Evaluating Human Understanding of a Mixed Reality Interface for Autonomous Robot-Based Change Detection

Christopher Reardon[1], Kerstin Haring[1], Jason M. Gregory[2], John G. Rogers[2]

***Abstract*— Online change detection performed by mobile robots has incredible potential to impact safety and security applications. While robots are superior to humans at detecting changes, humans are still better at interpreting this information and will be responsible for making critical decisions in these contexts. For these reasons, robot-to-human communication of change detection is a fundamental requirement for successful human-robot teams operating in such scenarios. In this work we seek to improve this communication, and present the results of a study that evaluates the interpretability of autonomous robot-based change detections conveyed via mixed reality to untrained human participants. Our results show that humans are able to identify changes and understand the visualizations employed without prior training. Our analysis of the limitations of this initial study should be constructive to future work in this domain.**

I. Introduction

In safety- and security-critical domains, important information can be conveyed by changes in the environment. For example, the presence of a new object in an environment could indicate a survivor's activity in a search and rescue mission, or adversarial activity in a military mission, depending on the object and scenario. A robot equipped with perception sensors and change detection capabilities can detect metric-based changes, but may not understand the saliency or relevancy of these detections with respect to the mission context and objectives. A human teammate, on the other hand, might accidentally overlook these changes, especially in visually complex environments or cognitively taxing situations, but is typically more proficient at reasoning about changes and taking the appropriate action.

A fundamental requirement of this human-robot team is effective communication of information between agents. To this end, an agent's performance is directly impacted by their ability to understand the information provided to them by another agent; a robot can enable – or hinder – human decision making in response to a detected change depending on how the information is presented [1]. Already, we have presented a mixed reality-based system that combines an autonomous robot and augmented reality head-mounted device (AR-HMD) for facilitating communication of robot-based change detection [2]; however, presenting these detections in a way that maximizes human understanding is a significant challenge. Fig. 1 shows an example illustration of point-based change detection; this paper seeks to address the problem of how a robot can convey detected changes to a human teammate.

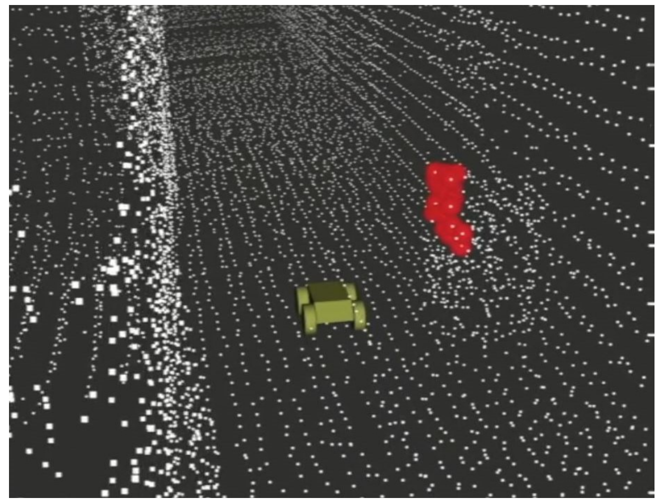

Fig. 1. A robot constructs a relatively dense point-cloud-based map of the environment, and detects changes in the point cloud (shown in red) from a previous reading. This paper seeks to address the problem of how an autonomous robot can convey point-based change detections for interpretation by a human teammate.

In this work, we strive to improve human-robot communication by empirically evaluating the comprehensibility of our change detection interface using data from our system operating in the real world. We present results from a user study consisting of humans viewing changes detected and indicated by an autonomous robot. The participants are not provided with any *a priori* or supporting information aside from the visualizations offered by the robot. The quality of communication achieved by our interface is then quantified by the human's scene understanding as measured by their ability to correctly identify changes in the environment.

II. Related Work

Generally speaking, novelty detection refers to the identification of differences between *test* data and *training* data or a *model* learned from that data [3]. Environmental change detection is an application of novelty detection where an initial model of the scene is built and used to compare against new "test" data in order to identify physical changes. An autonomous robot can perform change detection by navigating through an operational environment and analyzing the streaming data from its sensors with respect to the stored, *a priori* model. Autonomous robot-based change detection is useful in a range of domains including inspection [4],

[1] University of Denver, Denver, CO USA {christopher.reardon, kerstin.haring}@du.edu
[2] U.S. Army Research Laboratory, Adelphi, MD USA {jason.m.gregory1, john.g.rogers59}.civ@mail.mil

surveillance [5], [6], safety and security [7], and robust outdoor navigation [8]. Change detections from an autonomous robot can be made actionable by human teammates if they are communicated effectively, which has previously been performed using mixed reality systems [9], [10].

