

Demonstrating Trustworthiness in Open-Loop Model Mediated Teleoperation for Collecting Lunar Regolith Simulant*

Joe Louca^{1,2}, Aliz Zemeny¹, Antonia Tzemanaki² and Romain Charles¹
joe.louca@bristol.ac.uk

Abstract—Teleoperated robotics will be an essential tool to support upcoming lunar exploration and in-situ resource utilisation activities. However, the communication delays between Earth and the Moon makes operating these robots extremely challenging. Model-Mediated Teleoperation (MMT) is a method of controlling these remote systems in perceived real-time, via a simulation, but is dependent on the accuracy of its model. In this work, a computationally efficient model of lunar regolith was implemented in an open-loop MMT system. The behaviour of the virtual model was compared with its physical equivalent during manipulation tasks. The model predicted the outcome of a regolith simulant scooping task with sufficient accuracy to be considered effective and trustworthy 100% and 92.5% of the time, respectively. Pouring actions were less accurate, but trustworthiness and effectiveness can still be ensured by restricting the orientation of the end effector whilst carrying simulant material. Simulated haptic interactions were representative of the real-world during simple, linear tasks (pressing and dragging), but not during more complex motions. This simulation could be adapted to account for reduced gravity, to form a delay-robust lunar MMT system, or to build operators' trust in the system by familiarising themselves in a low-risk virtual world.

I. INTRODUCTION

Alongside a boom in lunar lander missions this decade, several public and private organisations are now researching how best to extract valuable resources, such as oxygen and water, from readily available materials such as lunar regolith [1], [2]. Remote manipulation of lunar regolith will be an essential step in these activities, as regolith would first need to be collected from the Moon's surface [3]. This could be achieved by using a teleoperated robotic manipulator.

To interact with the task environment during teleoperation, the operator uses an interface (local subsystem) which sends commands to a remote subsystem. The significant distances between these subsystems pose a major challenge in Earth-Moon teleoperation, as they result in noticeable delays, e.g., 2.6 sec round-trip delay time if travelling at the speed of light [4] or much longer, practically, due to relays and data processing. The negative impacts of signal delays on performance are widely reported, e.g., reduced accuracy [5], increased mission completion time [6], greater incidence of task errors [7], and reduced trust in the system [8].

*This work was funded by the European Space Agency, EPSRC (grant no. EP/V519650/1) and Thales-Alenia Space (UK).

¹European Centre for Space Applications and Telecommunications, European Space Agency, Harwell, Didcot, UK.

²School of Engineering Mathematics and Technology, University of Bristol, UK and the Bristol Robotics Laboratory, UK.




Fig. 1. Teleoperation system used in these experiments. Camera IDs, positions and orientations are indicated by green circles and arrows. Physical and virtual camera views are shown in Fig. 2.

In this context, *trustworthiness* is defined as the degree to which the system behaves as expected without negative consequences [9]–[11]. Conversely, we define *effectiveness* as the degree to which the system produces a desired result. A system will only be put to use if its users perceive it to be trustworthy, regardless of whether it is technically safe and effective [12]. This perception can be encouraged by demonstrating to the user that the system will behave as expected, avoiding negative consequences, and, hence, is worthy of their trust. This could be achieved, for example, by allowing users to familiarise themselves with the system in a low risk virtual environment [13].

Another method of developing trust is to improve the user's situational awareness [14]. However, the lighting conditions on the lunar surface are often insufficient to provide clear visual feedback [15]. In terrestrial systems, supporting visual feedback with haptic feedback leads to improvements in several performance metrics [16], [17], but delayed haptic feedback introduces instabilities which reduces performance and trust [8]. A system which offers the benefits of haptic feedback, while accommodating the unavoidable Earth-Moon communications delays, would, therefore, be a valuable tool for future lunar regolith handling tasks.

One possible solution, which provides a low risk environment for training and improves situational awareness through delay-robust feedback, may be to use Model Mediated Teleoperation (MMT), where a remote system is controlled via a virtual intermediary [18], [19]. In such a system, the operator controls a local virtual robot and receives sensory feedback from a local virtual environment, which is updated based

on delayed sensor data from the remote system [19]. The remote subsystem (physical robot) receives commands from the simulation to mirror the behaviour of the virtual robot. As the haptic feedback is not directly coupled to the remote system, MMT ensures robustness to variable communication delays, making delays imperceptible to the operator [18] and providing more stable haptic feedback [19].

