Communicating Efficiently to Enable Human-Multi-Robot Collaboration in Space Exploration

Stewart Jamieson
Jessica E. Todd
sjamieson@mit.edu
jetodd@mit.edu
MIT-WHOI Joint Program in Oceanography/Applied Ocean Science and Engineering
Cambridge, MA, USA

Jonathan P. How
jhow@mit.edu
Department of Aeronautics and Astronautics, MIT
Cambridge, MA, USA

Yogesh Girdhar
ygirdhar@whoi.edu
Applied Ocean Physics and Engineering, WHOI
Woods Hole, MA, USA

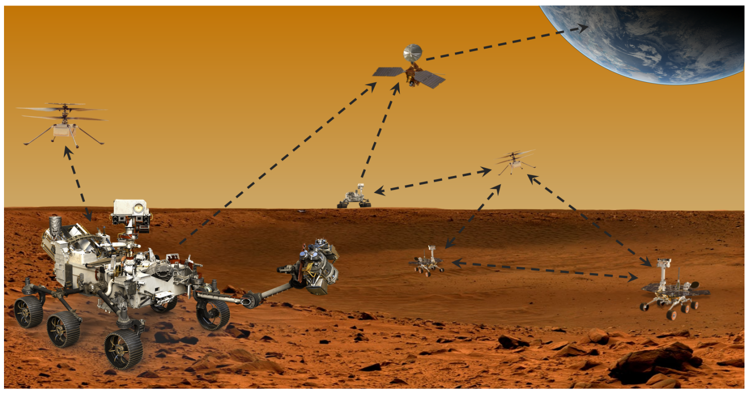

Figure 1: Future exploration teams will be composed of many robotic agents that must communicate and cooperate to succeed.

ABSTRACT

The future of deep space exploration demands a shift in the current paradigm of robotic planetary exploration. Mission Operators on Earth spend days planning just short sequences of robotic exploration activities on Mars due to the limited-bandwidth, high-latency interplanetary communications and low levels of robot autonomy. We propose deploying robots that perform unsupervised semantic mapping of their environments as well as active learning to improve the efficiency of data exchange and activity planning between robots and human operators. Robots that build and transmit high-level semantic maps and learn by querying operators on the science objectives of the mission can improve mission productivity. Multi-robot space missions may utilize similar methods to efficiently communicate between different robotic assets. As robot exploration teams become larger and heterogenous, the sophistication of human-robot and multi-robot cooperation should grow with them.

SpaceCHI '21, Saturday May 14, 2021, Yokohama, Japan

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**; • **Hardware** → *Sensor applications and deployments*; • **Computer systems organization** → **Robotic autonomy**.

KEYWORDS

multi-robot systems, co-robotic systems, remote sensing, semantic mapping, active learning, deep space exploration

ACM Reference Format:
Stewart Jamieson, Jessica E. Todd, Jonathan P. How, and Yogesh Girdhar. 2021. Communicating Efficiently to Enable Human-Multi-Robot Collaboration in Space Exploration. In *Proceedings of SpaceCHI: Human-Computer Interaction for Space Exploration (SpaceCHI '21).* ACM, New York, NY, USA, 4 pages.

1 INTRODUCTION

Planetary robotic exploration is a time-costly endeavour. Mission controllers count their days in Sols, readjusting their sleep cycles to the Martian calendar, in order to curate daily navigation plans for Mars rovers that may well be rendered obsolete when the rover crests the next hill. The cycle of downlinking data, analyzing, activity planning, and uplinking a day's plan can take up to 3 Sols, and waiting for communication windows can lead to entire "restricted sols" where rover activities are severely limited [8]. Imagine a fleet of planetary robots that performs this daily planning on their own, querying Mission Control when they've found something of interest, exchanging scientific information, and refining their own scientific modelling with each new communication from its Science Operations team. More intelligent human-robot and multi-robot communication and cooperation will enable missions to achieve greater scientific returns with the same hardware, thus increasing the efficiency and return on investment of deep space exploration.

A key challenge in this paradigm shift is that a robot collects vast amounts of *data* about its environment, yet lacks its human operator's understanding of the scientific *context* needed to interpret that data [8] and make decisions about how to proceed. Deep space robots communicate without the advantage of modern high-speed communications that would facilitate the transfer of such data, and contend with transmission rates below 1 Mbps, latencies from minutes to hours, and intermittent channel availability. These communications limitations have impeded the development and deployment of effective autonomous exploration and multi-robot distributed exploration systems. This paper will explore how new paradigms for human-robot interaction, *human-robot cooperative planning* and *multi-robot federations*, enable communication bandwidth to be utilized far more efficiently in order to maximize the scientific return of the exploration mission achievable under given communication limitations.