Augmented, virtual, and mixed reality technologies have the potential to address critical communication issues by mediating human-robot interactions [11], [12]. Augmented reality visualizations have been shown to improve robot communication, when combined with complex natural language and gestures, regardless of human cognitive load in mixed reality systems [13]. Augmented reality has also been shown to be effective at efficiently communicating robot motion intent in human-robot joint tasks [14].

In this work we seek to characterize the comprehensibility of change detection communicated by an autonomous robot to an AR-HMD. The study presented here quantifies the performance of our previously proposed interfaces, including when all detected changes are shown to the human [15] as well as aggregated detection data [2]. We believe the quantitative analysis offered in this work will support future development of system improvements for conveying robot-based change detection in a more human-understandable fashion.

III. Approach

Our approach to communication of change detection information begins with an autonomous robot detecting changes in an environment relative to a prior model. The robot also connects with a human wearing an augmented reality head-mounted device (AR-HMD), and aligns its map with the AR-HMD's 3D map representation. This allows the robot to pass information about changes detected to the human, which are then visualized through AR. We discuss the specific details of this implementation below.

A. Change Detection

The goal of our change detection capability is to automate the process of identifying and locating geometric-based anomalies in the environment, such as the addition of new objects, with respect to a given reference model. To do this, a robot first creates a 3D model of the environment that represents the known state of the world. The robot can then continuously compare observations using its onboard sensors during any subsequent traversals, which can ultimately provide enhanced situational awareness and alleviate the required workload for a human to monitor an environment. Here, we provide a brief overview of our approach to change detection and direct the reader to our previous works for additional details [15].

Our change detection algorithm builds the 3D *model* and *test* representations of the environment using OmniMapper [16], [17], a pose graph-based Simultaneous Localization and Mapping (SLAM) algorithm that computes a solution to the robot's trajectory from the robot's LiDAR and inertial measurement unit (IMU) sensors. We also compute the homogeneous transformation matrix between the robot and human using the Iterative Closest Point algorithm [18] to provide a common coordinate frame that is necessary for accurately displaying visual information, such as change detection markers in the human's AR-HMD. Given a *model* point cloud, the robot detects changes in a *test* cloud using a set of difference segmentation functions and outlier filters implemented in Point Cloud Library [19]. For every point in the *test* cloud, the nearest point and corresponding distance in the *model* cloud is computed by constructing KD-trees that reduce the quadratic search complexity to $n \log n$, which supports real-time operations. If the distance between two corresponding points is greater than a threshold δ, it is denoted as a change and accumulated in the *change* point cloud. After all points have been evaluated, the *change* cloud is filtered to reduce noise by excluding any points with less than a threshold quantity λ within a radius r.

One of the paramount challenges for change detection solutions is overcoming the various potential sources of error, e.g., inherent range error from the sensor, quantization error from the SLAM algorithm, and distance computation error from variance in viewpoints between clouds. Error plays an important role in human understanding because it not only decreases the performance of the change detection algorithm, but further complicates the communication of mission-relevant information to the human due to degraded or unexpected visualization. In the presence of error, the visualization of detected changes may not align precisely with objects in the scene or could appear in void space, which can become distracting or confusing for the human. While it is possible humans may be capable of mitigating the effect of imperfect change detection through proper training and field experience, future work will seek to improve the quantitative performance of our change detection algorithm.

B. Aggregation of Change Data for Visualization

In order for changes to be interpreted by a human user, they must be presented to the human. Augmented reality has been shown to offer a situated and salient means of accomplishing interpretable information exchange from robots to humans [11]. We have previously demonstrated using AR technology to convey raw change detection information from a robot to its human teammate [15]. However, an open challenge exists in that raw changes detections generated from real-world field environments by robots with high-resolution sensors may be difficult for a human to interpret. This could be because of the quantity of changes, the significance of for example small changes below a threshold of interest (e.g., leaves blowing on a tree), or data noise from sensors.

