

Regular Article

NeBula: TEAM CoSTAR's Robotic Autonomy Solution that Won Phase II of DARPA Subterranean Challenge

Ali Agha¹, Kyohei Otsu¹, Benjamin Morrell¹, David D. Fan¹, Rohan Thakker¹, Angel Santamaria-Navarro², Sung-Kyun Kim¹, Amanda Bouman³, Xianmei Lei¹, Jeffrey Edlund¹, Muhammad Fadhil Ginting¹, Kamak Ebadi¹, Matthew Anderson¹, Torkom Pailevanian¹, Edward Terry¹, Michael Wolf¹, Andrea Tagliabue⁴, Tiago Stegun Vaquero¹, Matteo Palieri¹¹, Scott Tepsuporn¹, Yun Chang⁴, Arash Kalantari¹, Fernando Chavez¹, Brett Lopez¹, Nobuhiro Funabiki¹, Gregory Miles¹, Thomas Touma¹, Alessandro Buscicchio^{1,11}, Jesus Tordesillas¹, Nikhilesh Alatur¹, Jeremy Nash¹, William Walsh¹, Sunggoo Jung⁸, Hanseob Lee⁸, Christoforos Kanellakis⁹, John Mayo¹, Scott Harper¹, Marcel Kaufmann¹, Anushri Dixit³, Gustavo J. Correa¹⁰, Carlyn Lee¹, Jay Gao¹, Gene Merewether¹, Jairo Maldonado-Contreras¹, Gautam Salhotra⁵, Maira Saboia Da Silva¹, Benjamin Ramtoula¹, Seyed Fakoorian¹, Alexander Hatteland¹, Taeyeon Kim¹, Tara Bartlett¹, Alex Stephens¹, Leon Kim¹, Chuck Bergh¹, Eric Heiden⁵, Thomas Lew¹, Abhishek Cauligi¹, Tristan Heywood¹, Andrew Kramer⁶, Henry A. Leopold¹, Hov Melikyan¹, Hyungho Chris Choi⁸, Shreyansh Daftry¹, Olivier Toupet¹, Inhwan Wee⁸, Abhishek Thakur¹, Micah Feras¹, Giovanni Beltrame⁷, George Nikolakopoulos⁹, David Shim⁸, Luca Carlone⁴ and Joel Burdick³

¹NASA Jet Propulsion Laboratory, California Institute of Technology

²Institut de Robòtica i Informàtica Industrial, CSIC-UPC

³California Institute of Technology

⁴Massachusetts Institute of Technology

⁵University of Southern California

⁶University of Colorado Boulder

⁷École Polytechnique de Montréal

⁸Korea Advanced Institute of Science and Technology

Received: 15 February 2021; revised: 18 October 2021; accepted: 3 January 2022; published: 3 July 2022.

Correspondence: Ali Agha, NASA Jet Propulsion Laboratory, California Institute of Technology,

Email: aliagha@jpl.nasa.gov

This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Copyright © 2022 Agha, Otsu, Morrell, Fan, Thakker, Santamaria-Navarro, Kim, Bouman, Lei, Edlund, Ginting, Ebadi, Anderson, Pailevanian, Terry, Wolf, Tagliabue, Vaquero, Palieri, Tepsuporn, Chang, Kalantari, Chavez, Lopez, Funabiki, Miles, Touma, Buscicchio, Tordesillas, Alatur, Nash, Walsh, Jung, Lee, Kanellakis, Mayo, Harper, Kaufmann, Dixit, Correa, Lee, Gao, Merewether, Maldonado-Contreras, Salhotra, Da Silva, Ramtoula, Fakoorian, Hatteland, Kim, Bartlett, Stephens, Kim, Bergh, Heiden, Lew, Cauligi, Heywood, Kramer, Leopold, Melikyan, Choi, Daftry, Toupet, Wee, Thakur, Feras, Beltrame, Nikolakopoulos, Shim, Carlone and Burdick

DOI: <https://doi.org/10.55417/fr.2022047>

⁹Luleå University of Technology

¹⁰University of California, Riverside

¹¹Polytechnic University of Bari

Abstract: This paper presents and discusses algorithms, hardware, and software architecture developed by the TEAM CoSTAR (Collaborative SubTerranean Autonomous Robots), competing in the DARPA Subterranean Challenge. Specifically, it presents the techniques utilized within the Tunnel (2019) and Urban (2020) competitions, where CoSTAR achieved second and first place, respectively. We also discuss CoSTAR's demonstrations in Martian-analog surface and subsurface (lava tubes) exploration. The paper introduces our autonomy solution, referred to as NeBula (Networked Belief-aware Perceptual Autonomy). NeBula is an uncertainty-aware framework that aims at enabling resilient and modular autonomy solutions by performing reasoning and decision making in the belief space (space of probability distributions over the robot and world states). We discuss various components of the NeBula framework, including (i) geometric and semantic environment mapping, (ii) a multi-modal positioning system, (iii) traversability analysis and local planning, (iv) global motion planning and exploration behavior, (v) risk-aware mission planning, (vi) networking and decentralized reasoning, and (vii) learning-enabled adaptation. We discuss the performance of NeBula on several robot types (e.g., wheeled, legged, flying), in various environments. We discuss the specific results and lessons learned from fielding this solution in the challenging courses of the DARPA Subterranean Challenge competition.

Keywords: aerial robotics, exploration, extreme environments, GPS-denied operation, mapping, motion planning, subterranean robotics, legged robots, teleoperation, wheeled robots

1. Introduction

Robotics and artificial intelligence (AI) are transforming our lives, with a growing number of robotic applications ranging from self-driving cars (Yurtsever et al., 2020), search and rescue (Balta et al., 2017), healthcare (Qin et al., 2020), and humanitarian missions (Santana et al., 2007), to robots under water (Kinsey et al., 2006) and robots beyond our home planet on Mars (Sasaki et al., 2020; Bajracharya et al., 2008) and the moon (Ford et al., 2020). Autonomy and AI are empowering these robots to carry out missions autonomously, increasing efficiency with reduced human risk, saving lives, and accomplishing tasks that are often in hazardous environments too dangerous for humans.

Extreme environments: Underground environments are an important example of the type of terrain that imposes a lot of risk for humans, with a wide range of terrestrial and planetary applications. On Earth, autonomous underground exploration is a crucial capability in search and rescue missions after natural disasters, in the mining, oil, and gas industry, and in supporting spelunkers and cave rescue missions. One prominent example is the Tham Luang cave rescue (Figure 1), where the international community aimed at rescuing 13 members of a football team from 4 km inside a partially flooded cave. Drones equipped with thermal cameras have been flown over Tham Luang to detect possible access points, and an underwater robot was deployed to send information back on the water depth and condition of the cave. However, at that time, no technology existed to autonomously reach the people, map the cave, and scan for people deep underground.

Planetary applications: Beyond our home planet, the research community has identified more than 200 lunar and 2000 Martian cave-related features (Cushing, 2012). Caves and subsurface voids, in general, are of utmost importance in space exploration for several reasons. First, their stable, radiation-shielding environment and potential to act as volatile traps make them an ideal habitat candidate for future human exploration (Kesner et al., 2007; Titus et al., 2020). Second, planetary cave environments may harbor life due to their shielding from cosmic rays, and if there is life beyond Earth, deep planetary caves are one of the most likely places to find it. Third, exploring




Figure 1. Tham Luang cave rescue mission. Figures from Fedschun (2018) and Vejpongsa (2018).

caves provides an unprecedented opportunity for scientists to study planetary volcanic processes and the geological history of planetary bodies. These reasons, among many others, have made subsurface exploration one of the main next frontiers for space exploration (Stamenković et al., 2019; Touma et al., 2020).

While autonomy and AI technologies are growing fast, challenges still remain for operations in extreme environments. Technical challenges include perceptual degradation (darkness, obscurants, self-similarity, limited textures) in wholly unknown and unstructured environments, extreme terrain that tests the limits of traversability, mission execution under constrained resources, and high-risk operations where robot failure or component degradation is a real possibility. Most challenging of all, however, is the combination of the above features. Further work is needed to push the state of the art to enable systems that can robustly and consistently address these challenges simultaneously.

Contributions: In this paper, we discuss the NeBula autonomy solution and Team CoSTAR’s contributions towards addressing some of the challenges in robotic exploration of unknown extreme surface and subsurface environments. We discuss these technologies in the context of the DARPA Subterranean Challenge (DARPA, 2018a), where Team CoSTAR won the Urban competition and ranked second in the Tunnel competition. The videos in CoSTAR Team (2020a,b,c,d, 2021) depict some highlights of these runs. As we will discuss in Section 2, this competition pushes the state-of-the-art boundaries in extreme environment exploration in mobility, perception, autonomy, and networking. Specifically, we will discuss Team CoSTAR’s contributions in advancing the autonomy in the following fronts:

1. Resilient, modular, and uncertainty-aware autonomy architecture (Section 4),
2. State estimation in perceptually degraded environments (Section 5),
3. Large-scale positioning and 3D mapping (Section 6),
4. Semantic understanding and artifact detection (Section 7),
5. Risk-aware traversability analysis (Section 8),
6. Global motion planning and coverage/search behaviors (Section 9),
7. Multi-robot networking (Section 10),
8. Mission planning and system health management (Section 11), and
9. Mobility systems and hardware integration (Section 12).

2. DARPA Subterranean Challenge

The DARPA Subterranean or “SubT” Challenge (DARPA, 2018a) is a robotic competition that seeks novel approaches to rapidly map, navigate, and search underground environments (Figure 2). The competition spans a period of 3 years. The teams participating in the systems track, develop, and




Figure 2. The three subdomains in the DARPA Subterranean Challenge: tunnel systems, urban underground, and cave networks.

implement physical systems for autonomous traversal, mapping, and search in various subterranean environments, including mines, industrial complexes, and natural caves.

Illustrative scenarios: The primary scenario of interest for the competition is providing autonomous and rapid situational awareness in unknown and challenging subterranean environments. The layout of the environment is unknown, could degrade or change over time (i.e., dynamic terrain), and is either impossible or too high risk to send in personnel. Potential representative scenarios range from planetary cave exploration to rescue efforts in collapsed mines, post-earthquake search and rescue missions in urban underground settings, and cave rescue operations for injured or lost spelunkers. Additional scenarios include missions where teams of autonomous robotic systems are sent to perform rapid search and mapping in support of follow-on operations in advance of service experts, such as astronauts and search-and-rescue personnel. These scenarios present significant challenges and dangers that would preclude employing a human team, such as collapsed and unstable structures or debris, a presence of hazardous materials, lack of ventilation, and potential for smoke and/or fire.

2.1. Competition Rules

Competition structure: Each team has a fixed time window (1 hour) to carry out the mission. Each team deploys its robots to provide rapid situational awareness through mapping of the unknown environment and localization of artifacts (e.g., CO₂ gas source, survivors, electrical boxes). As the systems explore the environment, these situational awareness updates need to be communicated back to a base station, outside the challenging area, in as close to real time as possible. The urgency in completing the course objectives and providing near-real-time situational awareness updates is a consistent focus of the competition.

Scoring: The detailed scoring metrics are discussed in [DARPA \(2021\)](#). At a high level, each team gets 1 point per artifact, if the team can (i) reach, detect, and recognize the artifact, (ii) localize the object in global coordinates with less than 5 m error, and (iii) report the object to outside the course during the 1-hour mission period. There are 20 artifacts distributed throughout the course (see [Figure 3](#)), and each team has only 25 chances to report the artifacts. Hence, teams need to be careful in using their reporting budget. In addition to artifacts, the quality of a 3D map, environment coverage, and the traverse length are quantified as interim measures of success.

High-level rules: There is no prior map of the environment. No team member is allowed to enter the course prior to or during the competition. The mission length is limited to an hour. There is one single human supervisor outside the course, who is allowed to see the data coming from the robot and interact with them, if and when a communication link is established. Communication




Figure 3. The final event will comprise elements from all subdomains in a course to demonstrate the versatility of solutions developed. To motivate the efficient exploration and search of the environment, competitors score points by accurately reporting the type and position of artifacts distributed along the course.

between the operator and robots is typically limited to the areas near the course entrance, due to the scale, complexity, and communication-denied nature of the course.

