AIAA SciTech Forum
8-12 January 2024, Orlando, FL
AIAA SCITECH 2024 Forum

10.2514/6.2024-1747

To Boldly Go Where No Robots Have Gone Before – Part 4: NEO Autonomy for Robustly Exploring Unknown, Extreme Environments with Versatile Robots

Rohan Thakker*, Michael Paton*, Bryson Jones*, Guglielmo Daddi*, Rob Royce*, Michael Swan*, Marlin Strub*, Sina Aghli, Harshad Zade, Yashwanth Nakka, Tiago Vaquero, Joseph Bowkett, Daniel Loret, Jenny Zhang, Jack Naish, Daniel Pastor, Tristan Hasseler, Carl Leake, Benjamin Nuernberger, Pedro Proença, William Talbot, Kyohei Otsu, Andrew Orekhov[†], Philippe Tossi, Matthew Gildner, Abhinandan Jain, Rachel Etheredge, Matthew Travers,[†] Howie Choset[†], Joel Burdick[‡], Michel Ingham, Matthew Robinson, Masahiro Ono
These authors contributed equally to this work.
[‡] *Department of Mechanical & Civil Engineering, California Institute of Technology, Pasadena, CA, USA*
[†] *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA*
Others: NASA Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA.

This paper introduces NEO, a novel autonomy framework for controlling a versatile high-degree-of-freedom (DOF) robots such as EELS (a screw-driven snake-like robot), aimed at exploring unknown and extreme environments like the geysers of Enceladus or the subsurface oceans of icy worlds. Distinct from conventional Mars mission strategies, NEO embodies resilience, adaptivity, and risk awareness. NEO supports fault-aware perception using both exteroception and proprioception, inspired by a blind climber's feat of scaling El Capitan. NEO tightly couples planning, perception, and control, along with leveraging machine-learning-based methods for adaptation. Moreover, NEO incorporates risk-aware decision making with integrated task and motion planning under consideration of uncertainty, enabling autonomous adaptation of actions to mitigate risks and maximize mission success. This paper presents the architecture of NEO, along with experimental results showcasing these capabilities and discusses the potential for NEO in spearheading a new paradigm in space exploration.

Figure 1 Demonstrating a paradigm for autonomy algorithms to enable space exploration into unknown worlds beyond Mars by eliminating the need for a prior orbiter or a lander mission.

I. Introduction

As we shift our gaze beyond the familiar Martian landscapes towards the intriguing but challenging realms of our solar system, such as the active geysers of Enceladus or the subsurface oceans of icy worlds, a paradigm shift is required in our approach to robotic exploration. As outlined in the companion paper (Part 1), the essence of this shift lies in the symbiosis of a robot possessing versatile mobility and an intelligent, risk-aware autonomy system. In this paper, we spotlight NEO, a groundbreaking autonomy framework developed specifically to control high-DOF versatile robots such as EELS (Part 2) in a highly unknown, extreme environment. The conventional Mars Missions have paved the way for our understanding of autonomous robotic operations in extraterrestrial contexts. However, their modus operandi, largely built on conservative, risk-averse, and protective strategies, while effective in well-understood Martian environments, may find itself ill-equipped to handle the unprecedented uncertainties and challenges that lie in the exploration of unfamiliar worlds.

Traditionally, autonomy on Mars rovers has revolved around a protective paradigm. The primary objective is to ensure the safety and survival of the rover in an unexpected event or anomaly. Hence, the autonomy system is designed to halt activities and switch the rover into a safe mode, where it patiently awaits further instructions from Earth. However, in a context where communication delays can be substantial, and the environment is dynamically changing, this reactive approach might not suffice. In these situations, we propose a shift towards resilience, where the robot is designed to persistently pursue its mission, leveraging remaining capabilities to work around disruptions or failures.

Similarly, past rovers have functioned based on a deliberative paradigm, where activities are planned and scheduled well in advance, based on the available knowledge. While this strategy serves well for environments with minimal dynamism and change, it is likely to falter in the face of environments filled with unknown variables and rapid changes. We argue for an adaptive approach, where the robot can dynamically select activities and adjust its actions in response to real-time sensory inputs and events.

Finally, our conventional risk management strategy has been risk-averse, employing pre-programmed behaviors that ensure safety under any imaginable situations. This conservative approach, while offering a high degree of certainty and protection, may limit the potential for scientific discovery in environments that are vastly unknown and can't be accurately predicted beforehand. NEO proposes a risk-aware strategy, where the autonomy system can adjust its behaviors according to observed situations, carrying out an onboard risk-assessment, and making informed decisions based on the trade-off between risk and reward.

This paper delves into the inner workings of NEO's framework, discussing its fault-aware perception system with exteroception and proprioception, its integrated task and motion planning, and its adaptive strategy with tightly coupled planning, perception, and control. We also describe approaches for leveraging machine-learning techniques to enable generalized adaptation across a wide range of operating condition. Through the lens of NEO, we explore how a blend of resilience, adaptability, and risk-awareness can shape the future of robotic exploration in unknown, extreme environments. We present this work alongside three companion papers, where they provide a broad overview of the vision, technologies, scientific impacts, capabilities, and field test results of EELS [1], the development of EELS hardware [2], and active-skin propulsion of EELS [3].

II. Technologies and Experimental Results

Through the lens of NEO, we explore how a blend of resilience, adaptability, and risk-awareness can shape the future of robotic exploration in unknown, extreme environments. In this section, we describe how we translate each of these high-level concepts into functional and tangible capabilities provided by NEO, with demonstrations on hardware in real-world scenarios. Figure 2 displays an overview of the hardware platforms, the EELS 1.0 and EELS 1.5 robots, on which we deploy NEO. For more details with respect to the hardware platforms, we refer to our companion paper [2].

A. Resilience: Fault-Aware Perception with Exteroception and Proprioception

Operating in extreme environments such as icy terrains, or atmospheres filled with fog, dust, or snow, introduces significant opportunities for perceptual degradation, a common failure mode for many field robotic systems. To counteract these challenges, NEO employs a multi-modal sensor suite that includes stereo cameras, lidar, inertial measurement units (IMU), force-torque sensors, and absolute joint encoders. This variety of sensing modalities enhances the system's ability to perceive its surroundings accurately even under adverse conditions. We achieve much higher resiliency by exploiting the redundancy and heterogeneity in sensing modalities.

However, perceptual degradation in environments such as plumes of Enceladus or lack of geometric/appearance features can cause failures in the perception system. One of the central pillars of our design is an efficient mechanism