To address this challenge, we propose using data aggregation methods for the ultimate purpose of creating change visualizations that are more interpretable by human teammates without losing representational ability. We use an approach similar to what was previously presented as a concept [2] to cluster detected changes. In particular, we use Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [20]. We selected this method for aggregation after empirical comparison with other state-of-the-art clustering techniques, and note that it has several advantages that make it robust to real-world environments, including incorporation of noise and robustness to outliers, as well as handling a non-*a priori*-specified number of irregularly-shaped clusters.

After the change detections are clustered, we discard outliers and generate a visualization for each cluster. Because of the density of the point cloud sample corresponding to any changed object is relative to a number of factors, including the size of the object, the sensor resolution, and the distance and angle between the sensor-equipped robot and the object, the change detection points are only a small representation of the exposed surface of the changed object. For this reason, our visualization must approximate the location and size of the object without full knowledge of its shape or bounds. Therefore, for each change cluster we generate a spherical visualization, centered on the cluster centroid, with radius scaled by the intracluster distance of the change cluster.

The robot then publishes these changes in real-time to generate visualizations on the human-worn AR-HMD. Because of the online alignment of human and robot pose discussed in Sec. III-A, these change detections generated by the robot are able to be presented live in AR to the human teammate. In our experiments, we are then able to present both raw change detections and change detection clusters in this manner for comparison.

IV. Experiments

To provide an initial evaluation of the interpretability of change detection visualizations in augmented reality, a video-based study with crowd-sourced participants was conducted. The intent of this study was to determine how interpretable the change visualization methods proposed were to untrained human study participants. This will provide understanding for the representation and presentation of change information for future in-person studies and system design.

A. Hardware

Our change detection algorithm can be run on any robot that has an onboard LiDAR sensor, IMU sensor, and sufficient computing resources. In our experiments, we used a Clearpath Robotics Jackal mobile ground robot equipped with a Velodyne VLP-16 LiDAR, MicroStrain 3DM-GX4-25 IMU, and Intel Core i5-4570TE CPU. The robot was teleoperated by the experimenters around the environment, described in Section IV-B, to build the initial *model* and *test* clouds. Change detection was performed online by the robot and visualization was communiated to a human wearing a Microsoft Hololens AR-HMD via a Ubiquiti Bullet M5HP 5 GHz WiFi radio. The experimenters visualized the change detection markers in real-time and recorded videos of this visualization to present to participants in a user study.

B. Environment

Experiments were conducted in a large ($\sim 200m^2$) indoor laboratory setting (Fig. 2) to control for any changes that would be extraneous to the experiment. Three identically-sized $\sim 0.5m$ cubical cardboard boxes were used as the objects-of-interest that were subject to change. In each experimental scenario, zero, one, or two boxes were added to the scene. All boxes were placed in fixed, evenly spaced positions roughly equally distant from the human viewer. While no other objects were physically co-located with the boxes, the background and sides of the image had numerous visible features, such as cabinets, a door, and markings on the floor.

C. Experimental Procedure

Each experiment consisted of two phases: the collection of an *initial* point-cloud map by the robot, and an exploration phase where the robot collected a *change* model. In the *initial* phase, one, two, or – in the baseline case – three boxes were present. In between the phases, the experimenters added the missing boxes to their predetermined positions as described in Sec. IV-B. Then, during the *change* phase, the robot maneuvered through the setting autonomously creating a new 3D map of the environment and comparing that map against the map created during the *initial* phase.

Differences between the two maps were treated as changes, as described in the approach in Sec. III-A. For the change detection and noise-reduction thresholds described in Sec. III-A, we use $\delta = 10$ cm, $\lambda = 10$ and $r = 30$ cm. Through real-world experimentation and validation using multiple models of LiDAR sensors with varying resolutions and fields of view, we have empirically discovered that these parameter values work well for online, incremental change detection [15].

Changes were then processed and visualized via AR as described in Sec III-B. This mapping, change detection, and change visualization is performed online. Details of this online process are found in [15]. Briefly, we assume the robot and AR-HMD can communicate wirelelessly; in our system the Hololens connects directly to the robot via an ad-hoc network, as robust field-capable multi-agent networking is outside the scope of this paper. The Hololens' built-in 3D visual SLAM-based mesh map is converted onboard to a point-cloud representation and transmitted to the robot, where the change detection step is performed. Alignment of the robot and AR-HMD frames are also calculated and maintained by the autonomous roobt. This allows for appropriately-positioned visualizations to be shown to the human teammate via AR. Detected changes to converted to the different visualizations used in this experiment, and shared wirelessly back to the Hololens AR-HMD. The entire system operates online on the specified hardware in the experimental environment at approximately 1-3 frames per second, which is sufficient for this evaluation.