2.2. Technical Challenge Elements

The competition is intended to push the boundaries of the state of the art and state of the practice across various challenge elements, including austere navigation, degraded sensing, severe communication constraints, terrain obstacles, and endurance limits.

All-terrain traversability: The environment includes multiple levels, loops, dead ends, slip-inducing terrain interfaces, and a range of features and obstacles that challenge a robot's mobility. Examples of terrain elements and obstacles include constrained passages, sharp turns, large drops/climbs, inclines, steps, ladders, mud, sand, and/or water. The environments may also include organic or human-made materials, structured or unstructured clutter, and intact or collapsed structures and debris.

Degraded perception and sensing: The environment includes elements that range from constrained passages to large openings, lighted areas to complete darkness, and wet to dusty conditions. Such environments with limited visibility, difficult and expansive terrain, and/or sparse features can lead to significant localization error and drift over the duration of an extended run. Perception and proprioceptive sensors will need to reliably operate in these low-light, obscured, and/or scattering environments while having the dynamic range to accommodate such varying conditions. Dust, fog, mist, water, and smoke are among the challenging elements.

Constrained communication: Limited line of sight, radio frequency (RF) propagation challenges, and effects of varying geology in subterranean environments impose significant impediments to reliable networking and communications links. Teams in this competition consider innovative approaches to overcome these constraints, including novel combinations of hardware, software, waveforms, protocols, distributed or dispersed concepts, and/or deployment methods.

Endurance and power limits: To succeed in accomplishing the mission goals, teams need to be capable of a team-aggregated endurance of 60–120 minutes. This aggregate endurance requires novel

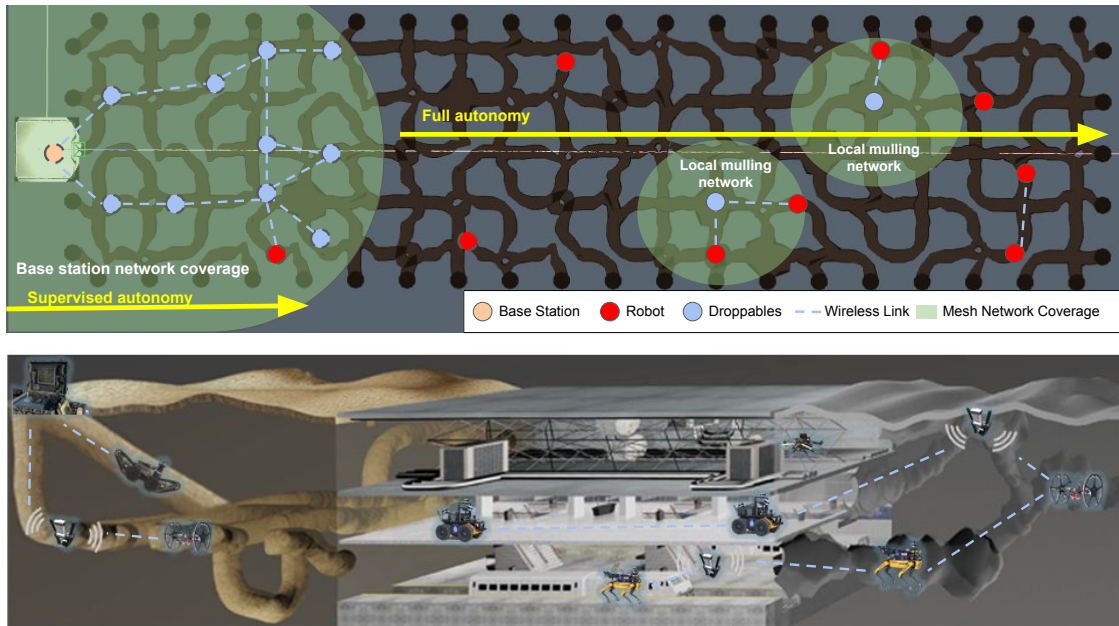


Figure 4. NeBula's concept of operations. Top: Bird's eye view of autonomy in cave. Bottom: Perspective view with our robots and different environments.

deployment concepts, energy-aware planning, heterogeneous agents of varying endurance, energy harvesting or transfer technologies, and/or a combination of various approaches to overcome the various challenge elements.

3. Concept of Operations

In this section, we briefly go over NeBula's concept of operations (ConOps) (Figure 4) for exploring unknown extreme terrains under time constraints. In order to simultaneously address various challenges associated with exploring unknown challenging terrains (Section 2.2), we rely on a team of heterogeneous robots with complementary capabilities in mobility, sensing, and computing.

Robot capabilities: Figure 5 shows the robots we have deployed. The capabilities of these robots drive the ConOps design process. Tables 1, 2, and 3 summarize our heterogeneous robot capabilities from mobility, sensory, and computing perspectives. Our ConOps induces specific mobility-sensor-computer combinations, defining the robots we deploy in the environment to satisfy the mission objectives. The payload capacity of each robot is directly correlated with the sensory capacity; larger payload capacity allows for a larger sensory suite. Also, the energy/battery capacity, and desired endurance on each robot is correlated with their processing capabilities; typically, larger robots are able to carry larger batteries and more powerful computing resources. Some of these mobility-sensor-computer combinations are discussed in Section 12.

ConOps: Our ConOps utilizes a heterogeneous set of platforms (see Table 1 and Figure 5). In the following, we describe several steps of an example illustrative mission ConOps.

1. *Vanguard operations:* As the robots enter the environments, they explore the frontier with vanguard robots with highly capable sensing for mapping and artifact detection.
2. *Mesh network expansion:* As robots start the mission they aim at maintaining communication with the human supervisor by creating and extending a wireless mesh network inside the environments of networking. Ground robots do so by deploying communication pucks like breadcrumbs, and aerial scouts can self-deploy for either comms relays or added sensing. Mission autonomy will decide where and how to deploy these breadcrumbs.

Table 1. Heterogeneous NeBula-Powered Mobility Modes

Robot Type	Deployed In	Energy capacity	Payload capacity	Comm	Speed	Mobility	Endurance
Legged robots	Urban	Mid	Mid	Mid	Mid	Mid	Mid
Hybrid (ground/aerial)	STIX	Low	Low	Low	High	Mid-High	Low-Mid
Wheeled	STIX, Mine, Urban	High	High	High	Low	Low	High
Drones	STIX, Mine	Low	Low	Low	High	High	Low
Tracked	STIX	High	Mid	Mid	Low	Low-Mid	Mid-High
Fast small rovers	Mine	Mid	Mid	Low	High	Low	Mid
Aggregated robot team	All events	Shared/Synergistic	Shared/Synergistic	Shared/Synergistic	Aggregated	Aggregated	Aggregated

Table 2. Heterogeneous NeBula Sensing Modalities (the values are based on our specific ConOps)

Sensors	Exteroceptive									Proprioceptive		
	Lidar	Vision	Radar	Thermal	Sonar	IR Depth	CO2/Gas	Wi-Fi	Sound	Contact/Force	Encoder	IMU
Accuracy	High	Mid	Low	Low	Low	High	Low	Low	Low	Low	Mid	High
Power efficiency	Low	High	High	Mid	Mid	High	High	High	High	High	High	High
Size/weight efficiency	Low	High	High	Low	Mid	High	High	High	Mid	Mid	Mid	High
Range and FOV	Mid	High	Low	High	Low	Low	Low	High	Mid	-	-	-
Dark/fog/smoke/dust	Mid	Low	High	High	Mid	Mid	-	-	-	-	-	-

Table 3. NeBula Processors

Processors	Micro-controllers	Snapdragon	Intel NUC	Nvidia Xavier	AMD
Compute	Low	Low	Mid	High	High
Power consumption	Low	Low	Mid	High	High
Size efficiency	High	High	Mid	Low	Low

**Figure 5.** Team CoSTAR's NeBula-powered robots.

3. *Leaving the mesh network*: The environment is highly communication denied. Due to the large scale, complexity of the environment, and capacity of robots to carry communication nodes, the mesh network reach is typically limited to the parts of the environments in the vicinity of the base station (i.e., environment entrance). Hence, the robots will leave the communication network range soon and, for the most part, carry out a fully autonomous mission.
4. *Autonomous mission*: Robots perform search, mapping, and exploration. Autonomous mission guides them to the rendezvous points to exchange information with each other, or they come back to the mesh network to exchange information with the base station.
5. *Dynamic task allocation*: Robots continue simultaneous frontier exploration. They autonomously monitor
 - (a) the state of the robot team, which includes (i) health, battery, and functionality level of the assets, (ii) robot locations, and (iii) the information value (e.g., the numbers of detected artifacts) on each robot;
 - (b) the state of the world, which includes what robots learn about the environment, e.g., geometric and semantic maps;
 - (c) the state of the mission, which includes (i) the remaining mission time, (ii) margin to the desired mission output, and (iii) acceptable risk thresholds; and
 - (d) the state of communication: (i) network connectivity, throughput, etc., (ii) how long each robot is out of the comm range, and (iii) location of comm nodes.

Given these states, the mission planner will decide to deploy new robots, or retask or reposition active robots in the environment.
6. *Team behaviors*: Vehicles and team formation are configured during the mission. Examples include
 - (a) Return to Base, when a battery swap is needed, optimal, and possible at base;
 - (b) Return to Mesh Network, to ensure the data are communicated, then continue;
 - (c) Explore Frontier, to continue as is, aggressively prioritizing coverage; and
 - (d) Act as a Data Mule, to retrieve data from a vehicle that cannot come back to the mesh network (due to limited battery, health, speed, etc.). Faster and healthier vehicles can act as data mules to carry the information between others agents and the mesh network.
7. *Heterogeneous coverage*: These behaviors continue until the entire course is explored. Due to the heterogeneous capabilities of the robots from mobility, sensing, and computation perspectives, the autonomy might dispatch different robots to the same parts of the course. This is to increase the confidence and coverage in mapping and artifact detection by providing multi-modal information (e.g., thermal, radar, etc.) about the environment elements. For example, the drone might have reached and searched parts of the course but, given its limited sensing capabilities, autonomy will dispatch a ground robot to get a second vote on an artifact before submitting it to the server. All the data are submitted to the server prior to the end of mission time.

4. NeBula Autonomy Architecture

Resiliency is a key requirement to enable a repeatable and consistent robotic autonomy solution in the field. *To enable a resilient autonomy solution, NeBula takes uncertainty into account* to cope with unmodeled and unknown elements during the mission.

Motivation/Insight: An important (if not the most important) cause of the brittleness of today's autonomy solutions is the disjointed design of various subsystems. Traditionally, when designing or advancing the performance of a certain module, the typical assumption is that the rest of the system functions properly and nominally. When it comes to real-world deployment, these assumptions typically break due to the discrepancy between the computational models and real-world models. This introduces uncertainty in the perception, inferences, decision making, and execution, potentially leading to suboptimal behaviors.

Key principle for resilient autonomy architectures: Focused on fielding autonomy in challenging environments, NeBula is built on the fundamental principle that “To achieve resilient

autonomy, the decision-making, inference, and perception modules must be reciprocal and tightly co-designed.” This implies reasoning over joint probability distributions across various component-level states as opposed to marginal distributions over a set of disjoint system states. Contrary to the typical sense \rightarrow infer \rightarrow plan \rightarrow act sequence in autonomy solutions, NeBula architecture is built on a plan-to-(sense \rightarrow infer \rightarrow act) loop, where the planner dynamically *plans for* the acquisition of sufficient sensory information and *plans for the quality of representation* required to enable resilient and uninterrupted missions within the prescribed mission risk thresholds. NeBula’s joint perception planning is formulated as an uncertainty-aware belief space planning problem. Belief captures the probability distribution over system’s states including the robot pose, environment state, measurements, team coordination state, health state, communication state, etc. Planning over joint beliefs and taking cross-component uncertainty into account (which describes the interaction of connected modules), NeBula allows for each module to be not only robust to uncertainties within its own subsystem but also resilient to uncertainties in the integration process, resulted from imperfections and off-nominal performance of connected modules.