The effectiveness of MMT is dependent on the accuracy of the model [18]. Model mismatches can lead to unsafe commands being sent to the robot, or instabilities occurring in haptic rendering during drastic local model updates. Detailed comparisons of the other control strategies proposed for delayed teleoperation are available in the literature [18]–[20].

For fine-grained, cohesive lunar regolith, it is too computationally intensive to compute each individual inter-particle interaction in real-time [21]. A MMT system for manipulating regolith would, therefore, require a more efficient method of approximating the behaviour of regolith, albeit at the cost of reduced accuracy. Pereira and Schmidt developed a computationally efficient model for haptically rendering lunar regolith [22], in which the fine-grained material is modelled as a collection of macroparticles. In our previous work, we verified this model against a physical flow rate experiment [23], showing high accuracy for large sample sizes (500 g), and, therefore, being a potentially suitable candidate for MMT.

In this work, we assess the sim-to-real gap of a lightweight model of lunar regolith with regard to visual feedback, end-effector interaction forces and overall functionality. Through this analysis, we determine whether this model is sufficiently accurate to be considered trustworthy and effective during a lunar regolith simulant handling task using a MMT system.

II. METHODS

In this section, first, we describe the teleoperation system used in these experiments (Section II-A). Then, we describe the simulant handling tasks used to assess the potential benefits of employing the virtual regolith model [22] for MMT (Section II-B). Finally, we outline the procedure for assessing the sim-to-real gap (Section II-C).

A. Teleoperation System Description

The remote subsystem used in our experimental setup included a robot arm (KUKA LBR iiwa7 R800), fitted with a 6-axis force-torque sensor (ATI Axia80-M8) at the wrist and a two-finger gripper (Robotiq Adaptive Robot Gripper 2F-85) as the end effector. The arm was mounted horizontally in the centre of the back plate of a 2 x 1 x 1 m box containing the task environment (Fig. 1). Three cameras (1 x Logitech BRIO and 2 x Logitech C270) captured video streams (640 x 480, 20 fps) of the task from different viewpoints (Fig. 2).

The local subsystem included a haptic feedback controller (Force Dimension sigma.7) and a Linux machine (Intel i7-7700k CPU 4.20 GHz, 16 GB RAM) running ROS-based control software and displaying the three video streams of the task. A simulation of the remote system in the task

environment, including the model of regolith, was built using the Chai3D library [24], and was run locally to display the virtual scene from the same angles as the physical cameras (Fig. 1). The simulation also computed the haptic interactions between the end effector and the virtual regolith, to compare against those sensed by the physical force-torque sensor.

There were three modes for controlling the manipulator.

- *Point-to-Point (PP)*: The robot moved between pre-planned positions along trajectories calculated by the MoveIt! path planning software, after the operator had initiated the movement.
- *Direct Teleoperation (DT)*: The operator used the haptic device to teleoperate the physical robot.
- *Open-Loop Model-Mediated Teleoperation (OL-MMT)*: The operator used the haptic device to teleoperate the virtual robot. The virtual movements were then replicated by the physical robot.

Each control mode was open-loop without signal delay, as opposed to conventional closed-loop MMT. This was because this study aims to assess the model’s sim-to-real gap, not the benefits of feedback.

B. Simulant Handling Tasks

Four tasks of varying complexity, inspired by sample curation tasks [25], were selected to validate whether this model is suitable for a MMT system. Each task was repeated 20 times. The DT and OL-MMT tasks were carried out by a volcanology expert (Author, AZ). The goal of the DT and OL-MMT trials was to examine the sim-to-real gap in response to a wider variety of motion profiles than in the PP tasks. This was not to compare results between different users, and thus, only a single operator was required.

1) *Press (PP)*: The robot with a spherical tool as end-effector pressed downwards 10 mm in the -Z direction (Fig. 2) onto the surface of the simulant. This one-dimensional test was selected to determine how accurately the model predicts haptic interactions.




Fig. 2. Camera views of the task in reality and simulation. Cartesian axis directions are shown in green. Positions and orientations of each camera can be seen in Fig. 1.

2) *Drag (PP)*: After inserting to a depth of 10 mm in the -Z direction, the robot dragged the spherical tool horizontally through the simulant 75 mm in the +Y direction. This was selected as a more advanced, two-dimensional test.

