active telepresence increased. This implies the level of autonomy for active telepresence needs to be carefully selected depending on whether the operator needs efficiency or accuracy in the task. We also found that the participants have varying camera selection preferences based on which part of the task they are performing, implying that it is necessary to personalize

¹Robotics Engineering Program, Worcester Polytechnic Institute, Worcester, MA 01609, USA {arvaliton, hbaez, ntharrison, jaroy, zli11}@wpi.edu

camera selection autonomy.

II. RELATED WORK

a) Supervising Robots using Multi-Camera Telepresence: Supervisory control relieves the teleoperator of the tedious and difficult control of robot motions and allows them to focus on high-level decision-making for the task. For patient care, supervisory control permits the nurse to focus on professional decisions and emotional care [8] instead of being distracted by robot teleoperation. The challenges of a supervisory control system arise when the operator trusts the robot system too much, leading to decreased situational awareness and increased operator error rates [7], [15]. The use of multiple active telepresence cameras may improve the teleoperators' situational awareness [10], yet supervisory robot control with a multi-camera robot platform is not a trivial task [16]. This is because switching between multiple camera views can trigger a saliency effect, in which operators may not be able to properly integrate information from multiple sources [8]. In addition, task performance will be lowered if the robot supervisors have to divide their attention [8], [16], such as if the participants have to manually select and control the telepresence cameras while focusing on task supervision.

b) Robot Autonomy for Assisting Active Telepresence: Robot autonomy for active telepresence can provide effective cognitive assistance to the operator, because cognitive workload can be decreased by reducing task demand [17]. Assistance for active telepresence can also improve participants' situational awareness, such that they will have better perception of features in the task workspace, comprehension of their meaning, and projection of future task states [18]. Thus far, assistance for active telepresence has mostly focused on autonomous control of a single telepresence camera. For structured tasks, visual servoing is used to keep the task goal in view [19]. For freeform tasks, optimal camera viewpoints can be autonomously determined using a well-engineered optimization function [12], [13], [20]. Prior research has shown that teleoperators benefit from both an egocentric and exocentric view of a remote workspace [8], which motivates the design of multi-camera telepresence system. The cognitive assistance for such systems has been limited to fusing information from multiple sensors (e.g., two camera views [21], point cloud and state estimation [22]) into one display [15]. Autonomous camera control and selection has been integrated to better utilize the complementary functions of various active telepresence cameras.

III. PROPOSED DESIGN OF TELEPRESENCE ASSISTANCE

The proposed active telepresence assistance is based on our prior human movement study on perception-motion coupling in active telepresence [23]. The prior study analyzed how people selected and controlled wearable cameras attached to their head, clavicle, dominant and non-dominant hands to perform different actions in a cup-stacking task.

From the prior study, we found that: 1) participants mostly prefer a combination of a "head camera" with large field of view, and a "hand camera" with large and dexterous

motion range; 2) participants prefer to use the head or clavicle cameras to observe reaching to or moving objects, and use the camera attached to one hand to observe the precise manipulation performed using the other hand; and 3) participants preferred not to view the wrist camera video feed while it moved. These findings inform our design of robot autonomy for camera selection and control for a humanoid nursing robot equipped with head and hand cameras.

a) Robot Autonomy for Camera selection: Overall, our proposed robot autonomy uses the robot *head camera* (C_H) to observe large robot motions and spatial relationships that need a global workspace view and uses the robot *right hand camera* (C_R) to observe precise motions and local task features (e.g., alignment of gripper and object). Consider a task of supervising a nursing robot to clean and organize a counter workspace: the robot needs to pick and place discarded objects into a bin and stack other objects at a designated location in a pre-defined order. For such a task, low-level errors (shown in Fig. 1) include misalignments between the gripper and the objects (*grasp errors*), and misalignments between the object in the robot's grasp and the location to stack or place (*placement errors*). The high-level errors include manipulating the objects in the wrong order (*order errors*) or sorting an object incorrectly (*action errors*). Beyond the example task above, the types of errors we consider are generally applicable to manipulation tasks in many domains. Although the detailed list of high-level and low-level tasks may be different depending on the application, in general our proposed camera selection uses head and hand cameras to detect high-level and low-level errors, respectively.