Illustrative example (perception-aware planning under uncertainty): Simultaneous localization and mapping (SLAM) is a fundamental problem in robotics that aims at simultaneously solving the localization problem (“where is the robot?”) and the mapping problem (“what does the environment look like?”). This is a very well-studied problem and it is well known that solving SLAM (i.e., incorporating joint probability distributions between localization states and environment states) typically leads to optimal and resilient inference, whereas solving localization and mapping separately and putting their solutions together is suboptimal and can lead to a brittle inference system. Analogous to SLAM philosophy, NeBula extends this concept from pure inference to “joint inference and decision-making” (Figure 6). For example, NeBula develops solutions where mapping and planning are solved simultaneously using SMAP (simultaneous mapping and planning) to achieve resilient traversability and risk awareness. Similarly NeBula develops SLAP (simultaneous localization and planning) solutions where localization uncertainty is taken into account in the planning phase using belief space planners. Solving these joint problems typically leads to behaviors where the autonomous system is intelligently planning proactive actions to improve the “inference quality” (e.g., world model or robot model) and reduce uncertainty to the levels necessary to achieve mission goals within the prescribed risk thresholds. This is in contrast to typical solutions where the relationship is one way and the inference module serves the decision-making modules, and decision-making components react to inference output.

Modularity and scalability: In addition to resiliency, NeBula focuses on a modular and scalable framework. This requirement is driven by missions carried out by a team of networked heterogeneous

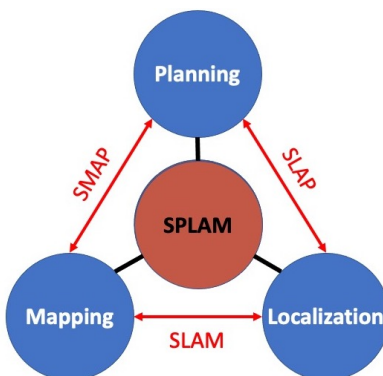


Figure 6. Illustrative example of joint inference and decision making in NeBula’s low-level navigation system. Denoting the state domain of the localization by pose x , mapping by map state m , and planning by policy parameters u , SLAM, SLAP, and SMAP aim at solving for (i.e., estimating or predicting) the joint distributions $p(x, m)$, $p(x, u)$, and $p(m, u)$. The full joint problem, SPLAM (simultaneous planning, localization, and mapping), solves for the probability distribution $p(x, m, u)$.

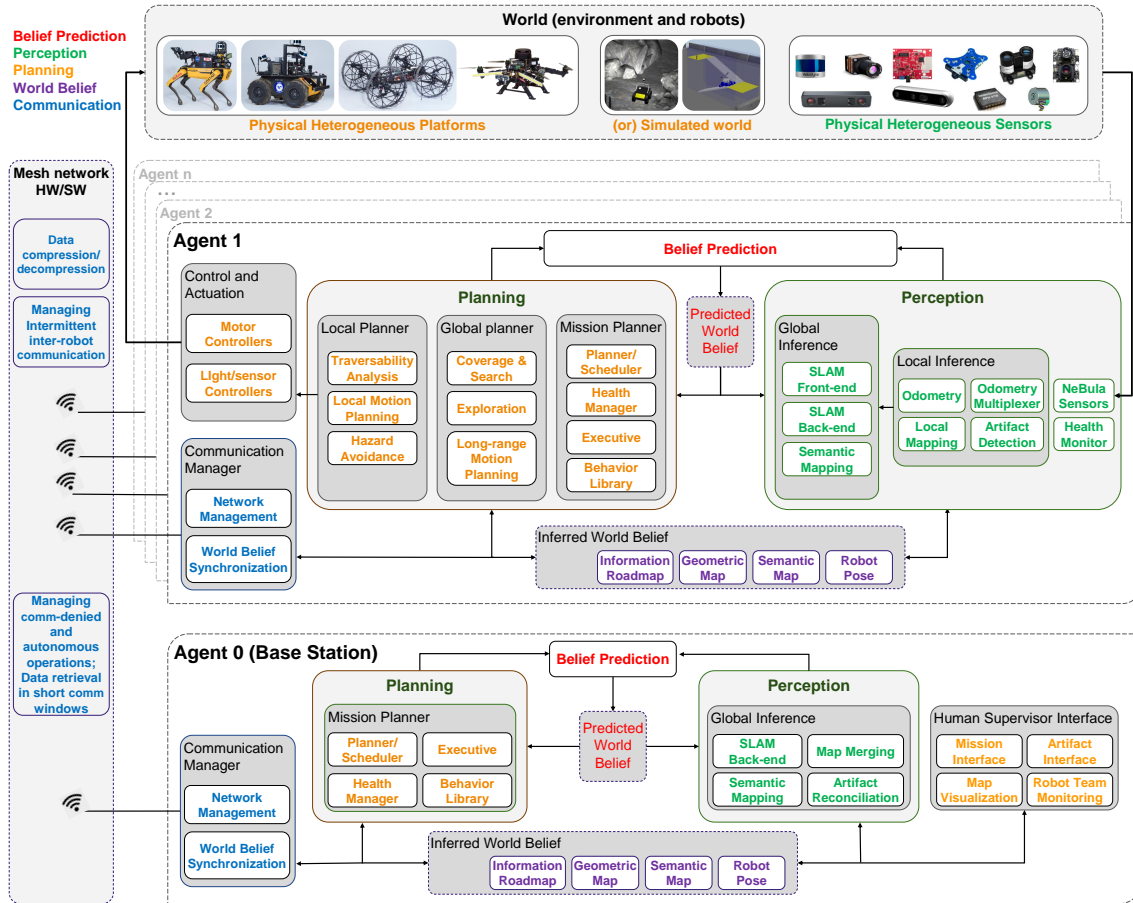


Figure 7. NeBula functional block diagram.

robots. Each robot has different mobility (e.g., wheeled, legged, aerial, hybrid), sensing, and computational capabilities (e.g., Section 12). NeBula provides appropriate abstraction to allow for reusability and agnosticism to the specific robot and hardware. Any low-level hardware-specific modules should be properly isolated to increase the reusability of software. Further, NeBula supports networked systems where agents are intermittently losing and reestablishing communications, and sharing knowledge with each other and with the base station, enabling large-scale environment exploration with a limited number of robots. Each robot has a level of local autonomy to act individually when it is disconnected from the rest of the team.

System architecture: Figure 7 illustrates NeBula’s high-level functional block diagram, and will serve as a visual outline of the sections of this paper. The system is composed of multiple assets: mobile robots, stationary comm nodes, and a base station, each of which owns different computational and sensing capabilities. The base station acts as the central component to collect data from multiple robots and distribute tasks, if and when a communication link to the base station is established. In the absence of the communication links the multi-asset system performs fully autonomously. The fundamental blocks are as follows:

- *Perception* (Sections 5, 6, and 7): Perception modules are responsible to process the sensory data and create a world model belief. The local perception modules (Section 5) provide the odometry and state estimation information needed for local navigation, such as state (pose, velocity) and traversability maps. The global SLAM module, in Section 6, tracks the robot’s position within a globally consistent frame while building a 3D map of the environment. The

semantic understanding and artifact detection module (Section 7) adds semantic information to the map and finds objects of interest from the environment, and in conjunction with the global localization module reports their location.

- *Planning* (Sections 8, 9, and 11): Planning modules will make onboard decisions based on the current world belief. The planning modules are composed of multiple layers. The highest layer is the mission planning module (Section 11), which runs a mission according to its specifications, generates global goals for each robot, and allocates tasks to different robots. The second layer is the global motion planning layer (Section 9), responsible for exploration, search, and coverage behaviors in global scale and large environments. It makes plans to safely move the robot to a goal assigned by the mission planner. It also enables autonomous exploration of the environment in order to increase the knowledge and confidence about the world belief. The third layer is the traversability and local navigation component (Section 8), responsible for analyzing how and with what velocities different terrain elements can be traversed. It quantifies the motion risk, and optimizes/replans local trajectories with high frequency to enable aggressive traversability in obstacle-laden and challenging environments, while ensuring the risk levels stay within the prescribed mission specifications. NeBula abstracts motion models, enabling the planning stack to be robot agnostic and to support heterogeneous mobility platforms.
- *World belief*: This block includes a probabilistic model of the world. It is jointly constructed by perception and planning modules, and enables a tight integration between these modules leading to perception-aware behaviors. World belief extends the traditional state database concept to a belief database, where we maintain probability distributions over various state domains as well as joint probability distributions across multiple domains. It includes belief over robot pose, environment map, mission state, system health, information roadmap, among other state domains. There are multiple variations of the world belief: (i) local to each robot, (ii) shared belief across robots, and (iii) predicted belief to assess future risk and performance to enable making perception-aware and uncertainty-aware decisions. During exploration tasks, robots develop their own local world models based on what they perceived with their limited sensor input. They generate the world model as an abstract representation of spatial and temporal information of their surrounding environment (e.g., maps, hazards) and internal state (e.g., pose, health). This world belief is internally predicted to enable uncertainty-aware decisions and actions based on this predicted model. The shared world model is synchronized among the robots and the base station using asynchronous bidirectional messaging with the publish/subscribe paradigm. The discussion of world belief is distributed across all sections of the paper.
- *Communications* (Section 10): When possible, communication modules synchronize the shared world models across the robots and the base station. To cope with the dynamic and unstable nature of the underlying mesh network, the communication manager is responsible to provide reliable message transfer with buffering, compression, and retransmission. The modules also provide capabilities to maintain a mobile ad hoc network using radio devices. Static communication nodes can be dropped from particular robots to help form a network.
- *Operations* (Section 11): Operations modules aid the human supervisor to effectively monitor the system performance and interact with it if and when communication links are established. One of the main roles is the visualization of complex world belief in a human-recognizable form. In the nominal operation scenarios, the human operator only interacts with the system by updating the world belief, when needed and when possible.

Over time, each layer adapts to the collected data as well as to improvements of models in other layers.

Current implementation status: NeBula is a growing and evolving framework. Its current version (in 2020) has been deployed in several large projects for terrestrial and planetary applications. It has enabled autonomous operations on various vehicle types including (1) wheeled rovers, (2) legged robots, (3) flying multicopters, (4) hybrid aerial/ground vehicles, (5) 1/5th-scale race cars, (6)

tracked vehicles, and (7) full-size passenger vehicles. We refer to Section 12 for a detailed description of our implementation on these platforms.

5. State Estimation

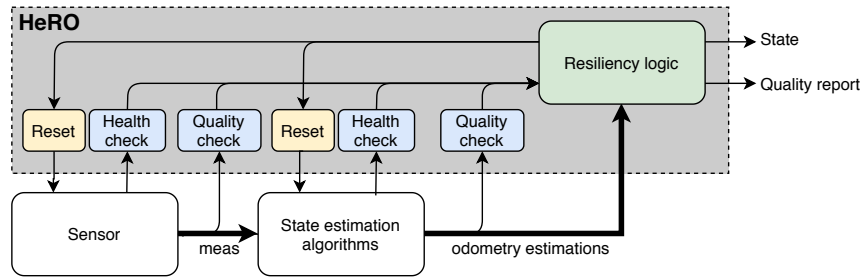
One of the fundamental components of the NeBula architecture, shown in Figure 7, is reliable state estimation under perceptually degraded conditions. This includes environments with large variations in lighting, obscurants (e.g., dust, fog, and smoke), self-similar scenes, reflective surfaces, and featureless/feature-poor surfaces. NeBula relies on a resilient odometry framework that fuses a set of heterogeneous sensors to handle these various challenges. This section briefly describes this odometry solution. (For more details, please see Santamaria-Navarro et al. (2019); Palieri et al. (2020); Kramer et al. (2020); Fakoorian et al. (2020, 2021); Tagliabue et al. (2021); Lew et al. (2019)).