3) *Scoop (PP)*: Using a hemispherical scoop tool, the robot collected regolith simulant from the surface. The pre-planned scoop movements consisted of a 180° rolling motion along the main axis of the scoop, starting with the scoop held at a 30° angle to the surface of the simulant. Next, the robot poured the scooped simulant into a tray on a set of scales using the inverse 180° rolling motion, starting from a horizontal position above the tray.

4) *Scoop (DT/OL-MMT)*: The operator teleoperated the scoop tool using the haptic interface to collect simulant from the surface (Fig. 1). This was carried out using either DT while observing physical camera views or OL-MMT while observing virtual camera streams (Fig. 2). Using a human-in-the-loop introduced more variation to the motion profile, compared with the PP tasks, in order to examine the model's behaviour in a range of situations. The mass of collected simulant was then recorded by using PP control to move the scoop over a tray on a set of scales and roll 180° to pour the simulant.

The spherical tool ($\varnothing = 4\text{cm}$) and the hemispherical scoop tool ($\varnothing = 4\text{ cm}$, length = 7cm) were 3D printed in ABS (acrylonitrile butadiene styrene). These tools were replicated in the virtual scene as rigid collision objects. 2 kg of regolith simulant (Exolith LMS-1) was placed in a 20 x 10 x 8 cm container, within the dexterous workspace [26] of the manipulator. The simulant surface was smoothed horizontal before each trial. LMS-1 was selected due to its well documented physical properties, and as the regolith model has previously been experimentally verified against this material [23]. The physical properties (density: 1,470 kg m⁻³, cohesion: 0.393 kPa, internal friction coefficient: 0.284) of this material were replicated in simulation using the regolith model, using macroparticles with radii = 7 mm [22], [23].

C. Measuring the Sim-to-Real Gap

1) *Visual Feedback*: Following the Scoop (PP) task, the length (distance in the X-direction) and width (Y-direction) of the resultant hole in the simulant were compared between simulation and reality. These measurements were taken at the widest points of the holes. In addition, a qualitative comparison of the regolith surfaces was made following the Press (PP) and Drag (PP) tasks. In a realistic lunar teleoperation application, in order to repeatedly carry out the collection task, the user would need to know where regolith remained in the scene, and from where it had been removed. Any indentations in the physical regolith surface should be present in the virtual scene, and if these indentations collapse in reality, this should also be represented virtually.

2) *Force Evaluation*: The sim-to-real gap of the model with respect to the haptic interactions between the tool and the simulant was determined by comparing the measured end-effector forces on the physical system against the virtual end-effector forces computed by the model. This was

assessed for each of the four tasks. Haptic feedback was not sent to the operator during the DT or OL-MMT control trials, so as to not impact the movement profile. The aim of these trials was to determine the accuracy of the computed haptic feedback, not to test whether operators found it helpful.

3) *Effectiveness and Trustworthiness*: To investigate the effectiveness and trustworthiness of the model, we assessed the overall functionality of the system during the scooping tasks in the virtual and physical worlds, using two criteria: 1) the total mass of simulant which was transferred to the collection tray, and 2) the success rates/negative consequences of the scooping tasks. We defined a successful scoop as simulant suspended by the scoop above the simulant bed, and a negative consequence as an unsuccessful scoop. Furthermore, the simulation represents the 'expected behaviour' while reality is the 'actual behaviour'. Therefore, based on the definitions of Section I, we quantified the effectiveness and trustworthiness of the system using the following methodology.

Using the notation, R and S , for success in real and simulated worlds, respectively, as well as $\neg R$ and $\neg S$ for failures, we can describe the four possible results: True Positive (TP), True Negative (TN), False Positive (FP) and False Negative (FN) as in Eq. 1-4.

$$S \wedge R \equiv TP \quad (1)$$

$$\neg S \wedge \neg R \equiv TN \quad (2)$$

$$S \wedge \neg R \equiv FP \quad (3)$$

$$\neg S \wedge R \equiv FN \quad (4)$$

A perfectly trustworthy system will exhibit only TP and TN, with zero FP or FN results, i.e., it will always behave (in reality) as expected (in simulation). FPs - where we expect success but it fails in reality - would be much more damaging to trust than FNs, as the result is a negative consequence [13].