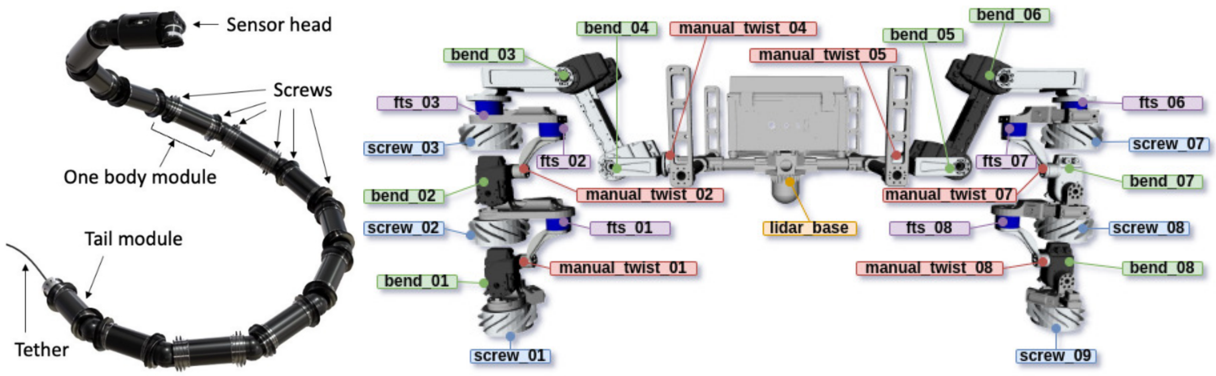

Figure 2 The EELS 1.0 and EELS 1.5 platforms, respectively

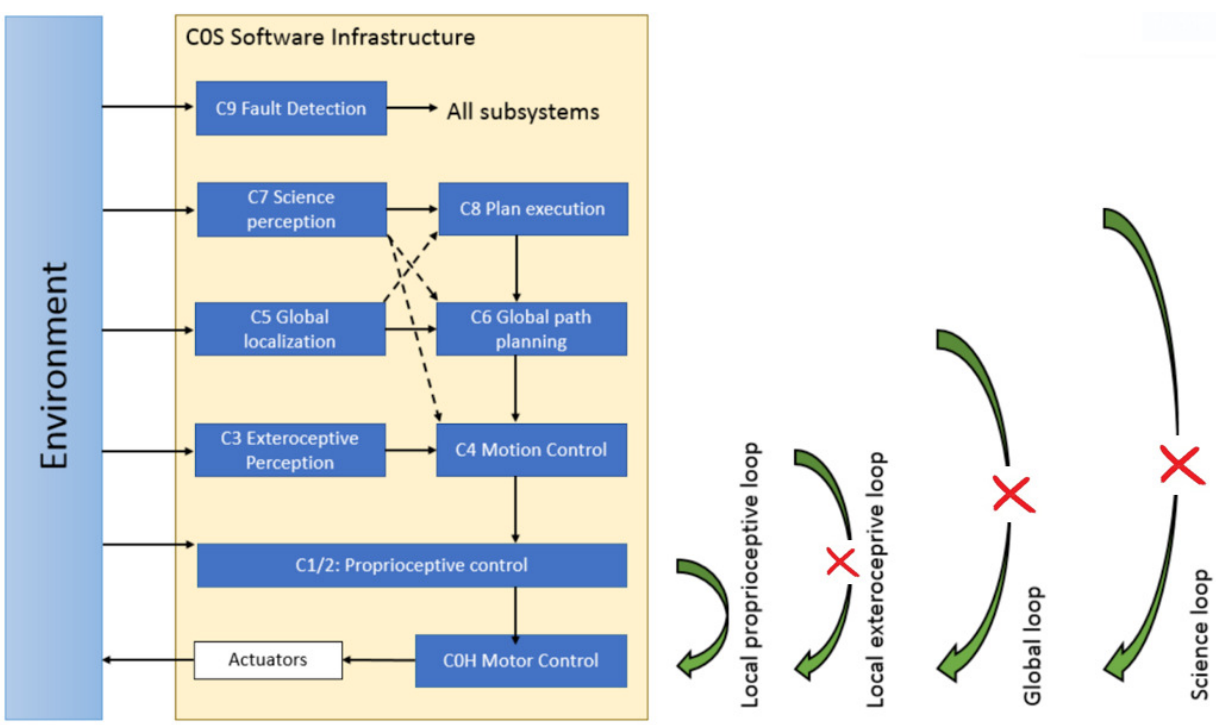

Figure 3 Module architecture for NEO, displaying the resilience proivded by the framework; when higher level modules fail, the mission continues

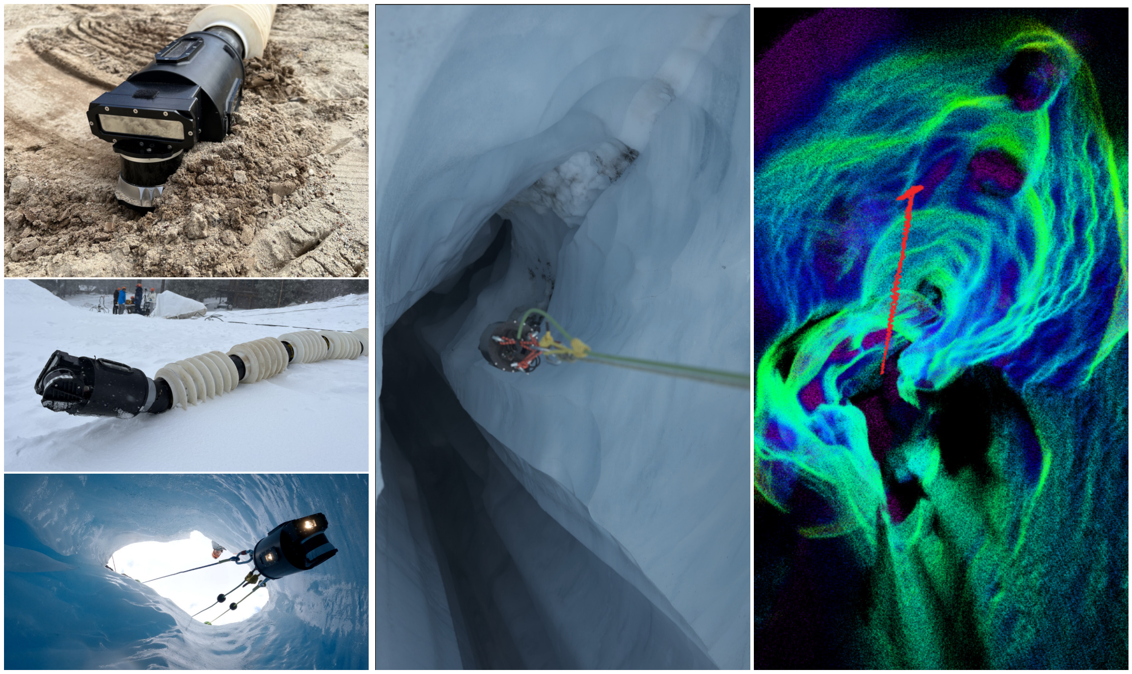

(a) EELS 1.0 sensor head deployed in extreme environments **(b) 3-D reconstruction of a Moulin at Athabasca Glacier using NEO's SLAM algorithm, SERPENT [4].**

Figure 4 NEO's Perception Capabilities

for failure-detection and uncertainty quantification [5]. NEO can identify and quantify sensor or system failures. This robust detection mechanism is crucial, as it facilitates the development of contingency plans and alternative navigation strategies on the fly, thereby enhancing the resilience of the system. Details of algorithms along with field deployment results of NEO's (Simultaneous Localization and Mapping) SLAM capabilities are covered in [4] and summarized in Figure 4.

Finally, despite multimodal redundant sensors and failure detection mechanisms, there is still a non-zero chance of getting a false negative perception failure event which can propagate through the planning/control loop and violate the safety constraints of the system. To prevent this, we propose redundant loops with the lower-most acting purely on proprioceptive feedback ensuring the safety of the robot as shown in Figure 3. Drawing parallels with a blind climber's ascent of El Capitan, the EELS system utilizes proprioception as its last resort navigation strategy. While proprioception-based navigation might not yield optimal outcomes due to inherent drifts and lack of absolute positioning information, its deployment serves as a robust fallback mechanism. This redundancy allows the EELS platform to maintain acceptable navigation performance even when all other systems are compromised, mirroring the resilience of the blind climber who relies on tactile and proprioceptive feedback to navigate challenging terrains. Thus, in the face of multiple navigation challenges, EELS demonstrates resilience and adaptability, where system compromises on slight optimality rather than the complete feasibility of the mission avoiding complete halts [6].