For our experiments, a human wears the AR-HMD during multiple change detection scenarios. Videos of the AR feed, including both real video and visual augmentations, were captured using the HoloLens built-in video capture capability. Three types of visualizations were generated for comparison in a user study: 1) visualization of the raw point-change data as fixed-size red spheres (Fig. 2a), 2) visualization of the change detection clusters as blue spheres

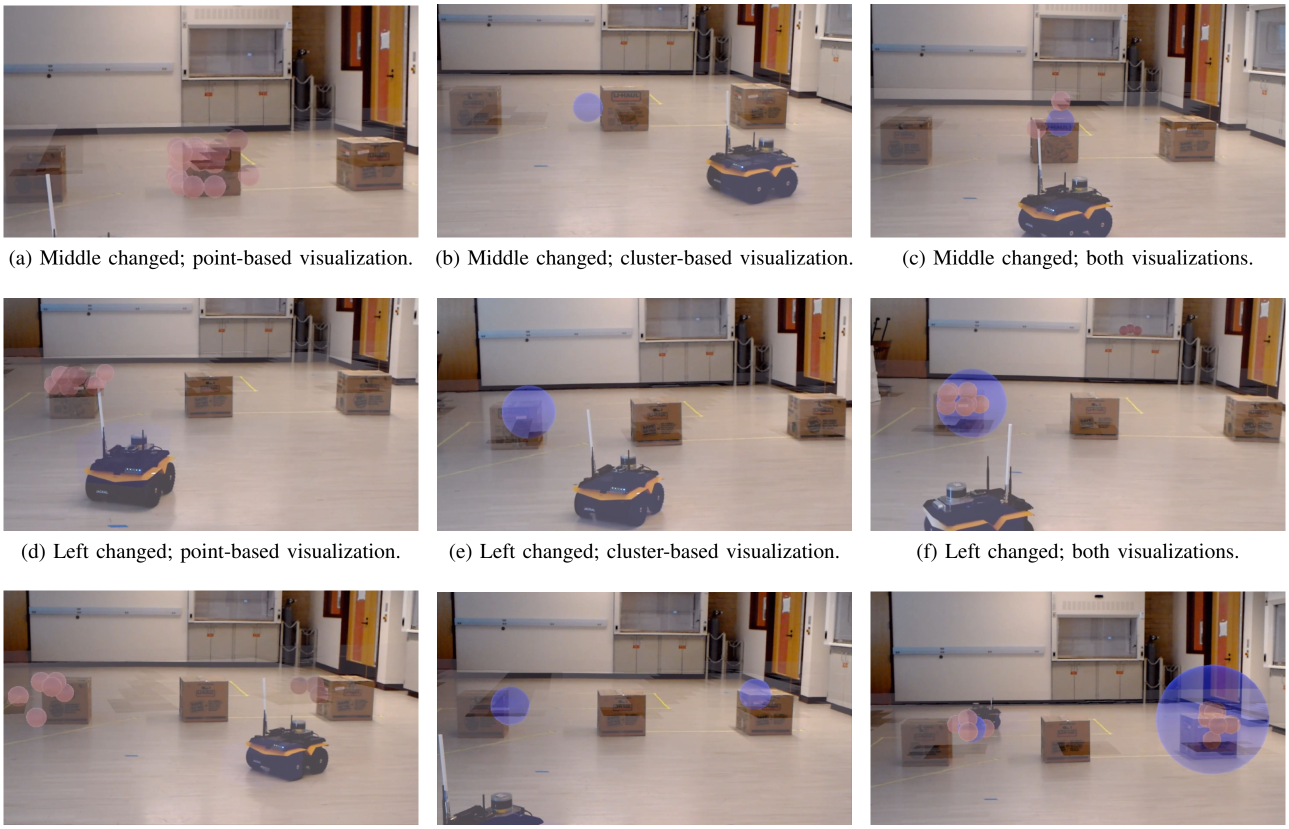

(a) Middle changed; point-based visualization. (b) Middle changed; cluster-based visualization. (c) Middle changed; both visualizations.

(d) Left changed; point-based visualization. (e) Left changed; cluster-based visualization. (f) Left changed; both visualizations.

(g) Left + right changed; point visualization. (h) Left + right changed; cluster visualization. (i) Left + right changed; both visualizations.