Conversely, a perfectly effective system must produce only TP results, with zero incidences of FPs, i.e., success in simulation never results in failure in reality. Effectiveness is unconcerned with FNs, as these still result in success, albeit unexpectedly.

In order to determine the incidences of TP, FP and FN, the scoop task was carried out under two conditions: the operator views either the physical scene (DT) or the virtual scene (OL-MMT), as in Fig. 2. Each scoop task was repeated 20 times, alternating between DT and OL-MMT to avoid any training-based bias. The incidence rates of TP, FP and FN were used to quantify how trustworthy and/or effective this model is for this scooping task.

III. RESULTS

A. Visual Feedback

Fig. 3 illustrates the visual outcomes of the PP tasks. In each of the tasks, the end effector left an indentation in the surface of the physical regolith simulant. This indentation was similar for the real and virtual worlds for each task,

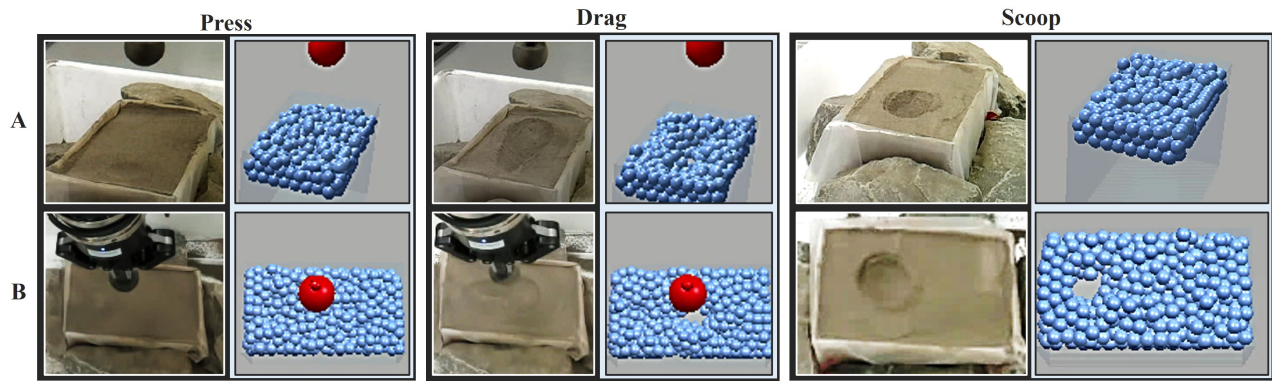


Fig. 3. Comparison of the resultant indentations in the physical (left) and virtual (right) regolith simulant following each of the three PP tasks. Top and bottom rows are the views from cameras A and B, respectively.

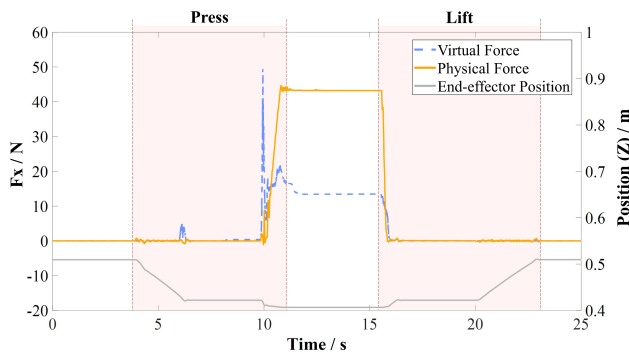


Fig. 4. Comparison of end effector forces in the physical and virtual worlds in relation to the position during the Press (PP) task.

but the discretisation of simulant into macroparticles in the simulation meant that some of the finer details seen in the physical world were lost. Differences between simulation and reality become clearer as the camera zooms in.

The resultant hole diameter in the simulant following each Scoop (PP) task was on average 27% wider and 57% greater in length in the physical test compared with the simulation. The mean physical width and length of the holes were 5.51 ($\sigma = 0.11$) cm and 5.76 ($\sigma = 0.16$) cm, respectively, vs 4.41 ($\sigma = 0.59$) cm and 3.77 ($\sigma = 0.60$) cm in the virtual world.

B. Force Evaluation

Fig. 4 and 5 compare the end effector forces sensed on the physical robot against those calculated by the virtual model, for the Press (PP) and Drag (PP) tasks, respectively. In the press task, the initial peak force in simulation (49 N) was comparable to the peak physical value (45 N). However, the virtual force manifested as a short impulse which reduced to 14 N, while the physical value was sustained at 43 N. This occurred because the virtual macroparticles were displaced, whereas the fine-grained physical material was compressed.