B. Adaptability: Tightly Coupled Planning, Perception, & Control with ML-Based Adaptation

High numbers of DOF from the hardware system make the robot versatile, however, to leverage its complete potential the robot also needs appropriate intelligence. Traditionally, autonomy is achieved by decoupling perception, planning, and control problems into separate modules to create a sense-plan-act loop. This makes the problem tractable but this decoupling doesn't generalize to traversal on extreme terrains e.g. while making a 5m long snake go into a crevasse on Enceladus or a steep slope on a crater on the Moon where the defining obstacles in the environment becomes challenging. As shown in Figure 5, viewing the eels as a rover might make you think that the system is over-actuated. However, EELS is closer to a robotic arm with each screw as an end effector (6 x 10 DOF) with 6 additional DOF for the floating base. Hence, our challenge is to control 66 DOF with just 30 actuators which makes it a massively underactuated control problem. NEO addresses this by formulating a loco-manipulation problem via tightly coupled perception, planning, and

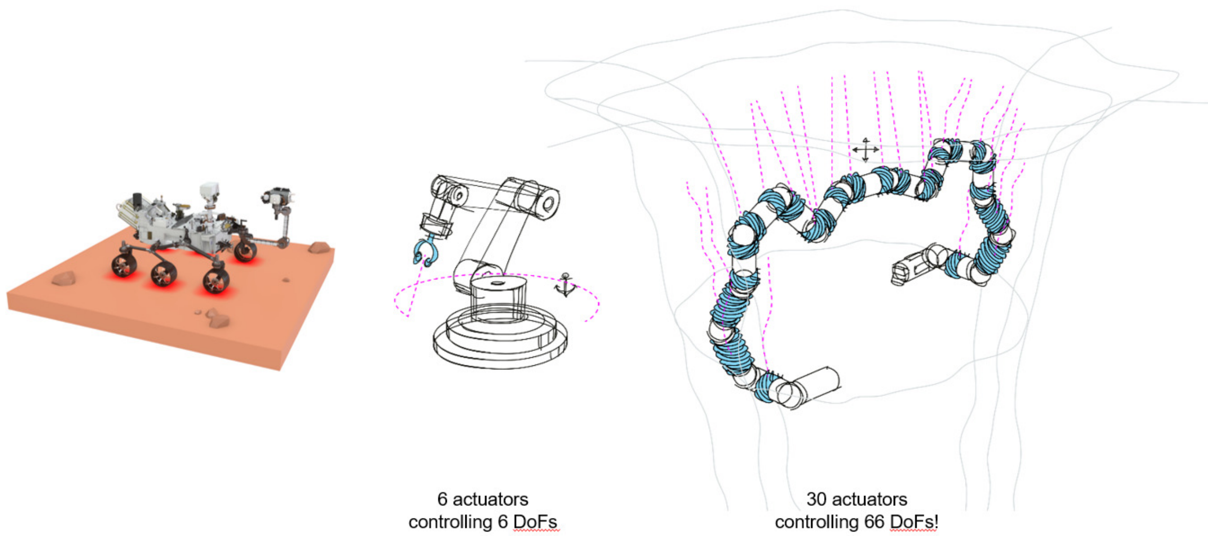

Figure 5 EELS as a Loco-manipulator: The Perseverance Mars Rover, a generic manipulator, and an EELS concept robot juxtaposed together

control loops. This is achieved by using a convex optimization formulation, specifically linear or quadratic programs. These algorithms provide completeness guarantees such that they find a solution if it exists and report a failure if it doesn't. Hence, the planner can run the controller library to check the validity of a state (i.e. whether it is an obstacle state or not).

A key component of technology that enables this is the simultaneous control of shape, contact, force, and motion (SCFM) of a platform, purely through proprioceptive sensing capabilities [7]. This control strategy enables vertical mobility by providing a method of locomotion that can navigate along undulating terrain, with emergent fault recovery behaviors. We deployed this capability at Athabasca Glacier in Alberta, CA, and were able to successfully achieve a controlled 1.5m descent into a moulin where flowing water was actively modifying the terrain. This deployment can be seen in figure 7.

We also suggest that the unique kinematic structure and form factor of the EELS robot allows for the NEO stack to adapt to a larger variety of tasks outside of locomotion and navigation. We propose that due to the large number of DOF within the EELS hardware platform, behaviors that transform the redundant DOF into a stable base allow for the remaining DOF to be leveraged as a manipulator, effectively allowing the EELS platform to operate as a mobile manipulator. A useful note for this is that all the algorithms and technology required to enable the vertical mobility motion control capability provided by SCFM directly translate to providing robust and capable manipulation and grasping capabilities for generalized purposes within NEO. We draw the comparison between the task of grasping an arbitrary object to the task of EELS pushing into its surrounding terrain to enable it to sustain its vertical position, where they are fundamental the same problem, but NEO exerts load outward, instead of inward to grasp. This concept of loco-manipulation, where capabilities of locomotion and manipulation are tightly coupled is packaged into the unique platform provided by EELS, where NEO can exploit the properties of the environment in which it is operating to adapt, enabling both mobility and interaction tasks to be performed.

Another approach we take is the leveraging of massively parallel simulation capabilities to generate large quantities of data for the training of reinforcement learning (RL) policies with emergent novel gait patterns. This approach to generalization of RL policies to extensively different conditions through scale and domain randomization has been demonstrated in [8] for quadrupedal locomotion; we utilize the same environment, the Nvidia Isaac Gym, for our work. The simulation environment Nvidia Isaac Gym was chosen due to the ease of simulating many environments in parallel, allowing for faster data set generation for the training of RL policies [9]. A key detail is that we enable the adaptation of common snake-inspired locomotion gaits (e.g. [10]) to arbitrary kinematic configurations of snake-like platforms. This hybrid approach to gait design enabled by curve-fitting RL policies provides intuitive and parametric authority over the gait behavior. Figure 8 describes the process we follow for generation and deployment of this approach. More details on the implementation of this approach can be found in our parallel work [11].

Figure 6 The EELS 1.0 robot traversing over mobility-stressing extreme terrains

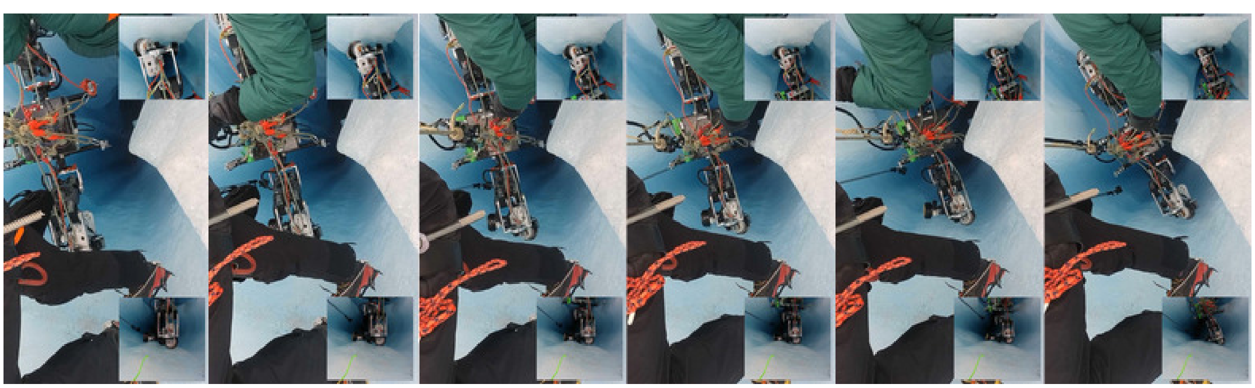

Figure 7 A set of consecutive images displaying a successful descension into a moulin in Athabasca Glacier, AB utilizing the SCFM control strategy