Objective: The objective of the state estimation pipeline is to utilize multi-modal sensing to determine the robot's state, producing resilient, high-rate, and smooth estimates in a probabilistic sense. A key aspect to the proposed approach is assigning each sensor output a *quality measure* that can be used to identify healthy or unhealthy measurements *before* they are fused. We start by introducing our notation: $\vec{p} \in \mathbb{R}^3$ (global position, x, y, z); $\vec{R} \in SO(3)$ (global orientation, which can be described with minimal representation ϕ, θ, ψ); $\vec{v} \in \mathbb{R}^3$ (body linear velocity); $\vec{\omega} \in \mathbb{R}^3$ (body angular velocity); $\vec{a} \in \mathbb{R}^3$ (body linear acceleration); $\vec{\alpha} \in \mathbb{R}^3$ (body angular acceleration); and $Q_i \in \{Good, Bad\}$ (quality of i , with $i \in \{\vec{p}, \vec{R}, \vec{v}, \vec{\omega}, \vec{a}, \vec{\alpha}\}$).

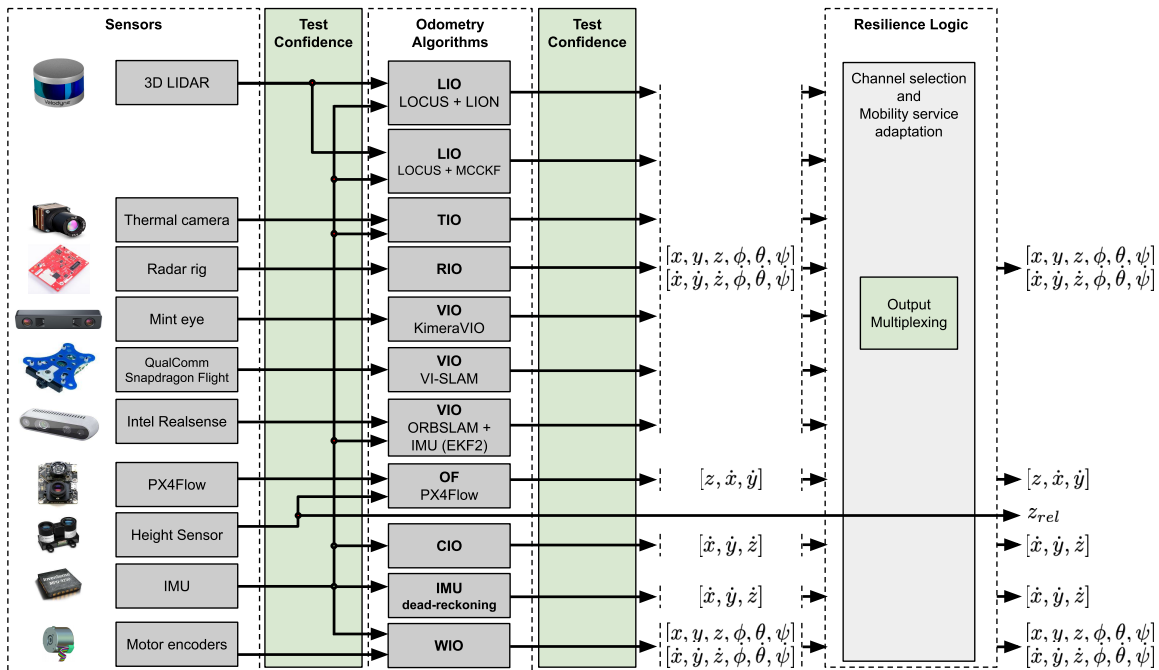
Note that we restricted the quality of the state to binary values although it can be easily generalized to higher resolutions and even continuous representations.

HeRO architecture: The proposed architecture, shown in Figure 8, considers redundancy and heterogeneity in both sensing and estimation algorithms. It is designed to expect and detect failures while adapting the behavior of the system to ensure safety. To this end, we present HeRO, the Heterogeneous and Resilient Odometry Estimator (Santamaria-Navarro et al., 2019): a framework of estimation algorithms running in parallel supervised by a resiliency logic. Resiliency logic has three main functions: (a) perform confidence tests in data quality (measurements and individual estimations) and check health of sensors and algorithms, (b) reinitialize those algorithms that might be malfunctioning, and (c) generate a smooth state estimate by multiplexing the inputs based on their quality. The output of this resiliency logic, which includes a *state quality* measure, is used by the guidance and control system to determine the best mode of operation that ensures safety (see, for instance, the NeBula interconnections between Perception and Planning modules in Figure 7). For example, guidance and control could switch to pure velocity control if the position estimates are unhealthy or issue a stop command if both position and velocity estimates are unreliable.

Heterogeneous complementary algorithms: In addition to selecting heterogeneous sensing modalities, HeRO uses heterogeneous odometry algorithms, e.g., LiDAR-inertial (LIO), visual-inertial (VIO), thermal-inertial (TIO), kinematic-inertial (KIO), contact-inertial (CIO), and RaDAR-inertial (RIO), running in parallel to decrease the probability of a state estimation failure. The key idea behind HeRO is that any single state estimation source can carry errors, either due to failures in sensor measurements, algorithms, or both, but having a complete failure becomes increasingly rare as the number of heterogeneous parallel approaches increases. HeRO is front-end agnostic, accepting various algorithmic solutions and with the ability to incorporate either tightly or loosely coupled approaches. However, to take advantage of all possible mobility modes, there is a need for estimating position, orientation, velocity, and, ideally, acceleration. HeRO is tailored to incorporate a vast variety of estimation algorithms. Figure 8 depicts the main sensor and algorithmic solutions developed and used by Team CoSTAR in the DARPA Subterranean Challenge. Our solution considers software-synchronized sensors (common clock synchronization after initialization), with extrinsic calibrations roughly obtained from the robot model designs and fine tuned used optimization approaches such as Kalibr (camera-IMU) (Furgale et al., 2013; Rehder et al., 2016) or LiDAR-align (LiDAR-LiDAR) (Millane and Taylor, 2019); or by aligning



(a) Main HeRO modules and interconnections



(b) Examples of HeRO sensors, algorithms (e.g., LIO: LiDAR-inertial, VIO: Visual-inertial, TIO: thermal-inertial, RIO: RaDAR-inertial, OF: Optical flow, CIO: contact-inertial or WIO: Wheel-inertial odometries) and observable state components

Figure 8. NeBula's state estimation architecture.

with the robot manufacturer frames. We rely on a variety of in-house and open-source algorithms for sensor fusion. A few examples are as follows: WIO uses an extended Kalman filter (EKF) to fuse the measurements of the wheel encoders and those from an inertial measurement unit (IMU). CIO also takes advantage of an EKF but this time including the modeling of the contacts (Lew et al., 2019). The optical flow (OF) approach, VIO, and thermal imagery fusion (TIO) utilize a combination of open-source and commercial solutions, including PX4Flow (Autopilot, 2021), ORBSLAM (Mur-Artal and Tardós, 2017), Qualcomm VI-SLAM (Qualcomm Technologies, Inc.), and the MiT KimeraVIO (Rosinol et al., 2020), ROTIO (Khattak et al., 2019), among others. The RIO algorithm is our own development presented in Kramer et al. (2020). Finally, the fusion of LiDAR scans with IMU measurements is done by combining our LOCUS scan matching (Palieri et al., 2020) with IMU measurements, either using our Kalman filter variant (AMCKKF, Fakoorian et al., 2021) or a factor graph optimization. As an example, in the following we describe the latter LiDAR-inertial odometer.

5.1. LiDAR-Inertial Odometry Estimation

LiDAR is one of the key sensors offering high range and accuracy. LiDARs also perform well in low-light conditions. Moreover, the combination with an IMU provides essential dynamic information and the registration of the gravity vector, which is of high importance for planning modules (e.g., computing traversability regions). There exist several methods exploiting LiDAR-IMU data fusion achieving remarkable accuracy (Shan et al., 2020; Ye et al., 2019; Hess et al., 2016). However, they do not consider potential failures of the fused sensing modalities, which are likely to be observed in real-world field deployments and can result in catastrophic degradation of the odometry performance if not robustly handled.

To enable reliable operation in extreme settings, our proposed LiDAR-inertial odometry estimation consists of three modules: (i) a LiDAR front end, which provides ego-motion estimation by analyzing LiDAR scans; (ii) an IMU front end based on pre-integration techniques; and (iii) a data fusion module, formulated as a factor graph optimization problem, which fuses the data provided by the front ends and also analyzes the LiDAR observability. This observability analysis is then used by HeRO to take informative decisions about using or not the LIO estimation. These modules are briefly described hereafter.

5.1.1. LiDAR Front End

LOCUS: NeBula's LiDAR-centered front end, referred to as LOCUS (Lidar Odometry for Consistent operations in Uncertain Settings) (Palieri et al., 2020), is a multi-sensor LiDAR-centric solution for high-precision odometry and 3D mapping in real time. LOCUS provides a generalized iterative closest point (GICP) (Segal et al., 2009) based multi-stage scan matching unit equipped with a health-aware sensor integration module for robust fusion of additional sensing modalities in a loosely coupled scheme. The architecture of the proposed system, depicted in Figure 9, is made of three main components: (i) a point-cloud pre-processor, (ii) a scan matching unit, and (iii) a sensor integration module.

The point-cloud pre-processor is responsible for the management of multiple input LiDAR streams (e.g., syncing, motion correction, merging, filtering) to produce a unified 3D data product that can be efficiently processed in the scan matching unit. The scan matching unit then performs a cascaded GICP-based scan-to-scan and scan-to-submap matching operation to estimate the six-degree-of-freedom (6-DOF) motion of the robot between consecutive LiDAR acquisitions.

The sensor integration module is a key component of the system to enable joint optimization of accuracy and robustness. In robots with multi-modal sensing, when available, LOCUS can use

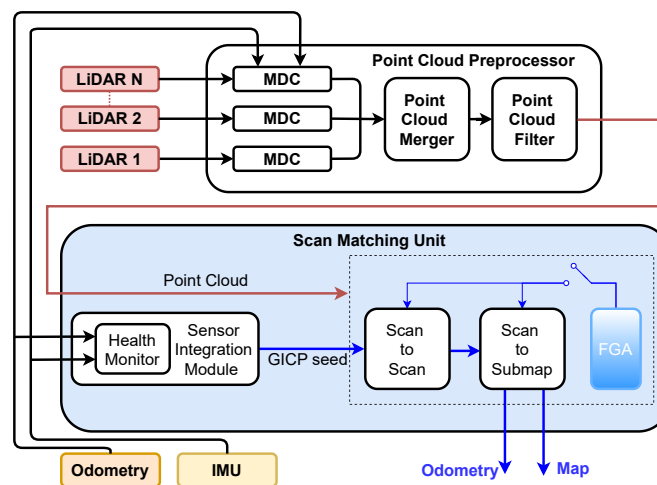


Figure 9. Architecture of the proposed system. In the block diagram MDC stands for motion distortion correction, while FGA stands for flat ground assumption.

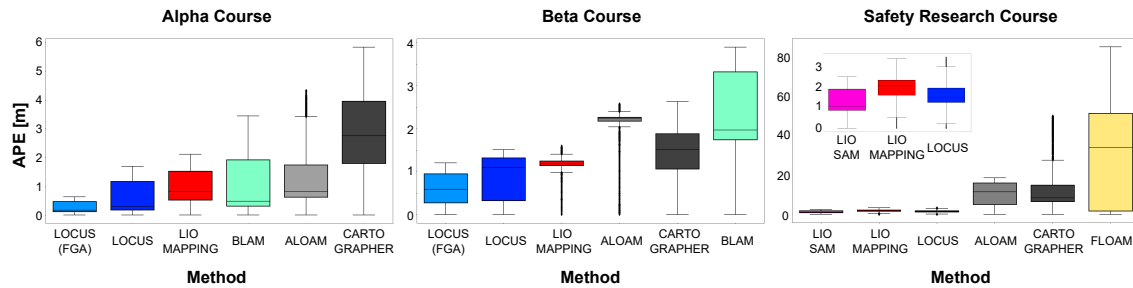


Figure 10. Boxplot visualization of the absolute position error (APE) computed for the different LiDAR-based methods on the perception-degraded underground test datasets: for clarity, only the best six algorithms in each dataset are shown. A larger drift is observed in the Safety Research Course as in this dataset points are not motion corrected and the course presents many perceptually challenging conditions such as harsh, unstructured terrains and long, featureless corridors: for this run, the inset reports a zoomed version of the performance of the best three algorithms.

an initial transform estimate from a non-LiDAR source to ease the convergence of the GICP in the scan-to-scan matching stage, by initializing the optimization with a near-optimal seed that improves accuracy and reduces computation, enhancing real-time performance. Multiple sources of odometry (e.g., VIO, KIO, WIO) and raw IMU measurements available on board are fed into a sensor integration module which selects the output from a priority queue of the inputs that are deemed healthy by a built-in health monitor, which prioritizes the order based on the expected accuracy of the methods. If the highest priority input is not healthy, then the next highest priority is used. If all sensors fail, the GICP is initialized with identity pose and the system reverts to pure LiDAR odometry. Notice how the confidence tests (health monitoring) depicted in Figure 8 are here incorporated within this sensor integration module. Palieri et al. (2020) provides more details on the system functioning.