In the drag task, similar Z-axis force profiles were elicited during the 'press' section. In the Y-direction, the physical forces reached and sustained ~ 15 N, whereas the model, again, predicted a peak of 12 N that was not sustained. This was also due to macroparticle displacement. Whilst moving in the +Y direction, during the 'drag' section of the task,

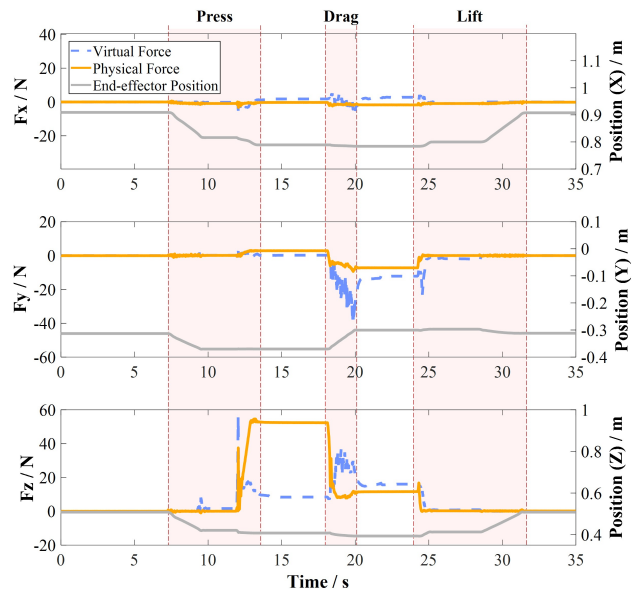


Fig. 5. Comparison of end effector forces in the physical and virtual worlds in relation to the position during the Drag (PP) task.

the model overestimated the forces along all axes; for the physical and virtual worlds, respectively, forces remained steady at -7 N and -12 N in the Y-direction and 11 N and 16 N in the Z-direction. Forces in the X-direction were smaller (physical: -1.7 N, virtual: +2.3 N).

Forces were also compared during the Scoop (PP) task, but these did not show any clear correlation between virtual and physical results.

C. Effectiveness and Trustworthiness

In the scoop task, less material was collected in the virtual environment than in the physical world, for each of the three control modes (Table I). In addition, when the task was carried out based on the virtual visual feedback (OL-MMT), on average 34.4% less simulant was collected in the real world than when physical visual feedback was used (DT). However, similar amounts of material were collected when using PP and DT control modes for the real and virtual environments.

TABLE I

MEAN MASS OF REGOLITH SIMULANT COLLECTED IN SIMULATION AND REALITY DURING SCOOPING TASKS USING DIFFERENT CONTROL MODES.

Control Mode	Mass Collected Mean (SD) / g	
	Real	Sim
PP	30.5 (1.7)	22.8 (2.5)
DT	30.5 (4.3)	19.5 (4.4)
OL-MMT	20.0 (6.0)	14.0 (8.5)

Across 40 trials (DT + OL-MMT), 92.5% were successful in both the virtual and physical worlds (TP). 7.5% were successful in the physical world, but not in the simulation (FN). No significant training effects were recorded over these trials. FNs occurred only in OL-MMT trials: 16, 17, and 20, and the mass collected showed no clear trend with sequential trial numbers. There were no incidences where a successful scoop in simulation resulted in a failure in the real world (FP), nor were there any incidences of failures in, both, the virtual and physical worlds (TN). Based on these trials, the OL-MMT system can be considered effective and trustworthy 100% and 92.5% of the time, respectively, for scooping simulant.

During the pouring motions, however, the difference between the virtual and physical worlds was more pronounced. The simulated scoop rotated 34.6° further than the physical equivalent before simulant material fell from the scoop. Following this, the virtual simulant fell more readily, completing the pour with 37.5° less final rotation than with the real simulant (Fig. 6). This means that during pouring, rotations $>32.6^\circ$ but $<67.2^\circ$ produce a FN outcome, whereas rotations $>86.9^\circ$ and $<124.4^\circ$ produce a FP outcome.




Fig. 6. Mean (SD) start and finish end effector roll angles during pouring motion in reality (left) and simulation (right).