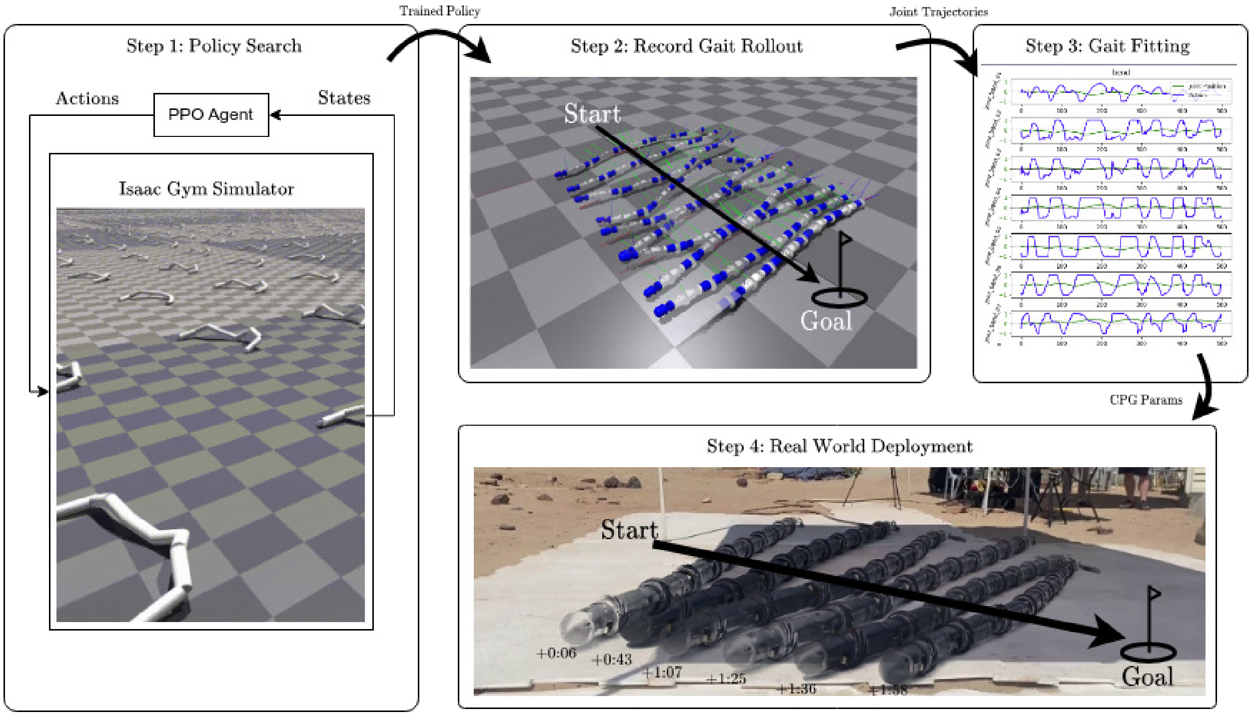

Figure 8 Process for training robust reinforcement learning policies and converting them to hybrid control approaches [11]

Another direction of research to combine reinforcement learning with physics-based approaches to the get best of both worlds is by use of control-barrier functions [12]. Our team is working on combining these techniques using the risk-matrix framework to achieve optimal adaptation to uncertainties while providing safety guarantees which will be covered in a future publication.

C. Risk-Awareness: Integrated Task and Motion Planning under Uncertainty

Following the risk-averse paradigm of Mars Missions in unknown environments like Enceladus would result in a highly conservative behavior possibly to the extent at which the robot might never move to avoid risk. To this end, NEO is equipped with a risk awareness that not only manages the risk vs rewards trade-off but judiciously performs information-gathering actions based on the environment's geometry, the state uncertainty growth rate, and the mission's risk posture.

In current missions, risk is manually evaluated by system engineers on-ground before up-linking the exact command sequence. In NEO, we bring similar reasoning based on the risk matrix that looks at the likelihood and consequence of a failure event to quantify the risk of the mission take appropriate actions to reduce it. Figure 9, shows a sample mission scenario of how the robot leverages versatility and adaptation to reduce risk of falling into a crater.

Classically, this is done in a two-step hierarchical process, one system-level autonomy layer performs the scheduling of discrete activities and then a motion planner/controller plans over the start/goal conditions set by those activities. While this decoupling makes the problem tractable, however, the solutions often suffer from sub-optimatliy. In a highly uncertain environment, this decoupling can also result in failure to find a feasible solution even when one exists. The general formulation of such a problem is a Partially Observable Markov Decision Process (POMDP) in hybrid continuous (robot motion actions) and discrete (activities/behaviors) action spaces which is prohibitively computationally intractable.

NEO addresses these challenges by using a Chance Constrained Task and Motion Planning (CC-TAMP) formulation in a Model Predictive Control (MPC) paradigm. Our comparisons show that this approach outperforms traditional two-stage deterministic formulations and POMDP methods in terms of solution quality and computational efficiency. Figure 10 demonstrates the need for this capability. The scanning behaviors allow the robot to create a map of the terrain

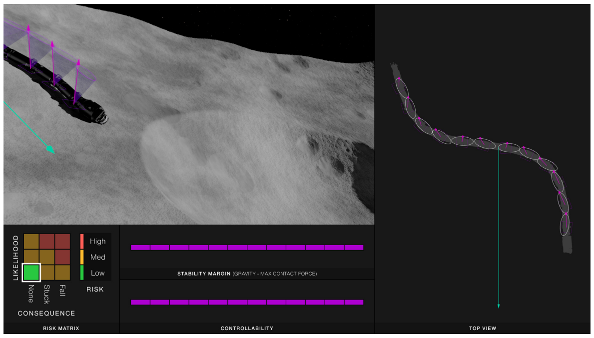

(a) Robot is moving forward using its screw-based locomotion at low since it has sufficient stability margin and control authority to move.

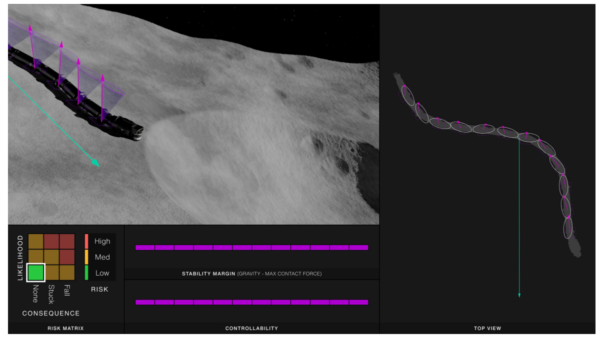

(b) Robot approaches a crater (negative obstacle) i.e. is not detected by the perception system.

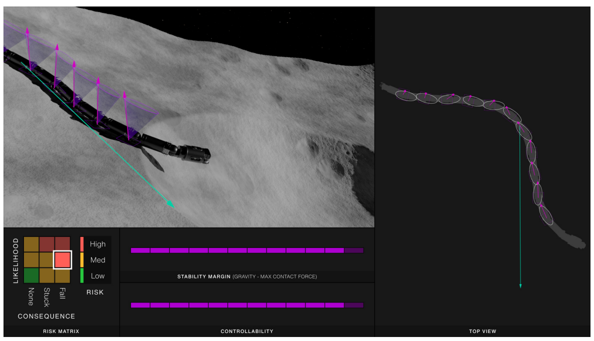

(c) Robot loses contact of the first module with the terrain and detects that using proprioceptive sensing and infers that the stability margin to resist the forces of gravity is has gone done.

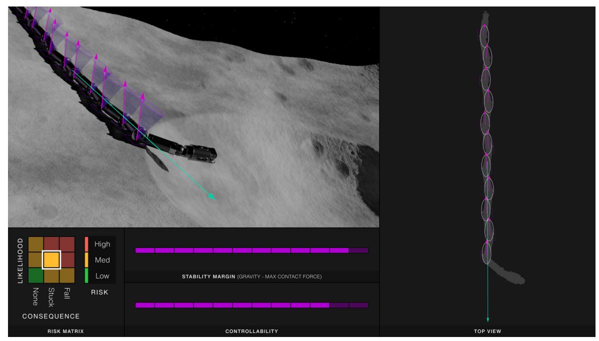

(d) Robot adapts its shape/pose to gain a higher stability margin which results in loss of control authority for the screw-based mobility. Note that the consequence of this is way lower than falling down the crater. This allows the robot to gather more sensor data and time to compute the next risk-free action.