Fig. 2: Objectives for the autonomous camera control.

b) Robot Autonomy for Camera Control: Our proposed autonomous camera control adopts an autonomous dynamic camera control framework [12], [13], [24] which selects the optimal robot arm configuration based on a weighted sum of objectives (minimizing joint and end-effector velocity, acceleration, and jerk; avoiding joint limits; avoiding self-collision). Within this optimization framework, we introduced additional camera positioning objectives based on our prior study on the vision-motion coupling in active telepresence [23]. These additional objectives determine the dynamic camera motions: 1) **Hand-Tracking:** Minimizing the distance between the manipulator (left) hand and the camera z-axis, which intends to center the manipulator in the camera view; 2) **Camera Roll:** minimizing the dot product between the camera frame x-axis and the world frame z-axis, which intends to align the camera view with gravity-direction; 3) **Camera Height:** maintaining the hand camera position to be close to the head camera position in

the gravity-direction; 4) **Viewing distance**: maintaining the distance from the hand camera to the manipulation hand to be close to the viewing distance empirically determined in our pilot study; and 5) **Side View**: minimizing the dot product between the camera frame z-axis and the world frame y-axis, such that the hand camera tends to view the robot’s sagittal plane. Fig. 2 illustrates these camera control objectives. The relative weights of these objectives were selected by empirical testing in the current implementation, yet the weights can also be learned from expert demonstrations using inverse reinforcement learning [25].

IV. EXPERIMENT

a) Robot Platform: Fig. 3 shows the robot platform and its operator console. The robot platform consists of a humanoid torso (Rethink Baxter) with a pair of two-finger soft grippers (UBIROS GentleDuo). The available telepresence cameras include: one 180° fisheye head camera, (ELP-USBFD01M) fixed on the robot torso, and one Intel RealSense D435 depth camera attached to the robot hand not used for manipulation. The operator console enables the participants to manually control the right hand camera viewpoint using an HTC Vive controller and switch between head and hand camera views using the trigger on the Vive controller. It also allows the participants to switch camera views by pressing the ‘S’ key on the keyboard.




Fig. 3: Robot platform (left) and operator console (right).

b) Task: The participants were asked to supervise the robot to perform a task similar to that described in Section III. Specifically, the robot needs to pick up five cups in order from A to E. The cups to stack and to place into the bin were randomly decided for each trial. For each trial, we intentionally introduced a random number of low-level and high-level errors (ranging in total from 11 to 15) to the robot actions. Low-level errors include *grasp errors* that misaligned the gripper with respect to a cup when picking up and *stack* or *bin errors* that misaligned the cup in hand with the stack or the bin, respectively. High-level errors include *order errors* when the robot picks up a cup out of order, and *action errors* when the robot places a cup to be stacked into the bin, or vice versa. Before the robot executes any pick, stack, or place action, the participant needs to either confirm the robot action by answering “yes” to the *action confirmation question* displayed in the operator console terminal or reject the robot action by answering “no” if they notice any errors that may lead to failure of the robot action. The participant also needs to indicate the level of confidence in their judgement using a 1-5 Likert scale, with 1 for “I have no idea if this position is correct or not” and 5 for “I am completely sure that

this position is correct or incorrect”. Note that the robot will correct the errors when executing the actions, regardless of the participant’s response. This is because the performance of the supervisory task is evaluated by whether the participants can accurately detect the errors with a high level of confidence.

c) Experiment Procedure: We recruited $N = 14$ healthy participants (7 females and 7 males, age = 20.7 ± 1.5) who are mostly engineering students. The participants took a verbal math test (24 questions, single- and double-digit addition and subtraction) before the experiment. We recorded the total time for them to answer all questions correctly to provide a baseline of their performance for the secondary task. Participants were then introduced to the task and watched a demonstration of the robot performing the task to be supervised. Participants were encouraged to switch the camera view between the cameras with the ‘S’ key during the demonstration to familiarize themselves with the interface and the camera control autonomy. After the demonstration and practice, the participants were asked to report their current level of trust in the autonomy for autonomous camera control on a 1-5 Likert scale. The participants were then taught to use the Vive interface to control the robot’s right hand and switch cameras using the Vive controller. Participants moved to the evaluation phase if they were able to demonstrate their understanding of the Vive interface controls by aiming the camera at specific points in the workspace and identifying the letters on the cups.