IV. DISCUSSION AND CONCLUSION

In this work, an efficient model of lunar regolith simulant was implemented in an OL-MMT system to manipulate a physical sample, for the first time. The experiments in this work aimed to explore whether the assumptions of the model, developed in [22] and scaled up in [23], are sufficiently accurate to be an effective and trustworthy tool when implemented in MMT. This work builds upon these previous studies by quantifying the sim-to-real gap with respect to haptic interactions of the end effector with lunar regolith simulant.

The results of the PP trials showed that the visual feedback provided by the virtual model provides qualitative information on how the end effector has manipulated the regolith

simulant surface, allowing the operator to monitor where the simulant remains in the environment in order to guide subsequent collection motions. The results of the scoop (PP, DT, OL-MMT) trials demonstrated that this feedback is sufficiently accurate to determine whether simulant has been successfully collected. In fact, there were zero FP cases, and, hence, using this model in a MMT system is unlikely to mislead the operator into believing they have successfully completed the task, when in reality it had failed. However, the model was not perfect – 7.5% of outcomes were FN. To mitigate this, [27] has suggested that delayed physical camera streams of the task could be provided alongside the virtual visual feedback, to reassure operators that the task is proceeding as expected and to continuously update the virtual model.

The results show that the model tends to underestimate the amount of simulant collected, i.e., it was less successful, akin to a partial FN. The impact of this outcome on the effectiveness and trustworthiness of the system would depend on the operational requirements of the task. Assuming there is no penalty for collecting too much material, this would not result in a negative consequence, and would consequently not be detrimental to trustworthiness. However, this model would not be appropriate if precise quantities are required unless the overestimation could be predicted and accounted for. Furthermore, 34.4% less simulant was collected in the OL-MMT trials than DT. Therefore, there must be a trade-off between the less efficient collection during MMT and the operational challenges of delayed teleoperation [8].

During pouring, the virtual and physical systems behaved differently. The virtual simulant poured suddenly over a narrow range of angles, whereas the physical simulant flowed more gradually over a wider angle range. This inaccuracy occurred because the simulant was discretised into macroparticles. A similar effect was seen in our previous work [23], when pouring small samples of macroparticles. However, the range of pouring angles in the simulation lies in the middle of the range for the physical system, suggesting that the friction and cohesion assumptions of the model accurately represent the physical regolith simulant. To ensure that the model is trustworthy, the end effector angle should be limited to $>124.4^\circ$ when pouring is desired, or $<32.6^\circ$ in transit.

For the simple tasks (Press and Drag), the interaction forces computed by the simulation were comparable to those sensed on the physical robot. However, the discretisation of the virtual simulant causes sharp changes in forces as particles are displaced. This could be addressed by reducing the macroparticle size, or increasing their friction (against particles and tools), so that they are not so readily displaced.

These results support our previous experimental verification of the model [23]; this model is accurate on a large scale, when the total simulant volume is far greater than the macroparticle size, but is less realistic on a small scale.

A. Limitations and Future Work

For the DT and OL-MMT trials, the inverse kinematics model was solved relative to the manipulator wrist. The offset

distance from this point to the scoop centre caused large tool movements in response to small rotations of the haptic device. Adjusting this to, instead, rotation around the tool centre point can reduce these movements.

Trust in delayed systems can be built by training in simulation with increasing risk and realism [13]. Although the model does not output realistic interaction forces during complex movements, it does visually represent the real world. Therefore, this should be explored as an early training tool to familiarise operators with the system. Future work should also incorporate this model into a closed-loop MMT system that updates its local simulation based on delayed sensor data, and compare its performance and trustworthiness against other control strategies for delayed teleoperation.

A single lunar mare regolith simulant (LMS-1) was used in these trials, but further testing with highland simulants would support upcoming exploration missions [28]. Moreover, the effects of reduced gravity, vacuum conditions must also be investigated.

The granularity of the model's feedback requires further improvement, whilst limiting the number of graphically rendered particles to ensure that it remains computationally lightweight [23]. To address this, the size of macroparticles could be dynamically adjusted based on their proximity to the tool, in order to provide more granular visual and haptic detail at points of interest in the scene. This would also enable greater volumes of regolith to be rendered at haptic rates (1 kHz [29]). Nonetheless, this model would be a valuable tool for overcoming the challenges of delayed teleoperation on the lunar surface, or for training operators ahead of lunar exploration activities.