Figure 9 NEO's On-board Risk Matrix-based Adaptation in a Lunar Crater Navigation Mission Scenario.

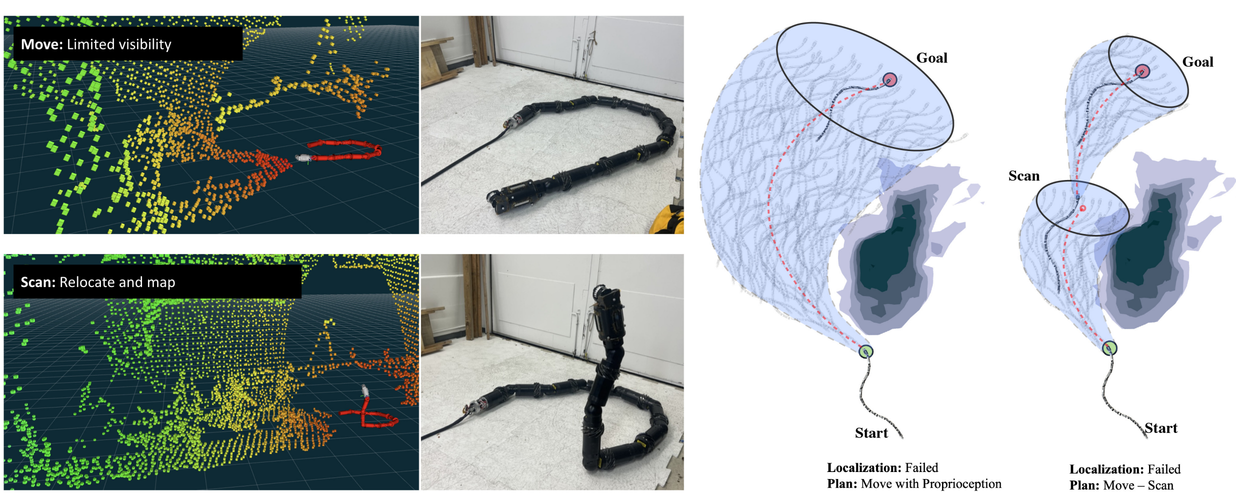

Figure 10 NEO's Risk-aware Integrated Task and Motion Planning Framework in a Moulin Avoidance Mission Scenario. Joint reasoning about moving and scanning activities along with risk-awareness allows the robot to come up with more optimal paths while satisfying safety constraints.

by lifting its head up but the moving can be done purely based on proprioception. The uncertainty in the localization of the robot grows as a function of the distance traveled using proprioception only. Hence, the decision of "when to scan" is highly dependent on the motion, similarly the decision of "how to move" is highly dependent on the scheduled scan time. Our proposed solution jointly solves both these problems in a single integrated optimization problem that is formulated as Mixed-integer Linear Program (MILP). Further details and quantitative results can be found in [13], [14].

III. Process and Lessons Learned

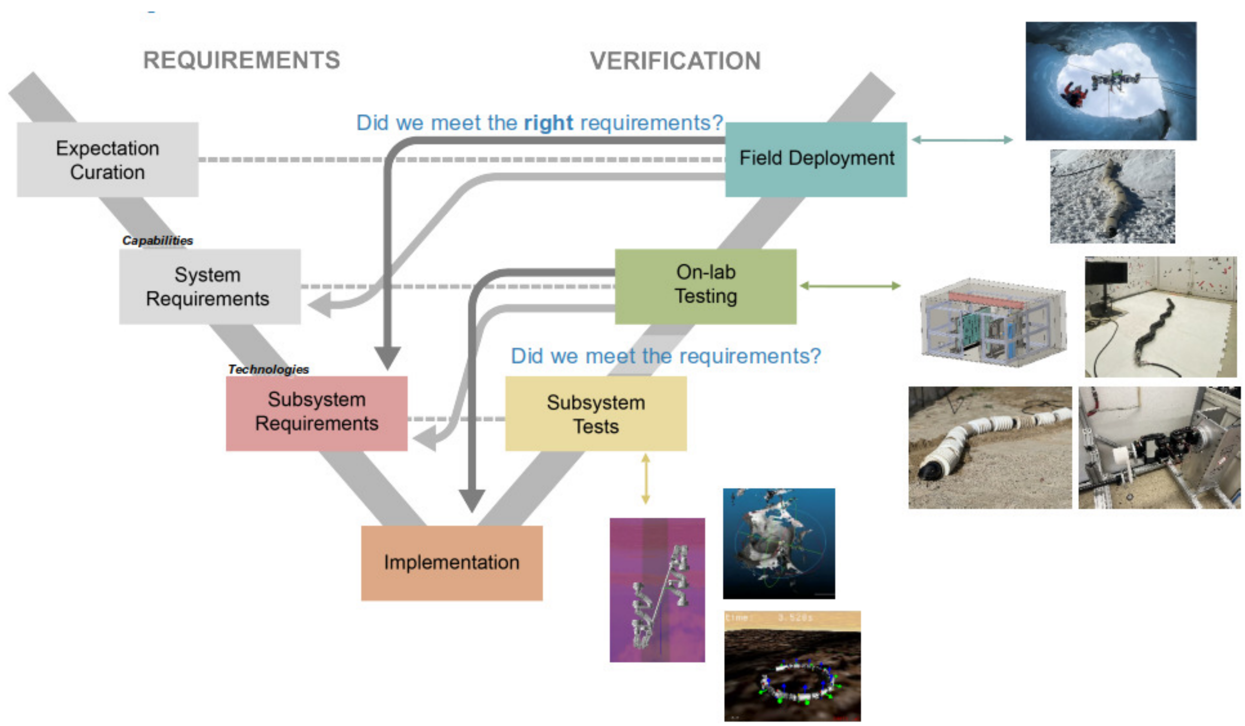

Figure 11 Overview of the NEO agile development, verification, and validation process.

In a similar vein to the paradigm proposed for autonomous exploration of unknown worlds, we apply the same principles to our processes for technology development and iteration of system requirements to enable such a platform. The concept of adaptability is deeply rooted in our methodologies of agile project planning and execution, which emphasizes flexibility, rapid iteration, and continuous improvement. We emphasize these as key components that have been the cornerstone of many successful pieces of technology development over the past three decades, specifically within high-quality, robust, software products.

In contrast to many traditional approaches to mission concept development, technology demonstration, and flight readiness preparation, where each phase often takes many months and commonly years of development and evaluation before continuing, we take on a shorter-horizon approach. This approach is motivated by the recognition that the complexity of a system required to be capable of achieving mission requirements, such as the exploration of an unknown world, contain a design space so vast it is inefficient and often infeasible to analyze all possible outcomes ahead of time. As such we instead establish our mission-level requirements that fully capture the end capabilities that must be present for a successful demonstration (L1 and L2 requirements), utilizing those to motivate and inform our selection and iteration of sub-system and component-level technology requirements (L3 and L4 requirements). Figure 11 displays our approach to continuous development, evaluation, and iteration with respect to our component and system-level technology and requirements, where subsystem-level tests are performed on the timescale of days, and on-lab level tests are performed on the timescale of weeks. Focusing on the speed of iteration and development momentum enables us to discover requirements and limitations that are challenging to forecast ahead of time, especially when breaking ground on a mission concept that has never been previously approached.