2 RECENT WORKS

A popular approach to increase scientific returns is "opportunistic science", where the science team specifies simple tests that the robot autonomously performs to decide whether to activate other science instruments [5]. For example, the Mars Science Laboratory Curiosity uses the AEGIS system to detect rocks matching pre-specified characteristics while moving; this enables the rover to autonomously target other sensors at such targets and have the data ready for the next communication cycle, avoiding hours of back-and-forth interplanetary communication [7].

More complex models enable robots to understand enough about their environment and scientific objectives to autonomously determine targets of interest to visit. This "autonomous planning" vastly outperforms human remote control in achieving well-specified scientific objectives, and requires no human-robot interaction, making it far more scalable. For example, Bayesian networks have been used by an analog Mars rover to model geological science objectives as well as the spatial distribution of geological phenomena, enabling the rover to autonomously collect a large number of scientifically valuable observations in short time with no human oversight [1]. However, while a team of scientists can work to specify the science objectives that guide this autonomous planning, the rover itself cannot yet be deployed with the domain knowledge and judgment of even one scientist. Thus a fully autonomous robot's inflexible and imperfect understanding of the scientific objectives may lead it to ignore unexpected high-value targets, or spend too much time observing relatively low-value targets.

A key motivation for crewed missions has been that humans are adaptable and can communicate plans and ideas, whereas robots can only collect and transmit raw data [13]. This is rapidly changing, as novel *unsupervised semantic mapping* systems enable robots to autonomously understand and recognize patterns in their operating environment and more efficiently communicate with other robots and humans [12]. Semantic maps compactly represent the variety and spatial distribution of unique phenomena in the environment using a few discrete labels; for example, in a colour-coded semantic map of Mars, one colour would represent sandstone while another would represent basalt rocks. New labels can be learned *in situ* by the robot and, since many natural phenomena are strongly spatially correlated, these maps tend to be highly compressible and much smaller than natural images while still being human-interpretable [9]. Human operators can easily choose which parts of a semantic map to request visual observations of and use those images to inform their decisions on where to send the robot next. Alternatively, if the robot can be taught how the learned semantic labels correlate to high-level mission objectives, it can autonomously handle low-level planning.

3 MAXIMIZING COMMUNICATION EFFICIENCY IN COOPERATIVE PLANNING

Systems for interplanetary human-robot communication and control are often characterized through the lens of *sliding autonomy* [4], presented visually in Fig. 2, left. Greater autonomy is appealing because it enables greater mission returns when communication is highly constrained or unavailable, and provides easier scalability to large multi-robot deployments. While sliding autonomy is a useful paradigm in many areas of human-robot interaction, this perspective reinforces the false beliefs that more autonomy is always better and higher levels of autonomy should mean less human-robot interaction [2].[1] A better way to characterize an exploration system is by its ability to leverage communication bandwidth to help in achieving the science objectives, as seen in Fig. 2, right. For example, when communications are increasingly constrained, remote controlled systems have decreasing performance, whereas systems without human-robot interaction, such as waypoint following robots, will have constant performance. This new perspective highlights that developing more sophisticated techniques that enable humans and robots to communicate more efficiently is potentially more important than only expanding what robots can do autonomously.

As an example, consider how perfectly specifying the scientific objectives in advance of the mission is impossible when little is known *a priori* about the exploration environment, presenting a barrier to fully autonomous planning. Human-robot cooperative planning overcomes this by having the robot and human exchange just enough information that the robot understands the science objectives well enough to go where the human would have instructed it if they could access all the robot's data. We propose that the best way to implement this paradigm is for the robot to handle semantic mapping and most low-level planning, but to query the operator with questions about the mission objectives. Semantic mapping means that the robot understands what it has found in its environment and can predict what it might observe in unvisited locations. This task is typically performed with machine learning algorithms, such as [15], and is well suited to the robot team member as it always has access to the most up-to-date data.

[1] Consider how humans, generally held as the gold-standard for autonomous agents, still make extensive use of human-human interactions in achieving complex tasks.

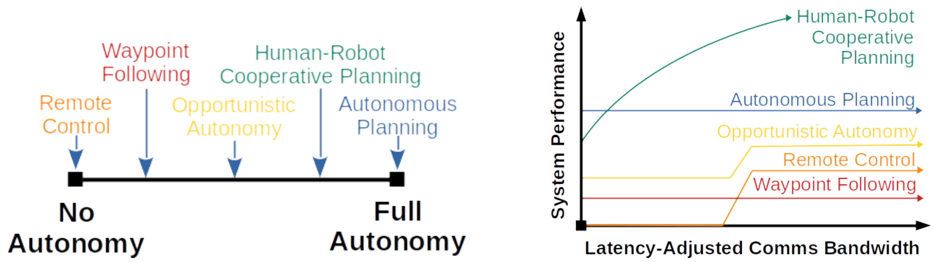

Figure 2: Left: *Sliding autonomy* ranks systems based on how much of their behaviour is fully autonomous. Right: We argue that a more useful perspective considers how well the system utilizes its communication budget to achieve mission objectives.