REFERENCES

- [1] I. A. Crawford, "Lunar resources: A review," *Progress in Physical Geography*, vol. 39, no. 2, pp. 137–167, 2015.
- [2] L. Schlüter and A. Cowley, "Review of techniques for In-Situ oxygen extraction on the moon," *Planetary and Space Science*, vol. 181, no. September 2019, p. 104753, 2020. [Online]. Available: <https://doi.org/10.1016/j.pss.2019.104753>
- [3] A. Austin, B. Sherwood, J. Elliott, A. Colaprete, K. Zacny, P. Metzger, M. Sims, H. Schmitt, S. Magnus, T. Fong, M. Smith, R. P. Casillas, A. S. Howe, G. Voecks, M. Vaquero, and V. Vendiola, "Robotic Lunar Surface Operations 2," *Acta Astronautica*, vol. 176, no. May, pp. 424–437, 2020. [Online]. Available: <https://doi.org/10.1016/j.actaastro.2020.06.038>
- [4] D. Lester and H. Thronson, "Human space exploration and human spaceflight: Latency and the cognitive scale of the universe," *Space Policy*, vol. 27, no. 2, pp. 89–93, 2011.
- [5] E. Ivanova, J. Eden, S. Zhu, G. Carboni, A. Yurkewich, and E. Burdet, "Short Time Delay Does Not Hinder Haptic Communication Benefits," *IEEE Transactions on Haptics*, vol. 14, no. 2, pp. 322–327, 2021.
- [6] M. C. Yip, M. Tavakoli, and R. D. Howe, "Performance analysis of a haptic telemanipulation task under time delay," *Advanced Robotics*, vol. 25, no. 5, pp. 651–673, 2012.
- [7] M. Perez, S. Xu, S. Chauhan, A. Tanaka, K. Simpson, H. Abdul-Muhsin, and R. Smith, "Impact of delay on telesurgical performance: study on the robotic simulator dV-Trainer," *International Journal of Computer Assisted Radiology and Surgery*, vol. 11, no. 4, pp. 581–587, 4 2016.
- [8] J. Louca, K. Eder, J. Vrubleviskis, and A. Tzemanaki, "Impact of Haptic Feedback in High Latency Teleoperation for Space Applications," *ACM Transactions on Human-Robot Interaction*, 2024.
- [9] A. Avizienis, J. Laprie, and B. Randell, "Fundamental Concepts of Dependability," no. 010028, 2001. [Online]. Available: <http://citeseer.ist.psu.edu/avizienis00fundamental.html>
- [10] K. Chopra and W. A. Wallace, "Trust in electronic environments," *Proceedings of the 36th Annual Hawaii International Conference on System Sciences, HICSS 2003*, 2003.
- [11] K. E. Schaefer, "The Perception And Measurement Of Human-robot Trust," Ph.D. dissertation, University of Central Florida, 2013.
- [12] D. B. Abeywickrama, A. Bennaceur, G. Chance, Y. Demiris, A. Kordoni, M. Levine, L. Moffat, L. Moreau, M. R. Mousavi, B. Nuseibeh, S. Ramamoorthy, J. O. Ringert, J. Wilson, S. Windsor, and K. Eder, "On Specifying for Trustworthiness," 2022. [Online]. Available: <http://arxiv.org/abs/2206.11421>
- [13] J. Louca, J. Vrubleviskis, K. Eder, and A. Tzemanaki, "Elicitation of trustworthiness requirements for highly dexterous teleoperation systems with signal latency," *Frontiers in Neurobotics*, vol. 17, p. 1187264, 2023.
- [14] A. Ezenyilimba, M. Wong, A. Hehr, M. Demir, A. Wolff, E. Chiou, and N. Cooke, "Impact of transparency and explanations on trust and situation awareness in human-robot teams," *Journal of cognitive engineering and decision making*, vol. 17, no. 1, pp. 75–93, 2023.
- [15] S. Coloma, C. Martinez, B. Can Yalcin, and M. Olivares-Mendezis, "Enhancing Rover Teleoperation on the Moon with Proprioceptive Sensors and Machine Learning Techniques," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 434–11 441, 2022.