LOCUS comparative results: We compare LOCUS with state-of-the-art LiDAR odometry methods (Shan et al., 2020; Ye et al., 2019; Hess et al., 2016; Zhang and Singh, 2014; Nelson, 2016) in extreme, perceptually degraded subterranean environments and demonstrate high localization accuracy along with substantial improvements in robustness to sensor failures. For the evaluation, we use the data collected in the Tunnel and Urban Circuit rounds of the DARPA Subterranean Challenge, from a wheeled ground robot carrying two Velodyne LiDARs, an IMU and running WIO onboard. To assess accuracy, we compute the absolute position error (APE) of the estimated robot trajectory against the ground-truth reference, for the different methods over the different runs, and report in Figure 10 a summary of the results. Throughout all the operations, LOCUS achieves highly accurate performance. To assess robustness, we analyze the flexibility of the various methods with respect to sudden failures of an input source by testing the following failure scenarios: (i) failure of WIO/IMU, (ii) failure of WIO, and (iii) failure of LiDAR. In these scenarios, tightly coupled approaches and methods designed with synchronized callbacks stop operating when an input is missing. In contrast, LOCUS consistently achieves reliable ego-motion estimation and mapping, demonstrating efficient handling of sensor failures in a cascaded fashion, behavior that is desirable to accommodate the unforeseen challenges posed by real-world operations where hardware failures are likely to happen, or sensor sources can become unreliable (see Palieri et al. (2020) for details).

5.1.2. IMU Front End

IMU pre-integration: The IMU front end is based on a pre-integration technique of the inertial measurements. This module makes use of the state-of-the-art on-manifold pre-integration theory to summarize the high-rate IMU measurements into a single motion constraint (Forster et al., 2015a,b) for the subsequent pose-graph optimization performed in the LiDAR-IMU data fusion algorithms. IMU pre-integration is also used to guarantee a pose and velocity estimate at high rate and low

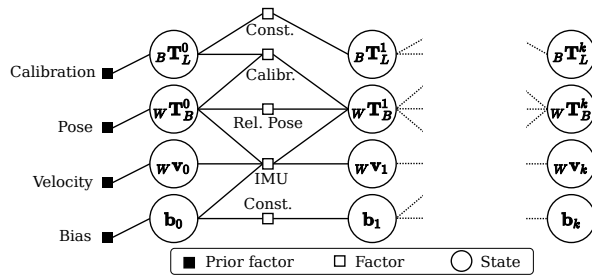


Figure 11. Factor graph representing state and measurements used by LION back end. Biases are denoted with $\mathbf{b}_j := \{{}_B \mathbf{b}^a, {}_B \mathbf{b}^g\}_j$.

latency, regardless of the time taken by the optimizer used in the back end of the sensor fusion algorithm.

5.1.3. LiDAR-IMU Data Fusion

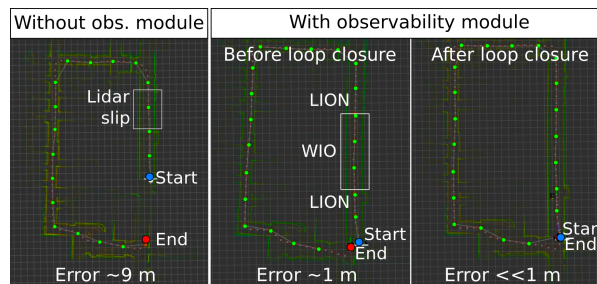
Smoothing framework (LION): The fusion of the relative ego-motion estimations, obtained from LOCUS and IMU pre-integration front ends, is performed via a fixed-lag smoother using a factor graph, as described in Tagliabue et al. (2021), where we introduce LION (LiDAR-Inertial Observability-aware Navigator for vision-denied environments). The state estimated by the proposed smoother consists of (a) the pose (position and attitude) ${}_W \mathbf{T}_B$ of the IMU-fixed reference frame B expressed in a slowly drifting inertial reference frame W ; (b) the linear velocity ${}_W \mathbf{v}$; (c) the IMU biases (${}_B \mathbf{b}^a$ for the accelerometer and ${}_B \mathbf{b}^g$ for the gyroscope), and (d) the extrinsic calibration ${}_B \mathbf{T}_L$ between the LiDAR-fixed frame L and the IMU frame B , introduced to reduce the effects of error in mounting the sensors, as well as to address the challenges in offline LiDAR-extrinsic calibration. A representation of the states and factors used in the factor graph with a window of k time steps can be found in Figure 11. Following Forster et al. (2015b), we model the smoothing problem using Georgia Tech Smoothing and Mapping (GTSAM) (Dellaert, 2012) and we solve the associated optimization with iSAM2 (Kaess et al., 2012).

Observability module: It is crucial for LiDAR-based estimation algorithms to determine if the geometry of the scene can well constrain the estimation of the translational motion, since long shafts and corridor-like structures can severely impact motion observability. Following Gelfand et al. (2003) and Bonnabel et al. (2016), we propose an observability metric, computed within LION architecture, which can inform the HeRO switching logic (Figure 8) about potential unreliability in the odometry output of LiDAR-based estimators. Such a metric is based on the condition number $\kappa(\mathbf{A}_{tt}) := |\lambda_{max}(\mathbf{A}_{tt})| |\lambda_{min}(\mathbf{A}_{tt})|^{-1}$ of the translational part \mathbf{A}_{tt} of the Hessian of the cost minimized by the point-to-plane ICP algorithm. The eigenvector associated with the smallest eigenvalue of \mathbf{A}_{tt} is the *least observable* direction for translation estimation. As a consequence, the larger the condition number $\kappa(\mathbf{A}_{tt})$ is, the more poorly constrained the optimization problem is in the translational part. More details are provided in our related work (Tagliabue et al., 2021).

Evaluation: We report the performance of the LiDAR-IMU fusion technique during the two tracks (A and B) of the Tunnel Competition at the DARPA Subterranean Challenge. The LOCUS output (selected, for this evaluation, to be the pure scan-to-scan matching from the GICP (Segal et al., 2009)) was fused at 10 Hz with the IMU, and consequently the fused output of LION, could be provided at up to 200 Hz. The sliding window of LION used here is 3 s and the factor graph optimization was tuned to use approximately 30% of one CPU core of an i7 Intel NUC. The root-mean-squared error (RMSE) for position (\mathbf{t} (m)) and attitude estimation (\mathbf{R} (rad)) and the percentage drift (\mathbf{t} (%)) are reported in Table 4, where we include a comparison with WIO and LOCUS (scan-to-scan), with the global localization algorithm LAMP (Ebadi et al., 2020) (presented in the following section) as ground truth. The results highlight that fusing inertial data with the odometry from the front end significantly reduces the drift of LiDAR's pure scan-to-scan matching. Additionally,

Table 4. Estimation error of wheel-inertial odometry, scan-to-scan matching, and LION for two runs of the two tracks of the Tunnel Competition, computed for one of the robots deployed.

Algorithm	Track A						Track B					
	Run 1 (685 m, 1520 s)			Run 2 (456 m, 1190 s)			Run 1 (467 m, 1452 s)			Run 2 (71 m, 246 s)		
	t (m)	t (%)	R (rad)	t (m)	t (%)	R (rad)	t (m)	t (%)	R (rad)	t (m)	t (%)	R (rad)
Wheel Inertial	130.50	19.05	1.60	114.00	25.00	1.28	78.21	16.75	0.99	6.91	9.79	0.12
Scan to Scan	105.47	15.40	0.90	18.72	4.11	0.18	56.6	12.14	0.79	4.55	6.45	0.27
LION	56.92	8.31	0.36	7.00	1.53	0.10	17.59	3.77	0.27	3.78	5.36	0.05

**Figure 12.** Comparison of the translation error in an office-like environment, with and without the observability module integrated in HeRO.

LION reliably estimates the attitude of the robot, and guarantees a gravity-aligned output provided at IMU rate. Last, in Figure 12, we report the performance of the observability module in an indoor, office-like environment, characterized by long corridors. The results show that when the observability module is not used (Figure 12, left), motion unobservability creates a *LiDAR slip*, producing a position estimation error of ≈ 9 m. When the observability module is used (Figure 12, left), the switching logic in HeRO switches to WIO instead of LION for the section of the corridor without LiDAR features. The total error is ≈ 1 m (“Before loop closure”). Improved state estimation (reduced drift in the output of HeRO) benefits the global mapping solution (Section 6), which can now correctly detect a loop closure (Figure 12, “After loop closure”), further reducing the drift.

5.2. Other Odometry Sources

Apart from the LiDAR-inertial odometry estimator, NeBula consists of other robust and resilient estimation algorithms developed to provide a robust state estimation while navigating in perception-challenging environments. Some examples are, for instance, a contact-inertial odometry estimation (Lew et al., 2019), where contacts are exploited to produce zero velocity updates into a Kalman filter that is integrating IMU measurements during a dead-reckoning situation, or a radar-inertial odometry (Kramer et al., 2020), which provides reliable ego-motion estimations even in the presence of obscurants thanks to the radar signal properties. The parallel combination of these heterogeneous estimation sources within the HeRO architecture provides a qualitative and robust state estimation that can be refined with a back-end algorithm providing global localization, as described in Section 6.

6. Large-Scale Positioning and 3D Mapping

NeBula’s SLAM solution, called LAMP (Large-scale Autonomous Mapping and Positioning), achieves low-drift, multi-robot, multi-sensor SLAM over large scales in perceptually degraded conditions. LAMP produces a consistent global representation of an unknown environment, along with the associated covariances to enable uncertainty-aware solutions across the NeBula system (Figure 7). In the context of the DARPA Subterranean Challenge, LAMP achieves the requirement

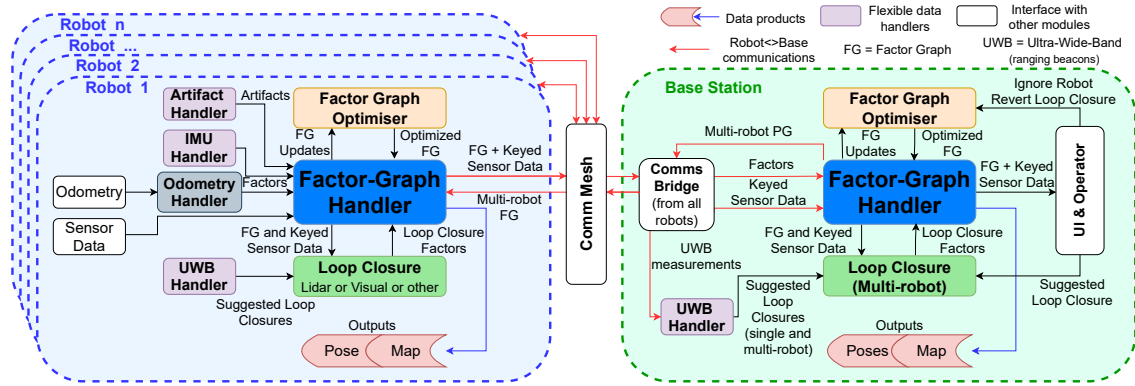


Figure 13. LAMP architecture. Each robot maintains its own factor graph (FG), which can then be fused with a multi-robot team on the base station, using a centralized architecture. The LiDAR and camera data that are used for loop closures and map building (white “Sensor Data” box) is labeled Keyed Sensor Data, being linked with a specific pose node. Handlers of artifact observations, IMU measurements, and ultrawideband (UWB) signals all process data to add constraints to the factor graph. Our robust optimization approach runs in parallel to optimize both the robot and multi-robot factor graphs.

for artifact localization error of less than 5 m over multiple kilometers of traverse. In this section, we will outline the architecture of LAMP, and then describe our approach to multi-sensor SLAM. Finally, we present results from representative field tests.