Fig. 2. Examples of the raw point visualization, aggregated cluster visualization, and combined visualization techniques used to indicate the robot's detected changes in a lab environment, in three different change configurations. In (a)-(c), the middle box changed from the initial model; the left and right boxes were present in the initial model. In (d)-(f), the left box changed. In (g)-(i), the left and right boxes were changes from the initial model.

of scaled size (Fig. 2b), and 3) visualization of both 1) and 2) together (Fig. 2c). Videos for each of these three types of visualizations, with all combinations of box locations, for zero (baseline), one, and two boxes changed were generated, for a total of 19 videos: 1 (no changes) + 3^2 (one change) + 3^2 (two changes).

D. User Testing

To test if humans can understand the changes in the environment as visualized, we conducted an evaluation with human subjects (approved by the University of Denver IRB 1770133-1). Due to restrictions in in-person experimental events imposed during the COVID-19 pandemic, an online study was necessary to perform this evaluation. In the online study 21 participants were shown 19 videos with either 0, 1, or 2 actual changes visualized (true positive) in either red (point-based), blue (cluster-based), or both spheres as shown in Figure 2 and several false positive objects highlighted by the algorithm. Two participants were excluded, one for not accepting to participate, and the other for completing the survey in 78 seconds. This left a total of 19 participants ($M_{age} = 36.1, SD_{age} = 7.7$, 12 Male, 7 Female). The participants' exposure to computer games was measured to gauge their experience as a very high or low exposure could influence how they interpret the visualization. On a scale from 1 ("never") to 5 ("daily"), none of the participants answered "never". The average for computer games was at $M_{games} = 3.7$ ($SD_{games} = 1.1$), which can be considered a slightly higher exposure than the Expected Value of 3. When divided into Action games ($M_{action} = 2.9, SD_{action} = 1.2$)), Strategy games ($M_{strategy} = 2.8, SD_{strategy} = 1.0$)), and Puzzle games ($M_{puzzle} = 2.7, SD_{puzzle} = 1.3$)), the average exposure is around the Expected Value and an analysis of variance shows no significant differences among or between the groups ($F(2, 54) = .17.p = .84$). This leaves the assumption that the visualization interpretation results are not significantly or systematically influenced by high or low computer game exposure. The videos were shown in random order, not altered and represented the algorithmic outcomes as described in III-A and as visualized in III-B to test the hypotheses on the visualization.

It was hypothesized that $\mathcal{H}1$: Humans can understand the changes detected without prior training and that $\mathcal{H}2$: understand changes with at least 90% accuracy. "Understanding" is defined as correct identification of a true positive using only the information provided by the robot. In the videos, participants only saw what the algorithm highlighted and did not have a-priori data about how the environment appeared before the change. Participants had no additional information that would otherwise facilitate decision making.

For each video, participants were asked to indicate which object was highlighted by the algorithm as a change. Then objects from the videos were listed and participants indicated on a rating scale (object "definitely did not change", "likely did not change", "likely changed", "definitely changed") how they perceived the visualization. With this scale, it could be identified how many objects participants correctly identified as well as the areas where they were unsure what true positive changes the visualization of the algorithm showed.

V. Results and Discussion

To evaluate $\mathcal{H}1$ we calculated the participants' hit rates, the proportions of trials in which the video stimuli was presented and the participants correctly responded. Out of the 152 possible answers to changes in the presented video, 67 (44%) have been answered absolutely correct at a 100% confidence level. This measure only includes when all participants unanimously answered that a stimuli is definitely (not) present. It does not include potential variations in the data when participants answered "likely (not)". To capture this data at the 90% confidence level, we calculated an average score for each question and out of the 152 possibilities, 117 (69%) changes were correctly identified with 90% confidence (Fig. 3).

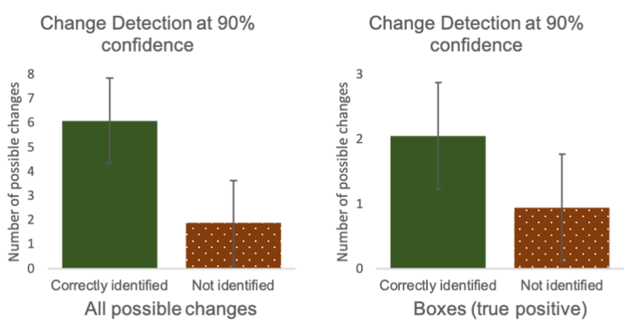

Fig. 3. The differences of correctly identified changes and not identified changes for all possibilities and looking at the possible true positives (Boxes) only. Error bars are standard errors.