- [16] I. El Rassi and J. M. El Rassi, "A review of haptic feedback in tele-operated robotic surgery," *Journal of Medical Engineering and Technology*, vol. 44, no. 5, pp. 247–254, 2020. [Online]. Available: <https://doi.org/10.1080/03091902.2020.1772391>
- [17] V. Pruks, K. H. Lee, and J. H. Ryu, "Shared Teleoperation for Nuclear Plant Robotics Using Interactive Virtual Guidance Generation and Shared Autonomy Approaches," *2018 15th International Conference on Ubiquitous Robots, UR 2018*, pp. 91–95, 2018.
- [18] X. Xu, B. Cizmeci, C. Schuwerk, and E. Steinbach, "Model-Mediated Teleoperation: Toward Stable and Transparent Teleoperation Systems," *IEEE Access*, vol. 4, pp. 425–449, 2016.
- [19] T. Hulin, M. Panzirsch, H. Singh, A. Coelho, R. Balachandran, A. Pereira, B. M. Weber, N. Bechtel, C. Riecke, B. Brunner, N. Y. Lii, J. Klodmann, A. Hellings, K. Haggmann, G. Quere, A. S. Bauer, M. Sierotowicz, R. Lampariello, J. Vogel, A. Dietrich, D. Leidner, C. Ott, G. Hirzinger, and A. Albu-Schäffer, "Model-Augmented Haptic Telemanipulation: Concept, Retrospective Overview, and Current Use Cases," *Frontiers in Robotics and AI*, vol. 8, no. June, pp. 1–22, 2021.
- [20] S. N. F. Nahri, S. Du, and B. J. Van Wyk, "A Review on Haptic Bilateral Teleoperation Systems," *Journal of Intelligent and Robotic Systems: Theory and Applications*, vol. 104, no. 1, 2022.
- [21] H. Otto, K. Kerst, C. Roloff, G. Janiga, and A. Katterfeld, "CFD-DEM simulation and experimental investigation of the flow behavior of lunar regolith JSC-1A," *Particuology*, vol. 40, pp. 34–43, 2018. [Online]. Available: <https://doi.org/10.1016/j.partic.2017.12.003>
- [22] A. Pereira and A. Schmidt, "Efficient Haptic Rendering of Regolith," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 975–13 981.
- [23] J. Louca, J. Vrubleviskis, K. Eder, and A. Tzemanaki, "Verification of a Virtual Lunar Regolith Simulant," *Frontiers in Space Technologies*, vol. 5, p. 1303964, 2023.
- [24] F. Conti, F. Barbagli, R. Balaniuk, M. Halg, C. Lu, D. Morris, L. Sentis, J. Warren, O. Khatib, and K. Salisbury, "The CHAI libraries," in *Proceedings of Eurohaptics 2003*, Dublin, Ireland, 2003, pp. 496–500.
- [25] R. Fukai, T. Usui, W. Fujiya, Y. Takano, K. i. Bajo, A. Beck, E. Bonato, N. L. Chabot, Y. Furukawa, H. Genda, Y. Hibiya, F. Jourdan, T. Kleine, M. Koike, M. Matsuoka, Y. N. Miura, F. Moynier, R. Okazaki, S. S. Russell, H. Sumino, M. E. Zolensky, H. Sugahara, S. Tachibana, K. Sakamoto, M. Abe, Y. Cho, and K. Kuramoto, "Curation protocol of Phobos sample returned by Martian Moons eXploration," *Meteoritics and Planetary Science*, vol. 59, no. 2, pp. 321–337, 2024.
- [26] A. Kumar and K. J. Waldron, "The workspaces of a mechanical manipulator," 1981.
- [27] R. Luz, J. L. Silva, and R. Ventura, "Enhanced Teleoperation Interfaces for Multi-Second Latency Conditions: System Design and Evaluation," *IEEE Access*, vol. 11, no. January, pp. 10 935–10 953, 2023.
- [28] M. Smith, D. Craig, N. Herrmann, E. Mahoney, J. Krezel, N. McIntyre, and K. Goodliff, "The Artemis Program: An Overview of NASA's Activities to Return Humans to the Moon," *IEEE Aerospace Conference Proceedings*, pp. 1–10, 2020.
- [29] B. Hannaford and A. M. Okamura, "Haptics," in *Springer Handbook of Robotics*, 2nd ed. Springer, Cham, 2016, pp. 1063–1084.