During the evaluation phase, the participants completed three trials of supervisory control tasks with different levels of active telepresence assistance: **Mode 1:** Manual camera selection and control, **Mode 2:** Manual camera selection and autonomous camera control, and **Mode 3:** Autonomous camera selection and control. The modes of active telepresence assistance were randomized for each participant. Before each trial, the experimenter reviewed the task instructions and camera controls for the selected mode, informed the participant which cups were assigned to the bin, and reminded the participant that the cups should be handled in alphabetical order. Along with the robot supervisory task, the participants also needed to continuously answer simple math questions verbally in addition to the action confirmation questions in the supervisory control task. The participants need to answer each math question within 10 sec, or the question was considered skipped. After each trial, the participant completed a NASA-TLX survey and customized questionnaire to report more details about their task performance and preference of the design features of the selected telepresence assistance mode. The participants also completed a customized questionnaire to compare the three modes and to indicate their preference and reasons.

d) Evaluation Metrics and Data Collection: Our evaluation considers the performance of both the primary and secondary tasks. For the supervisory (primary) task, we use the decision time to either confirm or reject as well as error detection accuracy as objective performance metrics. To compute these metrics, we recorded the time stamps for every camera switch, the participant’s decision time to confirm or




Fig. 4: Task performance compared between telepresence assistance **modes** and participant groups of different **genders** and **video Game Experience**.

reject each robot action, and the active camera view when they responded to each action confirmation question. We also recorded the robot action sequence, the participant’s answers (i.e., their judgement of the actions), and their self-assessed confidence level (1-5 Likert scale) in their answers. For the math (secondary) task, we counted the number of math questions asked, answered, and answered correctly. Our collected data also included the participant’s responses to the NASA-TLX and customized questionnaires.

V. RESULTS

a) Task Performance: Our evaluation shows that the proposed telepresence assistance 1) can improve supervisory task performance by reducing decision time for error judgement, and 2) does not have a significant impact on the task accuracy, i.e., the correctness of error detection. The ANOVA analysis results in Fig. 4 (left) show that participants generally have better performance in Modes 2 and 3 than in Mode 1, though there is no significant difference between performance of Modes 2 and 3. We also compared task performance between participant groups with higher and lower video game experience. We found that *the proposed telepresence assistance has significant effects on mitigating the impacts of gender and video game experience factors*. Shown in Fig. 4 (middle), participants who considered themselves to be “Proficient” or “Experienced” in video games (labeled “Hi”) have significantly shorter decision times than those who considered themselves to have less game experience (labeled as “Lo”). However, a significant difference does not exist when these participants are provided with telepresence assistance in both Modes 2 and 3. A similar significant difference in performance exists between male and female participants without telepresence assistance (Mode 1), but doesn’t exist between the groups in Modes 2 and 3. Among our participants, video game experience was significantly correlated with gender ($R = .85$).

b) Camera Usage: We analyzed which camera participants used for observing large robot motions and for identifying the different kinds of errors. A diverse camera usage across participants implies that it is necessary to personalize the autonomy for camera selection. Our analysis assessed camera usage in Modes 1 and 2, in which manual camera selection is available. We focused on which camera is selected when the robot is moving the cup large distances or when participants answer action confirmation questions with an error present (see error type definitions in Section IV).

We determined camera usage in two ways: the active camera at answer time (shown in Fig. 7 (right)) and camera selection behavior. We calculated camera usage behavior by

computing the proportion of time (decision time or periods between questions when TRINA is moving large distances) using the *head camera* (C_H). We classified the camera usage into different ‘Types’. Type 1 usage means the participant uses C_H very little ($\leq 20\%$ of the time), while Type 5 usage means the participant uses C_H a lot ($> 80\%$ of the time). The thresholds for Types 2, 3 and 4 are evenly distributed between 20% and 80%.




Fig. 5: C_H usage types for different Manipulation Actions

Fig. 5 compares the distributions of camera usage Types in Modes 1 and 2. For each kind of error, as well as periods where the cup is being moved large distances, it shows the mean and standard deviations of the participants’ likelihood to employ each camera usage Type. We noticed that for some low-level (grasp, stack) errors, participants dominantly prefer to use the C_R (Type 1 usage), which is to some extent consistent with the findings in our prior work [23]. However, we also noticed diverse camera usage for observing TRINA moving a cup or identifying high-level (order or action) errors.