Learning the mission objectives during the mission using queries is feasible because the robot only needs to ask the operators questions that help it determine which trajectories are best aligned with the mission objectives, so that it makes the same decisions that they would. As long as the number of unique types of phenomena encountered increases sublinearly over time, only a small subset of data needs to be sent to the operator to learn how those phenomena relate to the mission objectives. In ocean exploration analog experiments, this approach to human-robot cooperative exploration was found to result in vastly more productive missions than those based on the more traditional approach of pre-planned waypoints [11]. It was also found that some algorithms for choosing which questions to ask the operator resulted in greater performance over certain ranges of communication bandwidth levels, demonstrating the importance of evaluating a novel system's performance as a function of bandwidth.

4 THE FUTURE: SCALING THROUGH MULTI-ROBOT FEDERATIONS

With the upper limit of landable rover mass being reached, it is likely current monolithic systems such as the Curiosity rover will be replaced by smaller, cooperative units with distributed functionality. This trend is being seen in current and planned space exploration missions. Building on the technology demonstration of the Ingenuity helicopter accompanying the Mars Perseverance rover, future Mars rovers will be accompanied by a scouting helicopter that can assist in identifying science targets, retrieval of scientific caches for analysis and return to Earth, and path planning and obstacle avoidance [3]. Larger heterogeneous robot teams equipped with specialized functions will enable distributed exploration of even more diverse environments, such as subsurface caves and lavatubes on the Moon or Mars [6, 10]. To coordinate such teams, where robotic assets carry different sensor and instruments suites and are designed with different mobilities, pertinent science and mapping information will need to be communicated succinctly and efficiently between assets. Lightweight aerial vehicles might provide high level mapping to ground assets that can make more computationally intensive planning decisions, and direct aerial and ground vehicles to new sites of interest. Outer planet worlds like Titan and Europa offer exciting science opportunities for these types of robotic exploration [13] however the extreme communication latencies involved in such missions will make efficient exploration and communication between robots and operators even more critical mission requirements.

In order to efficiently meet science objectives, multi-robot missions (e.g. [6, 16]) will require effective means of communicating key science data between assets without relying on constant human input. Thus we extend our earlier concepts to the human-multi-robot setting: the team (human and robots) should share just enough information among each other that each robot performs actions the humans would command if they had access to the data of every robot. For the robots to agree on their optimal actions, it is sufficient that they reach consensus on the global semantic map, their individual positions on that map, and the mission objectives. Multi-Robot Simultaneous Localization and Mapping enables robots to reach consensus on their relative locations and the global semantic map [14], while consensus on mission objectives requires that answers to queries are shared between robots.

These conditions ensure that the exploration task is efficiently distributed so that the scientific return of N robots is about N times the return of a single robot. However, we believe superlinear scaling is possible by leveraging multi-robot *federations*; this describes a system in which the robots actively plan to help each other explore more efficiently. For example, one robot may take an opportunity to move closer to another in order to help it construct its semantic map, or may ask the operators a question meant to help another robot. As robotic explorers develop a more sophisticated understanding of their environment and objectives, and teams become heterogeneous, the possibilities for effective collaboration will expand exponentially.

REFERENCES

[1] Akash Arora, Robert Fitch, and Salah Sukkarieh. 2017. An approach to autonomous science by modeling geological knowledge in a Bayesian framework. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3803–3810. https://doi.org/10/ggb45r

[2] Jenay M. Beer, Arthur D. Fisk, and Wendy A. Rogers. 2014. Toward a framework for levels of robot autonomy in human-robot interaction. *Journal of human-robot interaction* 3, 2 (July 2014), 74–99. https://doi.org/10/gdz57h

[3] J. Delaune, R. Brockers, D. S. Bayard, H. Dor, R. Hewitt, J. Sawoniewicz, G. Kubiak, T. Tzanetos, L. Matthies, and J. Balaram. 2020. Extended navigation capabilities for a future mars science helicopter concept. In *2020 IEEE Aerospace Conference*. IEEE, Big Sky, MT, USA, 1–10. https://doi.org/10/gh2944