A key enabler to our approach to continuous integration, testing, and deployment, is reducing the barrier to performing experiments. A major component of this is the leveraging of simulation tools. Specifically, we utilize the existing Dynamics And Real-Time Simulation (DARTS) environment for simulating the multi-body dynamics of our systems [15]. DARTS provides useful tools for the modeling of contact interaction and interfaces for repeatable experiment setup and execution. Leveraging simulation-in-the-loop (SITL) testing is orders of magnitude more time and resource-efficient than performing a hardware test, and while there are some physical phenomena that prove challenging to properly capture within a simulation environment, many common software bugs can be caught with a single developer in only a few seconds, ensuring that can be caught before an integrated hardware test, is. It is important to note that the interface abstraction between DARTS and NEO is intentionally designed to be the same as between NEO and physical hardware platforms; this lowers developer friction when rapidly iterating on algorithms to evaluate their performance in both environments. Figure 12 displays the EELS 1.0 and EELS 1.5 robots being simulated within DARTS in various terrains. Specific technical details regarding the algorithms and architecture of the DARTS environment utilized are reported in [16]

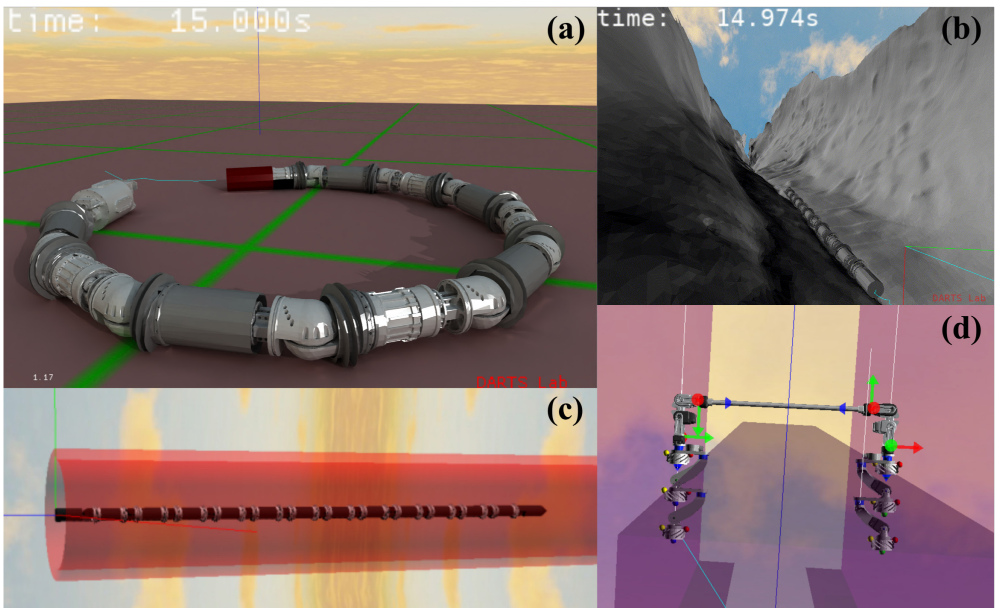

Figure 12 Different terrains within the DARTS environment, with EELS 1.0 and EELS 1.5 models being tested [16]

Another important portion of this effort to strive for speed of iteration is to inform changes that need to occur with longer lead-time modifications, such as hardware changes to the EELS platforms. Traditional hardware development cycles often take months or even years depending on the platform scope, but by maintaining a tight feedback loop between algorithm development, system performance, and hardware design, we were able to gain new hardware capabilities on the order of days and weeks instead of months or years. The speed at which we design experiments, execute them, and analyze results directly leads to our ability to inform hardware modifications.

To allow rapid iteration of software, we followed the architecture shown in Figure 13. The key idea is to impose minimal structure in the high-level architecture to keep things flexible and implement the core capabilities/algorithms as middleware-agnostic core libraries (e.g. python modules or cmake libraries). These core libraries are used to instantiate controllers/planners/etc. via managers who have a dependency on the middleware and are responsible for piping the relevant data in and out of the core libraries.

A key portion of hardware that was iterated on through this process was gait design for vertical mobility, where the procedure typically begins through analysis in simulation environments, where various gait patterns can be tested and analyzed without the need for physical prototypes. While some characteristics were relatively simple to downselect from with pure kinematic analysis or simulation tests, other aspects, such as the intricacies of screw-ice interaction were not feasible to capture within a simulation where the ability to run close to real-time is critical, and as such, required physical testing to inform what gait configuration would best accommodate the system-level requirements. Once a baseline is established, the first prototype is constructed and continuously tested in small-scale component tests, leading to larger

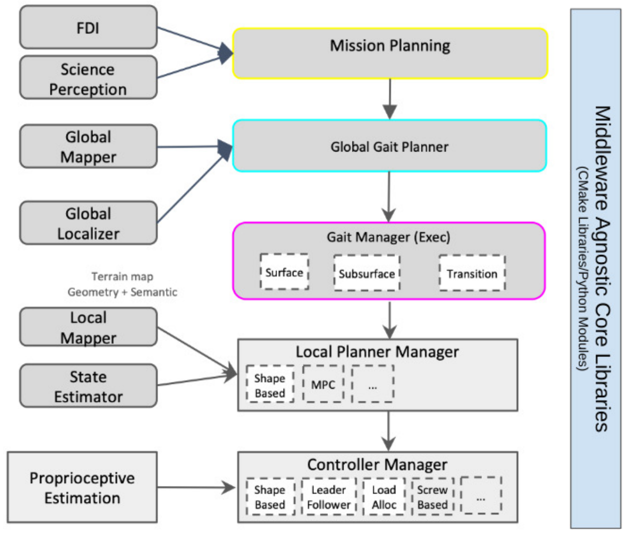

Figure 13 Software Architecture to enable Agile Autonomy Development by implementing core capabilities as middleware-agnostic core libraries.

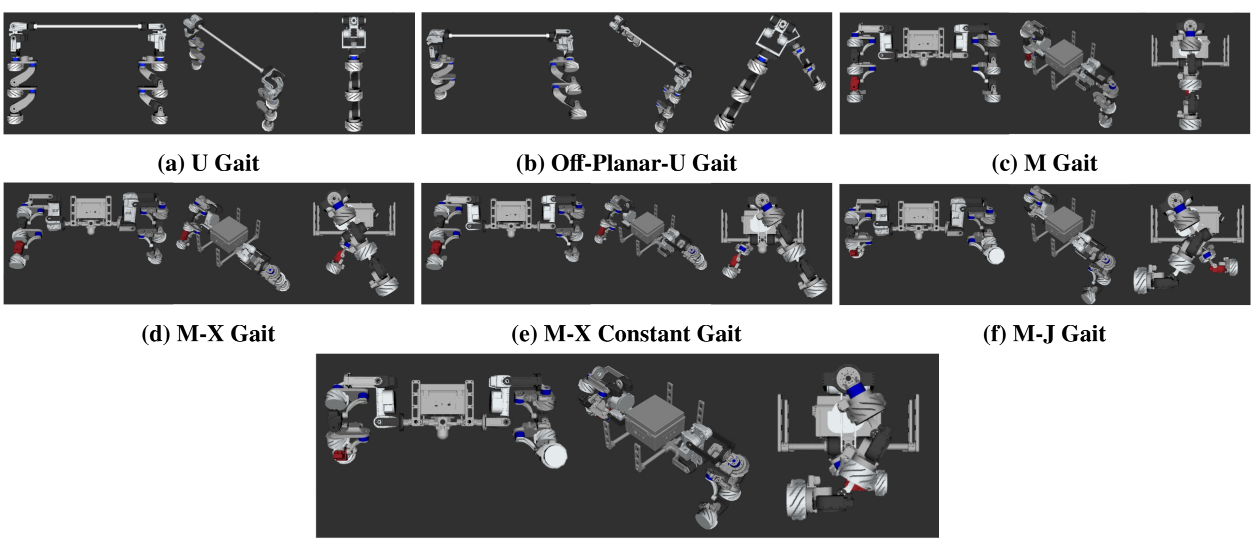

(a) U Gait

(b) Off-Planar-U Gait

(c) M Gait

(d) M-X Gait

(e) M-X Constant Gait

(f) M-J Gait

(g) M-J Constant Gait

Figure 14 Gaits iterated through during the development of the EELS 1.5 platform [7] (For each image: Left: Front View; Center: Isometric View; Right: Side View)

integration tests. Over the course of 3 months, we continuously designed experiments, made hardware modifications, and evaluated system-level performance with respect to vertical mobility to make numerous improvements to the gait. Figure 14 displays the numerous gaits iterated over for the EELS 1.5 platform, before arriving at the final configuration. More details on the gait design process can be found in [7]. Similarly, the experiments and technical analysis related to screw-ice interaction and active-skin mobility are explored in depth in our companion paper [3].