6.1. Subsystem Overview

As outlined in Figure 13, LAMP is a factor-graph-based SLAM solution, with the following key components: (a) an adaptable odometry input that can process individual or fused odometry sources, such as HeRO (Section 5), (b) a multi-modal loop closure module, based on LiDAR, visual, or semantic features, and (c) an outlier-resilient optimization of the factor graph, including multi-sensor inputs.

The flow of data starts with the odometry and sensor inputs, which add factors to the graph on the robot. Parallel processes then run loop closure searches and factor-graph optimization. Next, the graph is sent to the base station. The base station merges graphs from each robot into a common multi-robot graph that is further optimized with the addition of inter-robot loop closures. The main output products of LAMP are a set of poses describing the robot trajectory and the location of artifacts, as well as a point-cloud map.

Pose nodes and adaptable odometry input: To make the factor-graph optimization computationally tractable over large-scale, long-term multi-robot exploration, LAMP utilizes a sparse graph of pose nodes and edges (Figure 14). The edges are obtained from an accumulation of odometry measurements between two consecutive nodes (odometry edges) or from translation and rotation estimates between nonconsecutive nodes (loop closures edges, described below). A new pose node and linking odometry edge is created after traveling more than a threshold translation or rotation. To address the challenge of perceptual degradation for these odometry edges, we use HeRO (Section 5) as the input odometry source.

Multi-modal loop closures: A crucial capability to reduce the accumulated error in the robot trajectory is loop closure detection: the ability to correctly assert when a robot revisits a previously visited location or known landmark. Loop detection enables the computation of rigid-body 3D transformations between nonconsecutive pose nodes that can be added as loop closure edges in the factor graph (Figure 14). The multi-modal architecture of LAMP’s loop closure module (Figure 13) enables a robust and reliable system through the use of different sensing modalities. These loop closure sensing modalities include using LiDAR data (Ebadi et al., 2020), visual data (Rosinol et al., 2020), and semantic data (Ebadi et al., 2021).

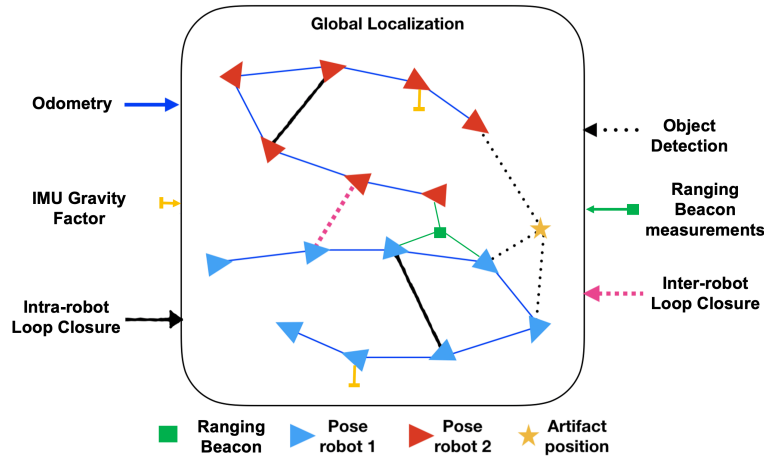


Figure 14. Graphical representation of the LAMP multi-robot pose graph, including odometry factors, inter- and intra- robot loop closures, gravity factors, ranging beacon measurements, and shared observations of landmarks (e.g., artifacts).

Multi-robot fusion: NeBula addresses the problem of multi-robot exploration of unknown environments by relying on LAMP’s multi-robot architecture. This architecture is centralized, to make use of agents with greater computational resources (such as a base station); however, for decentralized applications of NeBula we utilize the techniques described in [Choudhary et al. \(2017\)](#). In the centralized architecture the agent with the greatest computational capability serves as the central agent to fuse factor graphs constructed by individual robots into a consistent multi-robot graph, along with the associated sensor data ([Figure 13](#)). The factor graphs are fused using the same multi-modal loop closure modules as on the single robot, but instead of searching for intra-robot loop closures, these modules search for inter-robot loop closures. To further improve localization accuracy, we use the computational power of the central agent to perform batch loop closure analysis across the entire graph. This analysis identifies and computes additional inter- and intra-robot loop closures to add to the multi-robot graph. The updated multi-robot global graph is then optimized, and periodically sent back to the robots, for each agent to have a consistent global representation of the environment for global planning ([Section 9](#)).

Factor-graph optimization: Our factor-graph optimization (Kimera-RPGO, [Rosinol et al. \(2020\)](#)) uses a robust outlier rejection approach to reject the erroneous loop closures that can occur when operating in perceptually degraded conditions, such as with obscurants and self-similar environments. Kimera-RPGO rejects erroneous loop closures by finding the largest consistent set of loop closures for each set of single-robot and inter-robot loop closures, using a consistency graph and max clique search (an adaptation of [Mangelson et al. \(2018\)](#)). The loop closures that are not in the consistent set are discarded prior to optimization (see [Ebadi et al. \(2020\)](#) and [Lajoie et al. \(2020\)](#) for details). The updated factor graph is then optimized with a Levenberg–Marquardt solver that is implemented in GTSAM ([Dellaert, 2012](#)).

6.2. Additional Factors and Multi-Sensor Fusion

LAMP fuses multiple sensing inputs into the factor graph ([Figure 14](#)) to improve the robustness and accuracy of the SLAM solution. We present four examples here:

1. *IMU Gravity Factors:* When the robot is stationary, the accelerometers on the IMU can be used to obtain a robust estimate of the local gravity vector, which is added to the factor graph as a constraint on roll and pitch (yellow factors in [Figure 14](#)).




Figure 15. Example calibration to global frame from the DARPA Subterranean Challenge Urban Circuit. The coordinates of the survey prisms and reflective patches on the gate are provided and define the reference frame. A survey station measures the prisms to determine its pose in the global frame, after which it can measure the position of robots. LiDAR measurements of the reflective patches provide a yaw estimate and an IMU computes roll and pitch assuming gravity alignment of the reference frame.

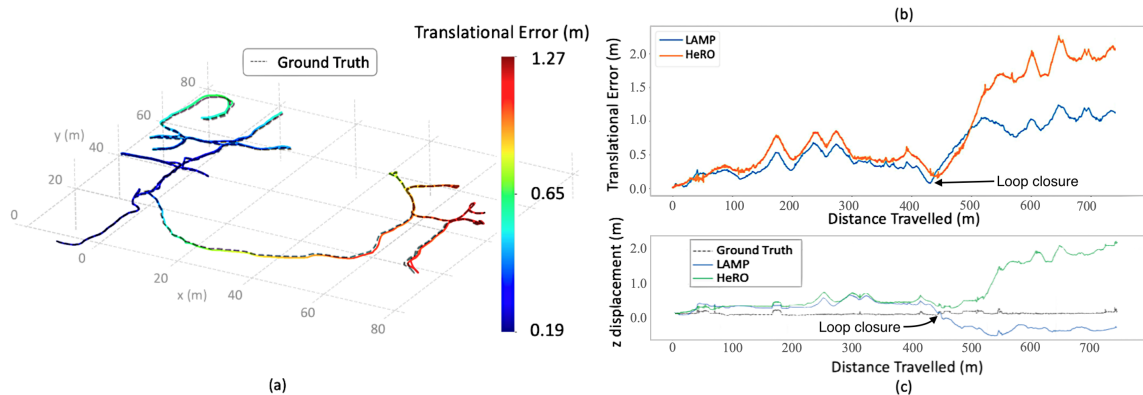
2. *Landmark Factors*: Measurements of distinct landmarks can either be used to detect loop closures or to directly provide constraints to the factor graph. These landmarks fall into two categories.
 - (a) *Deployed Landmark Factors*: These landmarks include visual beacons, ranging beacons, or retro-reflecting beacons and are deployed from a robot while exploring an unknown environment. For example, we have implemented deployable UWB ranging beacons in our system (see [Section 12.2](#) for hardware details, and [Funabiki et al. \(2020\)](#) for algorithmic details). The signals from the beacons robustly and efficiently identify loop closures, to seed LiDAR- or vision-based alignment computations for single- and multi-robot teams (green node and factors in [Figure 14](#)).
 - (b) *Environmental Landmark Factors*: Existing features in the environment such as signs, salient objects, and the shape of junctions (e.g., [Ebadi et al., 2021](#)) can be used as landmarks. For example, we use observations of specific objects, such as backpacks and fire extinguishers (called artifacts in SubT), with sets of range-bearing observations (dashed black lines in [Figure 14](#)) from the artifact relative-localization module ([Section 7](#)). By fusing the object observations into the factor graph we also ensure the most up-to-date global location of those objects for situational awareness (and scoring in SubT).
3. *Calibration Factors*: At the start of a mission, each robot is aligned with a global reference frame from a combination of LiDAR, IMU, and survey station measurements of the robot and a calibration gate (e.g., [Figure 15](#)). These initial calibration measurements, as well as any additional measurements generated during the mission, are added as constraints to the factor graph.

6.3. Metric and Semantic Map Generation

LAMP builds both a geometric and semantic global map from sensor measurements attached to the nodes in the factor graph. Both maps are built by projecting sensor measurements into the global reference frame by using the latest, optimized state of the associated pose nodes in the factor graph. For the geometric map, these sensor measurements (keyed sensor data in [Figure 13](#)) are either point clouds (from LiDAR, depth cameras, or visual feature tracking) or local occupancy grids. In particular, structures like confidence-rich occupancy grids ([Agha-mohammadi et al., 2019](#)) allow for encoding the environmental uncertainty, which then can be used for perception-aware coverage planning and enabling SMAP-like behaviors ([Heiden et al., 2017](#)). For the semantic map,

Table 5. Performance statistics for LAMP operating on single-robot benchmark datasets from different robots and in different environments (Figure 17)

Dataset	Dataset Characteristics			Absolute Transl. Error		Absolute Rot. Error	
	Robot	Environment	Dist. (km)	Max (m)	Mean (%)	Max (deg)	Mean (deg/m)
(a)	Husky	Tunnel	1.65	9.7	0.93%	5.3	0.006
(b)	Spot	Urban	0.65	2.2	0.46%	5.0	0.019
(c)	Husky	Urban	0.62	3.5	0.42%	3.9	0.011
(d)	Husky	Urban	0.75	0.9	0.19%	1.8	0.006
(e)	Spot	Cave	0.6	10.6	1.68%	6.0	0.020

**Figure 16.** Single-robot LAMP performance with loop closures and IMU factors in the Urban Circuit of the DARPA Subterranean Challenge: (a) LAMP localization accuracy compared to DARPA-provided ground truth. (b) Translational error against distance traveled for LAMP, and the HeRO odometry input to LAMP. (c) z trajectory against distance traveled for LAMP, HeRO, and the ground truth. IMU factors included in LAMP help to constrain attitude drift and achieved improved performance in z , and overall, compared to the input odometry.

the sensor measurements are descriptive observations, such as detections of distinct objects (e.g., backpacks, survivors), semantic classifications of 3D spaces (e.g., doorways, stairs), or ambient measurements (e.g., temperature, pressure, gas concentration). The resulting 3D semantic map provides rich situational awareness to the operator of the robotic team, and can be the critical output data product of the overall system. The semantic map is especially important in the context of SubT, where the semantic map primarily consists of the globally localized artifact observations (both objects and ambient measurements), which is exactly the information needed for scoring (Section 7).