Similarly, when we evaluated $\mathcal{H}2$ looking only at possible true positives on a 90% confidence level (i.e. only the boxes could be true positives), 68% were identified correctly. To identify how different those results for all 152 possibilities in the stimuli as well as for possible true positives (Boxes) only, a simple t-test compared the respectively correctly identified and not identified changes. For both, the results ($t(36) = 3.73, p = .0003$ and $t(36) = 7.37, p < .0001$) show that the correctly identified changes are significantly higher and not due to chance.

Combining the participants' interpretations into a binary score (i.e. combining "likely" and "definitely") and calculate correctly identified changes, the data shows that participants identify changes 75% correctly when the video stimuli shows one change, and only 49% correctly when two changes are visualized.

An analysis of variance was performed to identify potential differences between the three different stimuli types (point-, cluster-based, both) and did not show any significant differences between or within the three types.

These results confirm $\mathcal{H}1$ and $\mathcal{H}2$ in part only. Participants are able to identify changes and understand the visualization without prior training. The results also show that there is some confusion on how to distinguish false positive from true positive changes.

One of the limitations in this evaluation was that while the visual augmentations were positioned in three dimensions, because the videos were recorded with a fixed perspective (i.e., the AR-HMD wearer did not move), the visualizations may appear two-dimensional to the study participants. Minor errors in point cloud alignment discussed in Sec. III-A sometimes result in visualizations placed slightly to the side or above the detected change. Further, no training or *a priori* information was provided to the users, meaning that misinterpretations of the visualizations in the video stimuli are not corrected for in the obtained data. Also, as noted in Sec. IV-B, the environment contained numerous background objects which were included in the survey. Together, these issues may have exaggerated the likelihood of false positives (see Fig. 2c for example, where a red marker above the changed box visually intersects with the cabinets in the background). Some participants stated in their comments an unprompted interpretation of the visualization colors and the location of the points or clusters that was not mentioned at any point in the experimental description. This indicates that without any training on the current visualization, people might associate meaning to parts of the visualization that do not have any.

Overall, these results show a clear proof of concept that the visualization succeeds at highlighting environmental changes and can be understood by an untrained human user. For future studies, we are considering the implications and limitations to change the video stimuli for online studies [21] using moving perspective so that the visualization does not appear two-dimensional. We will also study the impact of better informed users by incorporating a short training video for the user study on how changes in the environment are visualized by the robot. A replication of this study in the future with in-person participants being able to perceive the visualization in three-dimensions is anticipated to introduce more clarity into the information the visualization displays. In future user study investigations we will introduce primary or secondary tasks in addition to the change detection interpretation to better understand the effects on situational awareness, performance, and mental workload [22], [23].

VI. Conclusion

In this work, we have presented an autonomous robot change detection and communication system and the results of a study that shows that humans are able to understand change detections communicated by a robot to a human via augmented reality. The results show no significant difference between the visualization strategies employed; however, we believe this could be induced by the fixed-perspective video collection method, combined with slight errors in point cloud alignment. In future work, we intend to address both the alignment issues as well as experiment with perspective-varying video collection. This work will also be highly informative to planned future in-person studies ultimately aimed at creating a more robust and interpretable change detection visualization system for enhanced human situational awareness.

References

[1] R. Parasuraman, M. Barnes, K. Cosenzo, and S. Mulgund, "Adaptive automation for human-robot teaming in future command and control systems," Army research lab aberdeen proving ground md human research and engineering . . . , Tech. Rep., 2007.

[2] C. Reardon, J. Gregory, C. Nieto-Granda, and J. G. Rogers III, "Designing a mixed reality interface for autonomous robot-based change detection," in *Virtual, Augmented, and Mixed Reality (XR) Technology for Multi-Domain Operations II*, vol. 11759. International Society for Optics and Photonics, 2021, p. 117590J.

[3] M. A. Pimentel, D. A. Clifton, L. Clifton, and L. Tarassenko, "A review of novelty detection," *Signal Processing*, vol. 99, pp. 215–249, 2014.

[4] S. Marsland, U. Nehmzow, and J. Shapiro, "On-line novelty detection for autonomous mobile robots," *Robotics and Autonomous Systems*, vol. 51, no. 2-3, pp. 191–206, 2005.