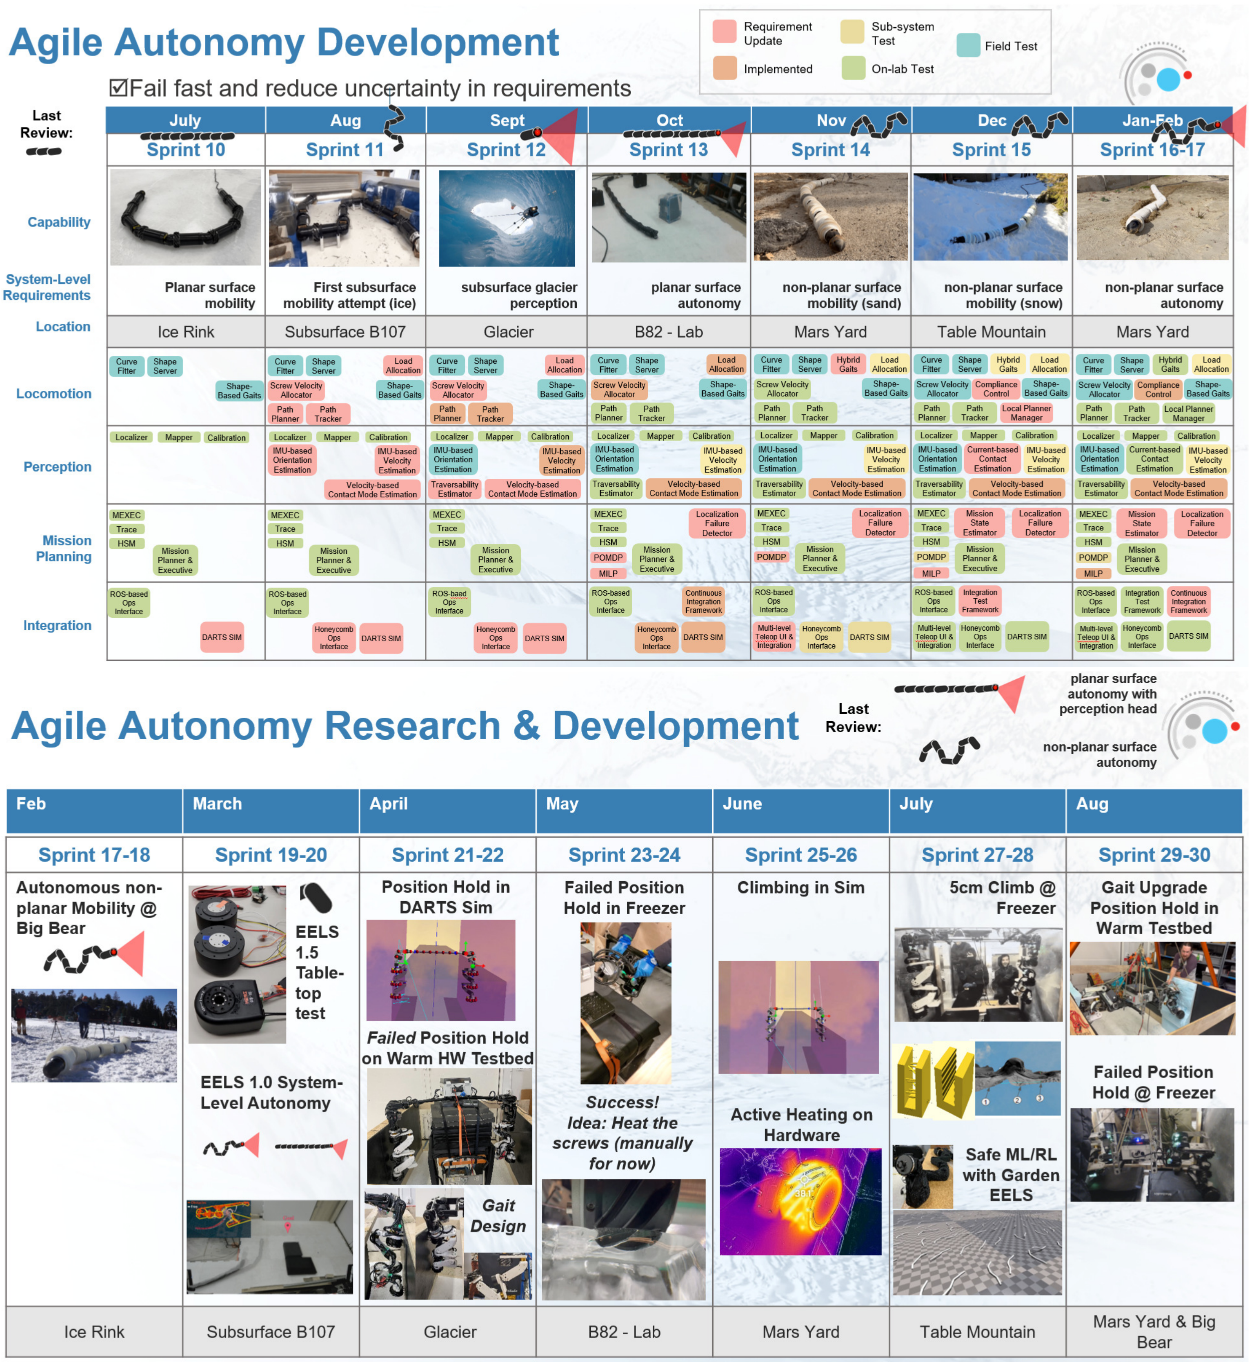

Figure 15 Example sprint planning boards from two progress reviews.

Such agile development practices have been followed in the IT industry and software startups for over a decade, our key innovation was to extend these practices to robotics development which involves iterations on hardware as well as software with feedback loops from one subsystem to the other. Figure 15 shows that the practices proposed in the section indeed led to high-speed agile development on hardware and software sub-systems leading up to a successful field deployment at a natural glacier.

IV. Conclusion

Extensive hardware evaluations on synthetic ice indoors, natural icy and sandy terrains outdoors, and vertical terrain features such as moulins within glaciers confirm that our strategies are not only resilient, adaptable, and risk-aware but also capable of handling mobility-stressing elements in extreme environments. These results underscore the potential of the EELS platform in paving the way for future exploratory missions in space. Exploring unknown, extreme environments necessitates a departure from conventional methodologies. Robots must be resilient, adaptive, and risk-aware. The NEO autonomy framework brings together a suite of technologies to equip robots like EELS with the capabilities necessary for robust exploration of such environments. The NEO-equipped EELS robot promises to spearhead a new space exploration paradigm of diving into the unknown, leading us into the future of space exploration.

Our ongoing research will continue to refine and enhance these capabilities. We foresee NEO enabling the exploration of even more extreme and challenging environments in the cosmos, paving the way for transformative discoveries in the search for extant life beyond Earth.