6.4. LAMP Performance

The performance of LAMP on a single robot dataset from a husky robot equipped with three LiDARs is demonstrated in Figure 16. LAMP achieves error at less than 0.2% of the distance traveled, with the IMU gravity factors assisting in reducing the z error in the latter portion of the trajectory (Figure 16(c)). Further single-robot tests are summarized in Table 5, from the five benchmark datasets shown in Figure 17. These benchmarks show LAMP achieving the accuracy better than 5 m on all other than the tunnel and cave datasets, which have motion-distorted LiDAR measurements. Multi-robot LAMP performance is demonstrated in Figure 18 with two huskies deployed in an urban environment. The results demonstrate the benefit of using UWB beacons (Figure 18(c)) compared to using pure LiDAR-based loop closures. Further results are presented in Section 13. Please refer to Ebadi et al. (2020) for detailed results in tunnel environments, including the impact of loop closures, Funabiki et al. (2020) for results using UWB beacons for a single robot, and Morrell et al. (2020) for single-robot cave exploration results.




Figure 17. Maps from the five LAMP single-robot benchmark datasets. (a) Tunnel dataset from Bruceton mine on a Husky. (b) Multi-level urban dataset from Satsop Power Plant on a Spot robot. (c) and (d) Urban dataset from Satsop Power Plant on a Husky robot platform. (e) Cave dataset from Valentine Cave on a Spot robot. The ground truth is provided by DARPA for (a)–(d) and by survey scans for (e).

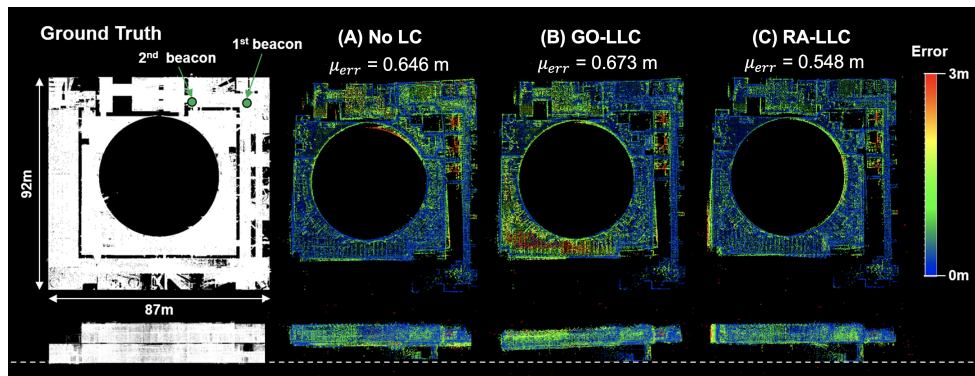


Figure 18. Multi-robot map accuracy results: (a) no loop closure (No LC); (b) geometric-only LiDAR loop closure (GO-LLC); and (c) range-aided LiDAR loop closure (RA-LLC). In this test, data from two robots in the Urban Circuit of the DARPA Subterranean Challenge are used, with two deployed UWB beacons (as shown on the left overlaid on the DARPA-provided ground truth map). μ_{error} indicates the mean map error to the ground truth map.

7. Semantic Understanding and Artifact Detection

Semantic understanding of the environment and detecting objects of interest and artifacts are important capabilities to enable higher levels of robotic autonomy in unknown environments. Semantic mapping and artifact detection are among the main components of the NeBula autonomy framework. This section discusses NeBula’s solution for detecting, localizing, and visualizing objects of interest on heterogeneous robots with different sensor configurations. Here, we focus on both

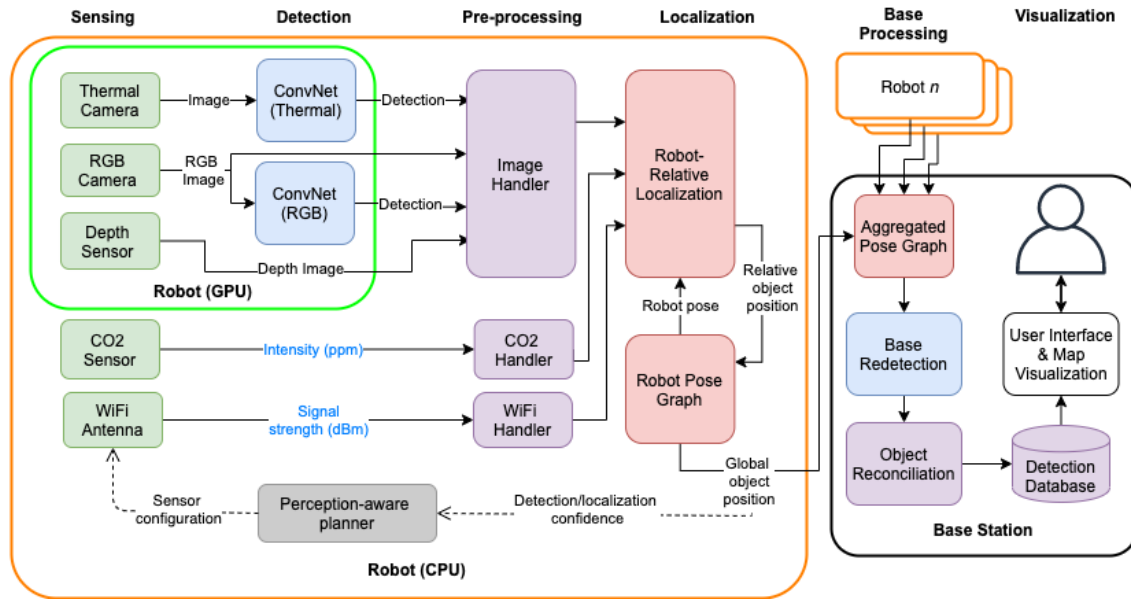


Figure 19. Multi-robot, multimodal object detection, localization, and visualization pipeline.

(i) static objects, with clear visual, thermal, or depth signature, and (ii) spatially diffused phenomena such as gas propagation and WiFi signal. The pipeline is explained in detail in [Terry et al. \(2020\)](#), with a summary and recent extensions detailed here.

Requirements: The object detection system needs to (i) make detections in real time across multiple sensor modalities, (ii) permit high-accuracy localization, (iii) adjust the sensor configuration based on the detection and localization confidence, and (iv) apply filtering to present the most likely detection candidates to the mission supervisor (when a communication link is established). While the method presented in this section is general, in the context of the DARPA Subterranean Challenge, we focus on a set of predefined object types including gas sources (e.g., CO₂ source) and man-made objects such as fire extinguisher, drill, rope, helmet, survivor manikin, backpack, vent, and cell phone ([Agrawal, 2019](#); [Orekhov and Chung, 2021](#)). Object signatures exist in one or more modalities: visual, thermal, depth, WiFi, audio, and chemical.

Architecture: [Figure 19](#) shows the proposed object detection, localization, and visualization pipeline. We break down the underlying object detection problem into two stages: (1) an image-based object detection pipeline to first find the object and (2) a relative localization filter applying projective geometry to the detection to estimate its position explicitly. By splitting the detection and localization tasks, we can utilize high quality and fast detection from existing algorithms and apply them to generic camera types in our relative localization approach. For temporally static objects, NeBula relies on multi-modal detection and, when available on robots, makes use of visual cameras, depth measurements, and thermal cameras.

Spatially diffuse phenomena: For temporally dynamic and spatially diffused phenomena, we rely on source-seeking methods based on gas sensors and WiFi sensing. The detection confidence provides the uncertainty assessment to the perception-aware planner. The planner motivates the sensor to adjust the configuration to make new measurements with higher fidelity, leading to more accurate detections. We discuss the three stages of detection, localization, and base station processing in the rest of this section.

Detection: For visually observable objects, detections are made in both the color and thermal spectra using state-of-the-art convolutional neural networks (CNNs). A CNN produces a bounding box on the image to pass to the relative localization module. NeBula relies on different CNN implementations to adapt to various processing capabilities. On ground robots, a YOLO Tiny ([Redmon](#)

and Farhadi, 2018; Bochkovskiy et al., 2020) variant is used, leveraging GPU hardware (e.g., Nvidia Jetson Xavier) to run in real time on multiple cameras (see Section 12). On drones, a MobileDet variant (Xiong et al., 2020) is used, modified to run on a Google EdgeTPU. To achieve sufficient detection performance for our specific application, we fine-tuned the detection networks with the appropriate domain-specific dataset. In the context of SubT, we have produced more than 40,000 annotations of the following objects: fire extinguisher (12.26%), drill (9.98%), rope (17.90%), helmet (19.29%), survivor (19.84%), and backpack (20.72%). Our training is focused on maximizing recall that increases true positives. This is followed by an outlier rejection method using range, color, and size filters, reducing the false positives. For more details on training methods on this data, see Terry et al. (2020).

Localization: The detection networks produce 2D bounding boxes (within image) that are combined with depth measurements, used to compute the position of the artifact relative to the robot. The depth measurements can be obtained from multiple sources: depth cameras (such as from an RGB-D camera), LiDAR scans mapped into the camera frame, or a size-based projection, where the depth is computed such that the bounding box, when projected to 3D at that depth, matches the expected size of the object. These methods are detailed in Terry et al. (2020). For robustness to sensor failure, each method is run in parallel, with the highest priority method (LiDAR, then depth, then size projection) used if the corresponding sensor is available. All methods jointly filter multiple detections to produce a combined relative location reported to the LAMP module (Section 6) to compute the global location before sending the report to the base station. The relative location can also be computed without depth measurements, as an additional back-up, and on systems without depth cameras or LiDAR (such as drones). In this case, we use a monocular-based tracking approach over an image sequence (detailed in Ramtoula et al., 2020).

Multi-modal multi-robot artifact reconciliation: The artifact reports from each of the robots are further processed on robots with more powerful computational resources (or on base station). These reports include (i) detection class, (ii) detection confidence, (iii) reference RGB/thermal image, (iv) bounding box, and (v) location estimate. The base station processes the reports, rejecting outliers and matching observations of the same artifact instances from other agents or previous visits. To reduce the number of false positive reports, a larger and more performant detection network (YOLOv4, Bochkovskiy et al., 2020) is used to update the detection confidence of each report. Then, the report is compared to previous reports of the same class to identify repeat observations and observations of the same artifact instance reported by other agents. This comparison uses both location proximity and a comparison of NetVLAD visual image descriptors (Arandjelovic et al., 2016). The highest confidence report of each object is then saved to a database and visualized for review by the human supervisor using the mission executive interface (see Figure 20, visualization block). When needed, to increase the confidence on the detected semantics, the perception-aware motion planner seeks new measurements, e.g., from a closer or better angle to the target, or by sending a different robot with complementary sensors to get multiple readings from the target.

7.1. Spatially Diffuse Localization

To detect and locate spatially diffuse phenomena, such as gas leaks and WiFi sources, the robotic team is leveraged as a mobile sensor network, with distributed and moving ambient sensor measurements. Signal strength (e.g., CO₂ concentration or WiFi RSSI) is recorded at every robot position and is (i) used to augment the global 3D semantic map (Figure 20, detection block) and (ii) processed to produce an initial location estimate at the area of peak signal strength. The combination of the spatially informative semantic map with an initial location estimate seeds a local search for source locations, based on the 3D geometry. Automation of this local search is ongoing work. In tests presented here, information is sent to the base station for displaying to the operator, who performs the local search on inspection of the metric-semantic map.