[5] P. Núñez, P. Drews, R. Rocha, M. Campos, and J. Dias, "Novelty detection and 3d shape retrieval based on gaussian mixture models for autonomous surveillance robotics," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2009, pp. 4724–4730.

[6] A. W. Vieira, P. L. Drews, and M. F. Campos, "Spatial density patterns for efficient change detection in 3d environment for autonomous surveillance robots," *IEEE Transactions on Automation Science and Engineering*, vol. 11, no. 3, pp. 766–774, 2014.

[7] M. Sturari, M. Paolanti, E. Frontoni, A. Mancini, and P. Zingaretti, "Robotic platform for deep change detection for rail safety and security," in *2017 European Conference on Mobile Robots (ECMR)*. IEEE, 2017, pp. 1–6.

[8] B. Sofman, B. Neuman, A. Stentz, and J. A. Bagnell, "Anytime online novelty and change detection for mobile robots," *Journal of Field Robotics*, vol. 28, no. 4, pp. 589–618, 2011.

[9] C. Reardon, K. Lee, J. G. Rogers, and J. Fink, "Augmented Reality for Human-Robot Teaming in Field Environments," in *International Conference on Human-Computer Interaction: Virtual, Augmented, and Mixed Reality*. Springer, 2019, pp. 79–92.

[10] J. M. Gregory, C. Reardon, K. Lee, G. White, K. Ng, and C. Sims, "Enabling intuitive human-robot teaming using augmented reality and gesture control," *arXiv preprint arXiv:1909.06415*, 2019.

[11] D. Szafir, "Mediating Human-Robot Interactions with Virtual, Augmented, and Mixed Reality," in *International Conference on Human-Computer Interaction*. Springer, 2019, pp. 124–149.

[12] K. S. Haring, V. Finomore, D. Muramoto, N. L. Tenhundfeld, J. Wen, and B. Tidball, "Analysis of using virtual reality (vr) for command and control applications of multi-robot systems," in *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[13] N. Tran, T. Grant, T. Phung, L. Hirshfield, C. Wickens, and T. Williams, "Get this!? mixed reality improves robot communication regardless of mental workload," in *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 412–416.

[14] E. Rosen, D. Whitney, E. Phillips, G. Chien, J. Tompkin, G. Konidaris, and S. Tellex, "Communicating robot arm motion intent through mixed reality head-mounted displays," in *Robotics research*. Springer, 2020, pp. 301–316.

[15] C. Reardon, J. Gregory, C. Nieto-Granda, and J. G. Rogers, "Enabling Situational Awareness via Augmented Reality of Autonomous Robot-Based Environmental Change Detection," in *International Conference on Human-Computer Interaction: Virtual, Augmented, and Mixed Reality*. Springer, 2020, pp. 611–628.

[16] A. J. Trevor, J. G. Rogers, and H. I. Christensen, "Omnimapper: A modular multimodal mapping framework," in *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2014, pp. 1983–1990.

[17] J. Gregory, J. Fink, E. Stump, J. Twigg, J. Rogers, D. Baran, N. Fung, and S. Young, "Application of multi-robot systems to disaster-relief scenarios with limited communication," in *Field and Service Robotics*. Springer, 2016, pp. 639–653.

[18] A. Segal, D. Haehnel, and S. Thrun, "Generalized-ICP," in *Robotics: Science and Systems*, vol. 2, no. 4, 2009, p. 435.

[19] R. B. Rusu and S. Cousins, "3D is here: Point cloud library PCL," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 1–4.

[20] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, *et al.*, "Scikit-learn: Machine learning in python," *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.

[21] D. Feil-Seifer, K. S. Haring, S. Rossi, A. R. Wagner, and T. Williams, "Where to next? the impact of covid-19 on human-robot interaction research," 2020.

[22] R. Parasuraman, K. A. Cosenzo, and E. De Visser, "Adaptive automation for human supervision of multiple uninhabited vehicles: Effects on change detection, situation awareness, and mental workload," *Military Psychology*, vol. 21, no. 2, pp. 270–297, 2009.

[23] R. van der Kleij, T. Hueting, and J. M. Schraagen, "Change detection support for supervisory controllers of highly automated systems: Effects on performance, mental workload, and recovery of situation awareness following interruptions," *International journal of industrial ergonomics*, vol. 66, pp. 75–84, 2018.