During the experiment, we noticed that some participants needed to switch back and forth between the cameras several times to detect some of the errors. A traditional solution to such behavior is to provide multiple telepresence camera views side-by-side [9]. However, multiple camera views run the risk of increasing the participant’s cognitive workload and distraction [15]. Our observations from the user study lead to the hypothesis that, if an operator is using a system with multiple active telepresence cameras and autonomous camera switching, it may be helpful to allow manual camera selection override capability.

c) Results from Survey: We used the Wilcoxon rank-sum test to analyze the survey responses. Across the three telepresence assistance Modes, we compared the participants’ perception of the task, trust in the autonomy for telepresence assistance, and confidence in the accuracy of their error detection. We found that: 1) Mode 3 has the most significant reduction in the participants’ *trust* in autonomy before and after using this mode; 2) *Confidence* in accuracy of error detection was significantly lower in Mode 3 than in Modes 1 and 2; 3) *Mental demand, physical demand, and effort* are

significantly higher in Mode 1 than in Modes 2 or 3; 4) There were no significant differences between Modes in terms of time demand, task success, or frustration. Note that we use $p < 0.05$ as the threshold for significance.




Fig. 6: Results from NASA-TLX survey and customized questionnaires.

Our survey results included participants' preference for telepresence assistance mode, choice of cameras for detecting different kinds of errors, and the reasons for their preference. Fig. 7 (left) shows the preferred cameras for observing each kind of action/error. Mode 2 was ranked most preferred by the majority of participants ($N = 10$). We found that the camera selection preference does not correspond to the demonstrated camera usage in Fig. 5. We hypothesize that improved autonomous camera selection may help operators match their camera selection to their preference.




Fig. 7: Participant preference and behavior in camera selection.

VI. CONCLUSIONS AND FUTURE WORK

We proposed a novel robot autonomy design to provide cognitive assistance to active telepresence based on our prior human movement study of vision-motion coupling in active telepresence [23]. Our user study found that our design for the autonomous camera control is generally preferred, though the design of the autonomous camera selection needs to be improved in future work, according to the diverse camera selection behaviors revealed in this study.