[4] M Bernardine Dias, Balajee Kannan, Brett Browning, E Gil Jones, Brenna Argall, M Freddie Dias, Marc Zinck, Manuela M Veloso, and Anthony J Stentz. 2008. *Sliding Autonomy for Peer-To-Peer Human-Robot Teams.* Technical Report CMU-RI-TR-08-16. Carnegie Mellon University, Pittsburgh, PA. 12 pages. https://www.ri.cmu.edu/publications/sliding-autonomy-for-peer-to-peer-human-robot-teams/

[5] Tara A. Estlin, Benjamin J. Bornstein, Daniel M. Gaines, Robert C. Anderson, David R. Thompson, Michael Burl, Rebecca Castaño, and Michele Judd. 2012. AEGIS automated science targeting for the MER Opportunity rover. *ACM Transactions on Intelligent Systems and Technology* 3, 3 (2012), 1–25. https://doi.org/10/gfvsn5

[6] W. Fink, V. R. Baker, D. Schulze-Makuch, C. W. Hamilton, and M. A. Tarbell. 2015. Autonomous exploration of planetary lava tubes using a multi-rover framework. In *2015 IEEE aerospace conference*. 1–9. https://doi.org/10.1109/AERO.2015.7119315 tex.ids= Fink2015.

[7] R. Francis, T. Estlin, G. Doran, S. Johnstone, D. Gaines, V. Verma, M. Burl, J. Frydenvang, S. Montaño, R. C. Wiens, S. Schaffer, O. Gasnault, L. DeFlores, D. Blaney, and B. Bornstein. 2017. AEGIS autonomous targeting for ChemCam on Mars Science Laboratory: Deployment and results of initial science team use. *Science Robotics* 2, 7 (June 2017). https://doi.org/10.1126/scirobotics.aan4582 tex.ids= Francis2017a tex.elocation-id: eaan4582 tex.eprint: https://robotics.sciencemag.org/content/2/7/eaan4582.full.pdf publisher: Science Robotics section: Research Article.

[8] Daniel Gaines, Robert Anderson, Gary Doran, William Huffman, Heather Justice, Ryan Mackey, Gregg Rabideau, Ashwin Vasavada, Vandana Verma, Tara Estlin, Lorraine Fesq, Michel Ingham, Mark Maimone, and Issa Nesnas. 2016. Productivity Challenges for Mars Rover Operations. In *Planning and Robotics Workshop of ICAPS 2016*. London, England, 12. https://trs.jpl.nasa.gov/handle/2014/46121

[9] Yogesh Girdhar, Levi Cai, Stewart Jamieson, Nathan McGuire, Genevieve Flaspohler, Stefano Suman, and Brian Claus. 2019. Streaming Scene Maps for Co-Robotic Exploration in Bandwidth Limited Environments. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, Montréal, Canada, 7940–7946. https://doi.org/10/ggb46q

[10] A. Husain, H. Jones, B. Kannan, U. Wong, T. Pimentel, S. Tang, S. Daftry, S. Huber, and W. L. Whittaker. 2013. Mapping planetary caves with an autonomous, heterogeneous robot team. In *2013 IEEE aerospace conference*. 1–13. https://doi.org/10.1109/AERO.2013.6497363

[11] Stewart Jamieson, Jonathan P. How, and Yogesh Girdhar. 2020. Active Reward Learning for Co-Robotic Vision Based Exploration in Bandwidth Limited Environments. In *2020 International Conference on Robotics and Automation (ICRA)*. IEEE, Paris, France. https://arxiv.org/abs/2003.05016

[12] Jeffrey W. Kaeli, John J. Leonard, and Hanumant Singh. 2014. Visual summaries for low-bandwidth semantic mapping with autonomous underwater vehicles. In *2014 IEEE/OES Autonomous Underwater Vehicles (AUV)*. IEEE, 1–7. https://doi.org/10/ggb452

[13] National Research Council. 2011. *Visions and Voyages for Planetary Science in the Decade 2013–2022*.

[14] Sajad Saeedi, Michael Trentini, Mae Seto, and Howard Li. 2016. Multiple-Robot Simultaneous Localization and Mapping: A Review: Multiple-Robot Simultaneous Localization and Mapping. *Journal of Field Robotics* 33, 1 (Jan. 2016), 3–46. https://doi.org/10/f78mgt

[15] John E. San Soucie, Heidi M. Sosik, and Yogesh Girdhar. 2020. Gaussian-Dirichlet Random Fields for Inference over High Dimensional Categorical Observations. In *2020 International Conference on Robotics and Automation*. IEEE, Paris, France. http://arxiv.org/abs/2003.12120 arXiv: 2003.12120.

[16] A. Tagliabue, S. Schneider, M. Pavone, and Ali-akbar Agha-mohammadi. 2020. Shapeshifter: A multi-agent, multi-modal robotic platform for exploration of titan. In *2020 IEEE aerospace conference*. 1–13. https://doi.org/10/gh295z