Acknowledgments

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

References

[1] Ono, M., Ambrose, E., Archanian, A., Balentine, D., Barchowsky, A., Botteon, K., Bowkett, J., Brown, K., Caballero, M., Cable, M., Carpenter, K., Chodas, M., Daca, A., Dacey, S., Daddi, G., de los Reyes, B., Dreger, H., Drevinskas, T., Etheredge, R., Feldman, J., Gardner, A., Gavrilov, P., Georgiev, N., Gildner, M., Goddu, A., Hasseler, T., Hockman, B., Holland, J., Ingham, M., Jain, A., Jasour, A., Jasper, J., Jin, C., Jones, B., Khanum, M., Kornfeld, R., Lopez, C., Loret de Mola Lemus, D., Malaska, M., Mark, G., Marteau, E., Melikyan, H., Morrell, B., Naish, J., Nash, J., Nakka, Y., Nuernberger, B., Pailevanian, T., Paster, D., Paton, M., Peticco, M., Proenca, P., Reid, W., Rieber, R., Robinson, M., Rodriguez, J., Royce, R., Shiraishi, L., Spilker, L., Strub, M., Swan, R. M., Talbot, W., Tempest, M., Thakker, R., Tosi, P., Tran, T., Vaquero, T. S., Veismann, M., Weber, J., Wright, M., Yearicks, S., and Zhang, J., "To Boldly Go Where No Robots Have Gone Before – Part 1: Adaptive EELS Robot to Spearhead a New Space Exploration Paradigm of Diving into the Unknowns," *AIAA SciTech*, 2024 Under review.

[2] Gildner, M., Georgiev, N., Ambrose, E., Pailevanian, T., Archanian, A., Melikyan, H., Loret de Mola Lemu, L., and Ono, M., "To Boldly Go Where No Robots Have Gone Before – Part 2: The Versatile Mobility of the EELS Robot for Robustly Exploring Unknown Environments," *AIAA SciTech*, 2024 Under review.

[3] Marteau, E., Tosi, L. P., Veismann, M., Gavrilov, P., Hockmann, B., Ono, M., Khanum, M., Maisha Abdelrahim, and Marvi, H., "To Boldly Go Where No Robots Have Gone Before – Part 3: The Screw Mobility System of EELS for Robust Surface and Subsurface Mobility on Highly Unknown Terrains," *AIAA SciTech*, 2024 Under review.

[4] Talbot, W., Nash, J., Paton, M., Ambrose, E., Metz, B., Thakker, R., Etheredge, R., Robinson, M., Ono, M., and Ila, V., "Analysis of ICP Covariance Estimation Methods in Perceptually Degraded Environments for the EELS Mission Concept," *International Conference on Intelligent Robots and Systems (IROS)*, 2023, under review.

[5] Santamaria-Navarro, A., Thakker, R., Fan, D. D., Morrell, B., and Agha-mohammadi, A.-a., "Towards resilient autonomous navigation of drones," *Robotics Research: The 19th International Symposium ISRR*, Springer, 2022, pp. 922–937.

[6] Lew, T., Emmei, T., Fan, D. D., Bartlett, T., Santamaria-Navarro, A., Thakker, R., and Agha-mohammadi, A.-a., "Contact inertial odometry: collisions are your friends," *Robotics Research: The 19th International Symposium ISRR*, Springer, 2022, pp. 938–958.

[7] Thakker, R., Jones, B., Nakka, Y. K., Aghli, S., Zade, H., Strub, M. P., Daddi, G., Royce, R., Paton, M., Swan, M., Gildner, M., Vaquero, T., Tosi, P., Veismann, M., Gavrilov, P., Marteau, E., Bowkett, J., de Mola Lemus, D. L., Hockman, B., Orekhov, A., Hasseler, T., Leake, C., Nuernberger, B., Proença, P. F., Reid, W., Talbot, W., Georgiev, N., Pailevanian, T., Archanian, A., Ambrose, E., Jasper, J., Etheredge, R., Roman, C., Levine, D. S., Otsu, K., Melikyan, H., Rieber, R., Carpenter, K., Nash, J., Jain, A., Shiraishi, L., Agha-mohammadi, A.-a., Travers, M., Choset, H., Burdick, J., and Ono, M., "Loco-Manipulation for Snake Robot Climbing through Simultaneous Shape-Contact-Force-Motion Control," *IEEE Transactions on Robotics (T-RO 2024)*, 2024 Manuscript in preparation.

[8] Rudin, N., Hoeller, D., Reist, P., and Hutter, M., "Learning to walk in minutes using massively parallel deep reinforcement learning," *Conference on Robot Learning*, PMLR, 2022, pp. 91–100.

[9] Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., and State, G., "Isaac Gym: High Performance GPU Based Physics Simulation For Robot Learning," *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. URL `https://openreview.net/forum?id=fgFBtYgJQX_`.

[10] Hirose, S., "Biologically Inspired Robots: Snake-Like Locomotors and Manipulators by Shigeo Hirose Oxford University Press, Oxford, 1993, 220 pages, incl. index (£40)," *Robotica (Cambridge. Print)*, 1994. URL `https://www.semanticscholar.org/paper/Biologically-Inspired-Robots%3A-Snake-Like-Locomotors-Owen/b84dfcd1184c570118c9ab77c335bc15863e531f`.

[11] Naish, J., Thakker, R., Rodriguez, J., Daddi, G., Orekhov, A., Zhang, J., Royce, R., Jones, B., Travers, M., Choset, H., , and Ono, M., "Reinforcement Learning-Driven Parametric Curve Fitting for Snake Robot Gait Design," *Learning for Dynamics and Control (L4DC 2024)*, 2024 Manuscript in preparation.

[12] Fan, D. D., Nguyen, J., Thakker, R., Alatur, N., Agha-mohammadi, A.-a., and Theodorou, E. A., "Bayesian learning-based adaptive control for safety critical systems," *2020 IEEE international conference on robotics and automation (ICRA)*, IEEE, 2020, pp. 4093–4099.

[13] Daddi, G., Vaquero, T., Jasour, A., Paton, M., Strub, M., Ingham, M., Corpino, S., Ono, M., and Thakker, R., "Risk-aware Integrated Task and Motion Planning for Resilient Mobility under Localization Failures with Chance Constraints," *IEEE Robotics and Automation Letters*, 2024.

[14] Vaquero, T. S., Daddi, G., Thakker, R., Paton, M., Jasour, A., Strub, M. P., Swan, M., Royce, R., Gildner, M., Tosi, P., Veismann, M., Gavrilov, P., Marteau, E., Bowkett, J., de Mola Lemus, D. L., Nakka, Y. K., Hockman, B., Orekhov, A., Hasseler, T., Leake, C., Nuernberger, B., Proença, P. F., Reid, W., Talbot, W., Georgiev, N., Pailevanian, T., Archanian, A., Ambrose, E., Jasper, J., Etheredge, R., Roman, C., Levine, D. S., Otsu, K., Yearicks, S., Melikyan, H., Rieber, R., Carpenter, K., Nash, J., Jain, A., Shiraishi, L., Robinson, M., Travers, M., Choset, H., Burdick, J., Gardner, A., Cable, M., Ingham, M., and Ono, M., "Autonomous Task and Motion Planning for Ice Worlds Exploration," *Science Robotics*, 2024.

[15] Jain, A., "DARTS - Multibody Modeling, Simulation and Analysis Software," *Multibody Dynamics 2019*, edited by A. Kecskeméthy and F. Geu Flores, Springer International Publishing, Cham, 2020, pp. 433–441.

[16] Hasseler, T., Leake, C., Gaut, A., Elmquist, A., Thakker, R., Swan, R. M., Ono, M., and Jain, A., "EELS-DARTS: A Versatile Multibody Dynamics Simulator for Autonomous Snake-Like Robots," *International Conference on Multibody System Dynamics*, 2024 Manuscript in preparation.