REFERENCES

- [1] H. F. Machiel Van Der Loos, D. J. Reinkensmeyer, and E. Guglielmelli, "Rehabilitation and Health Care Robotics," in *Springer Handbook of Robotics*, 2nd ed. Springer, Cham, 2016, ch. 64, pp. 1685–1728.
- [2] G. Adamides, G. Christou, C. Katsanos, M. Xenos, and T. Hadzilacos, "Usability guidelines for the design of robot teleoperation: A taxonomy," *IEEE Transactions on Human-Machine Systems*, vol. 45, no. 2, pp. 256–262, apr 2015.
- [3] M. A. Goodrich, J. W. Crandall, and E. Barakova, "Teleoperation and Beyond for Assistive Humanoid Robots," in *Reviews of Human Factors and Ergonomics*, 2013, vol. 9, ch. 5, pp. 175–226.
- [4] J. M. Beer, A. D. Fisk, and W. A. Rogers, "Toward a Framework for Levels of Robot Autonomy in Human-Robot Interaction," *Journal of Human-Robot Interaction*, vol. 3, no. 2, p. 74, jun 2014.
- [5] G. Niemeyer, C. Preusche, S. Stramigioli, and D. Lee, "Telerobotics," in *Springer Handbook of Robotics*, 2nd ed., B. Siciliano and O. Khatib, Eds. Springer, 2016, ch. 43, pp. 1085–1108.
- [6] J. Y. C. Chen, S. G. Lakhmani, K. Stowers, A. R. Selkowitz, J. L. Wright, and M. Barnes, "Situation awareness-based agent transparency and human-autonomy teaming effectiveness," *Theoretical Issues in Ergonomics Science*, vol. 19, no. 3, pp. 259–282, 2018.
- [7] A. E. Leeper, K. Hsiao, M. Ciocarlie, L. Takayama, and D. Gossow, "Strategies for human-in-the-loop robotic grasping," *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction - HRI '12*, p. 1, 2012.
- [8] J. Y. C. Chen and M. J. Barnes, "Supervisory Control of Multiple Robots," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 54, no. 2, pp. 157–174, 2012.
- [9] Z. Li, P. Moran, Q. Dong, R. J. Shaw, and K. Hauser, "Development of a tele-nursing mobile manipulator for remote care-giving in quarantine areas," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 3581–3586.
- [10] J. Chen, E. Haas, and M. Barnes, "Human performance issues and user interface design for teleoperated robots," *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, vol. 37, no. 6, pp. 1231–1245, 2007.
- [11] B. DeJong, J. Colgate, and M. Peshkin, "Mental transformations in human-robot interaction," in *Mixed Reality and Human-Robot Interaction*. Springer, 2011, pp. 35–51.
- [12] D. Rakita, B. Mutlu, and M. Gleicher, "An Autonomous Dynamic Camera Method for Effective Remote Teleoperation," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction - HRI '18*, 2018, pp. 325–333.
- [13] —, "Remote Telemanipulation with Adapting Viewpoints in Visually Complex Environments," in *Robotics: Science and Systems*, Freiburg im Breisgau, 2019.
- [14] S. Radmard and E. A. Croft, "Overcoming occlusions in semi-autonomous telepresence systems," in *2013 16th International Conference on Advanced Robotics (ICAR)*. IEEE, 2013, pp. 1–6.
- [15] C. W. Nielsen, M. A. Goodrich, and R. W. Ricks, "Ecological interfaces for improving mobile robot teleoperation," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 927–941, 2007.
- [16] S.-Y. Chien, Y.-L. Lin, P.-J. Lee, S. Han, M. Lewis, and K. Sycara, "Attention Allocation for Human Multi-Robot Control: Cognitive Analysis based on Behavior Data and Hidden States," *International Journal of Human Computer Studies*, vol. 117, pp. 30–44, 2018.
- [17] M. S. Prewett, R. C. Johnson, K. N. Saboe, L. R. Elliott, and M. D. Covert, "Managing workload in human-robot interaction: A review of empirical studies," *Computers in Human Behavior*, vol. 26, no. 5, pp. 840–856, 2010.
- [18] N. A. Stanton, P. M. Salmon, G. H. Walker, E. Salas, and P. A. Hancock, "State-of-science: situation awareness in individuals, teams and systems," *Ergonomics*, vol. 60, no. 4, pp. 449–466, apr 2017.
- [19] X. Sun, X. Zhu, P. Wang, and H. Chen, "A Review of Robot Control with Visual Servoing," in *2018 IEEE 8th Annual International Conference on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER)*. IEEE, jul 2018, pp. 116–121.
- [20] W. Bentz, S. Dhanjal, and D. Panagou, "Unsupervised learning of assistive camera views by an aerial co-robot in augmented reality multitasking environments," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3003–3009.
- [21] D. J. Rea, J. Wiebe, and J. E. Young, "Monocle : Interactive Detail-in-Context Using Two Pan- and-Tilt Cameras to Improve Teleoperation Effectiveness," in *IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, Lisbon, Portugal, 2017, pp. 962–967.
- [22] P. Marion, M. Fallon, R. Deits, A. Valenzuela, C. Pérez D'arpino, G. Izatt, L. Manuelli, M. Antone, H. Dai, T. Koolen, J. Carter, S. Kuindersma, and R. Tedrake, "Director: A User Interface Designed for Robot Operation With Shared Autonomy," Tech. Rep., 2015.
- [23] A. Valiton and Z. Li, "Perception-Action Coupling in Usage of Telepresence Cameras," in *IEEE International Conference of Robotics and Automation (ICRA)*, Paris, 2020, pp. 1–7.
- [24] D. Rakita, B. Mutlu, and M. Gleicher, "RelaxedIK: Real-time Synthesis of Accurate and Feasible Robot Arm Motion," in *Robotics: Science and Systems*, 2018, pp. 26–30.
- [25] M. Palan, N. C. Landolfi, G. Shevchuk, and D. Sadigh, "Learning reward functions by integrating human demonstrations and preferences," *arXiv preprint arXiv:1906.08928*, 2019.