Detection performance: Figure 21 shows examples of a true and false positive detection for each visual artifact type. We observed that spray paint markings and existing equipment in the

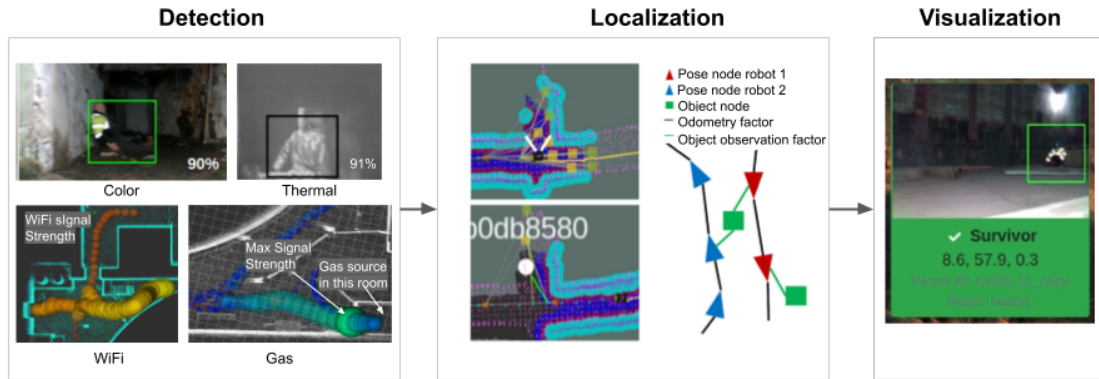


Figure 20. Visualization of representative parts of the object detection pipeline. In the left pane, a sample of detections for each of the four modalities (color, thermal, WiFi, and gas) is presented. In the middle pane, the addition of a natural landmark observation into the pose graph is depicted. The right pane shows the final operator view.

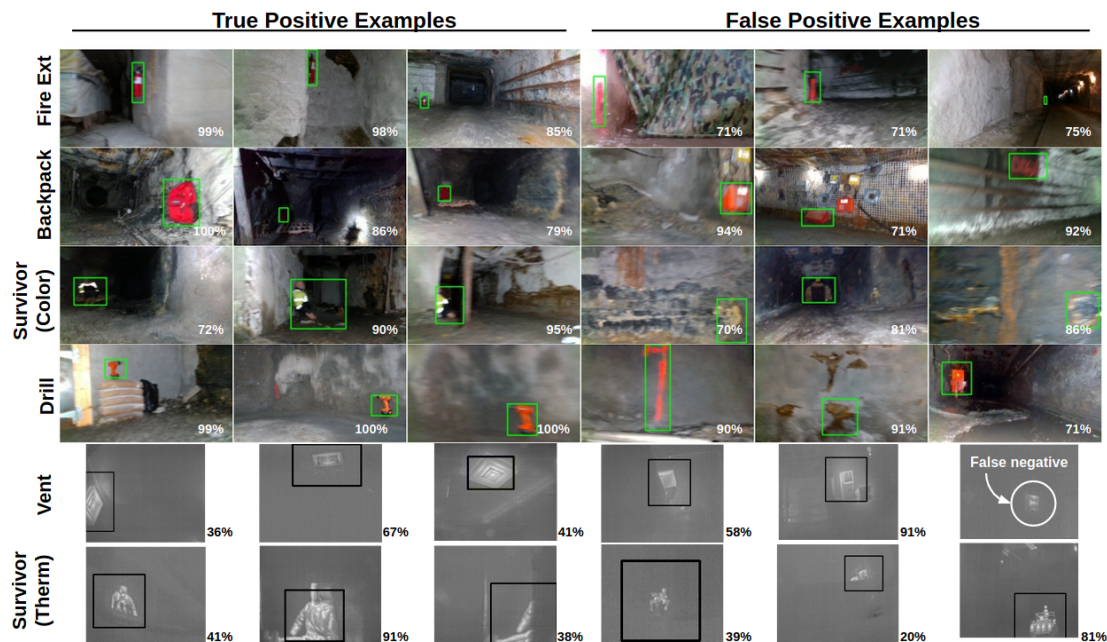


Figure 21. Examples of true and false positive detections of visual artifacts.

environments, which share the same gross features as the target objects, are incorrectly picked up. For spatially diffuse detection (gas and WiFi), we extrapolate the source location by measuring the signal strength gradient and move the robot in directions that increase the detection confidence (see Figure 20, detection block).

8. Risk-Aware Traversability and Motion Planning

A fundamental component of NeBula is its risk-aware traversability and motion planning (see Figure 7). This component, which we call STEP (stochastic traversability evaluation and planning), allows the robots to safely traverse extreme and challenging terrains by quantifying uncertainty and risk associated with various elements of the terrain. In this section, we briefly discuss NeBula's



Figure 22. Mobility-stressing elements commonly found during testing in various tunnel, urban, and cave environments, including 800 ft underground at Arch Mine in Beckley, WV, Valentine Cave at Lava Beds National Monument, CA, Satsop power plant in Elma, WA, and Mars Yard at JPL, Pasadena, CA.

traversability analysis, motion planning, and controls. For more details please see (Thakker et al., 2020) and Fan et al. (2021).


8.1. Design Philosophy

Challenges in extreme terrain motion planning: Unstructured obstacle-laden environments pose large challenges for ground roving vehicles with a variety of mobility-stressing elements. Common assumptions of a benign world with flat ground and clearly identifiable obstacles do not hold; environments introduce high risks to robot operations, containing difficult geometries (e.g., rubble, slopes) and nonforgiving hazards (e.g., large drops, sharp rocks) (Kalita et al., 2018; Léveill e and Datta, 2010). Additionally, subterranean environments pose unique challenges, such as overhangs, extremely narrow passages, etc. See Figure 22 for representative terrain features. Determining where the robot may safely travel has several key challenges: (i) localization error severely affects how sensor measurements are accumulated to generate dense maps of the environment; (ii) sensor noise, sparsity, and occlusion induces large biases and uncertainty in mapping and analysis of traversability; and (iii) the combination of various risk regions in the environment create highly complex constraints on the motion of the robot, which are compounded by the kinodynamic constraints of the robot itself.

System architecture: To address these issues, we develop a risk-aware traversability analysis and motion planning method, which (1) assesses the traversability of terrains at different fidelity levels based on the quality of perception, (2) encodes the confidence of traversability assessment in its map representation, and (3) plans kinodynamically feasible paths while considering mobility risks. Figure 23 shows an overview of the local motion planning approach. The sensor input (point cloud) and odometry are sent to the risk analysis module, evaluating the traversability risk with its estimation confidence. The generated risk map is used by hierarchical planners consisting of a geometric path planner and a kinodynamic MPC (model predictive control) planner. The planners replan at a higher rate to react to the sudden changes in the risk map. The planned trajectory is executed with a tracking controller, which sends a velocity command to the platform.

Robot agnosticism: Our approach is highly extensible and general to our different ground robot types, requiring only a change in the dynamics model of the system. Moreover, using this approach, we are able to specify a wide array of constraints and costs, such as limiting pitch or roll of the vehicle on slopes, preferring one direction of motion, keeping some distance from obstacles, fitting through narrow passages, or slowing down/stopping around risky areas. This flexibility has proven

Table 6. Specification of the Different Robots. Superscript on the robot type indicates location/source of data used for computing statistics: ¹DARPA SubT Urban Competition, ²DARPA SubT Tunnel Competition, ³Beckley Exhibition Coal Mine, WV.



Specification	Skid-steer ^{1,2}	Tracked ³	Ackermann ²	Quadruped ¹
Model	Husky A200	Telemax Pro	X-Maxx	Spot
Max. speed	1.0 m/s	1.1 m/s	22 m/s	1.6 m/s
Steering	Skid-steer	Skid-steer	Ackermann	Gait
Distance traveled	5.98 km	0.76 km	0.39 km	2.85 km
Avg. speed during traverse	0.93 m/s	0.57 m/s	0.83 m/s	0.65 m/s
Autonomous recoveries/km	6.9	0.16	0.0	13.0
Critical failures/km	0.2	0.0	0.0	1.1

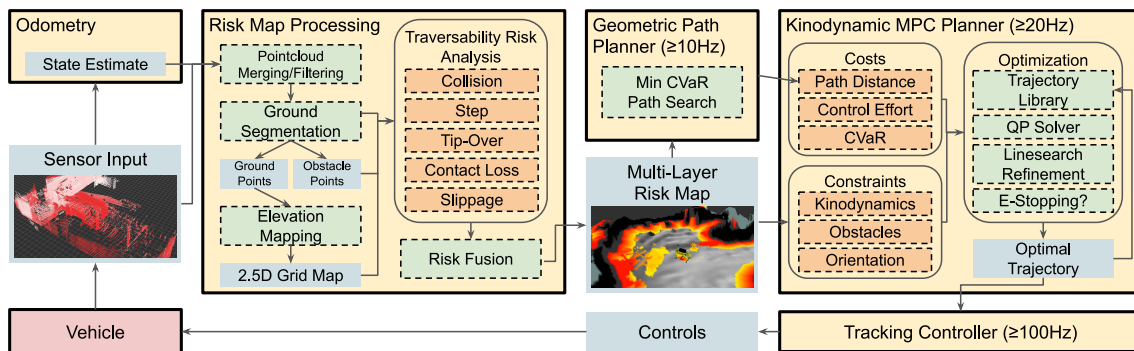


Figure 23. Traversability and motion planning architecture overview.

important in achieving robust navigation across the extreme traversability challenges encountered in highly unstructured environments (see Table 6).

8.2. Uncertainty-Aware Traversability

The importance of considering uncertainty: Precise traversability analysis and motion planning relies heavily on sensor measurements and localization. However, the quality of state estimation can often degrade, especially in perceptually challenging environments such as tunnels, mines, and caves. Additionally, sensors are subject to noise of various types, as well as occlusion, restricted field of view, etc. Therefore a key idea is to incorporate uncertainty awareness into our mapping for traversability and motion planning. We accomplish this by a multi-fidelity mapping approach in which we weigh more strongly high-confidence information from recent sensor measurements which are closer to the robot. Older and farther sensor measurements from the body of the robot are associated with higher uncertainties and decay more quickly. By aggregating these sensor measurements in an uncertainty-aware way, we create a robust and resilient belief-aware local map which is then used for traversability analysis. To reduce uncertainty and achieve higher levels of resiliency, we rely on multi-sensor high-FOV measurements (dense depth camera data and LiDAR




Figure 24. Comparison of the mean, VaR, and CVaR for a given risk level $\alpha \in (0, 1]$. The axes denote the values of the stochastic variable ζ , which in our work represents traversability cost. The shaded area denotes the $1 - \alpha\%$ of the area under $p(\zeta)$. $\text{CVaR}_\alpha(\zeta)$ is the expected value of ζ under the shaded area.

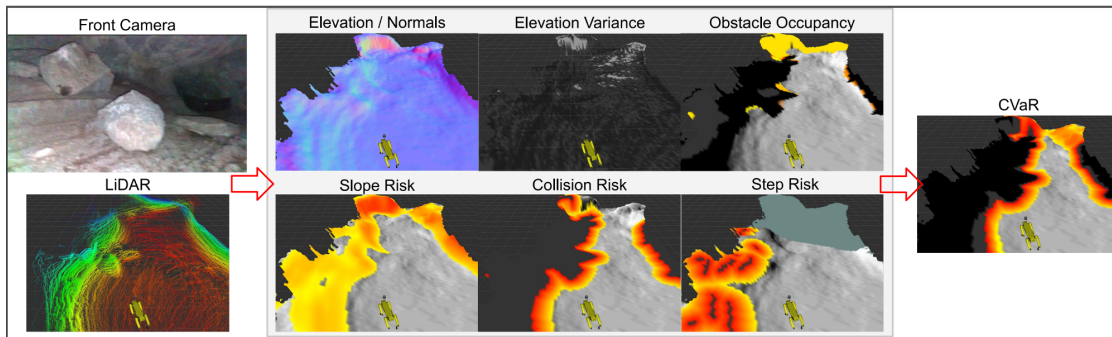


Figure 25. Traversability risk analyses which consider different sources of risk, as well as uncertainties, and fuse them into one CVaR metric costmap. Left: Raw sensor measurements are used to construct a point-cloud map. Middle: Statistical properties of the map are identified as various sources of traversability risk. Right: Risk sources are fused into one CVaR metric, which will be used for planning.

data) to efficiently update local maps and reduce sensitivity to localization uncertainties (similar, related work includes [Ahmad et al. \(2021\)](#) and [Hines et al. \(2021\)](#)).

Traversability risk analysis: To assess traversability risks, we utilize the constructed multi-fidelity local map. A ground segmentation method ([Himmelsbach et al., 2010](#)) is applied to the merged point cloud to filter obstacle and ceiling points. The ground point cloud is used to build a 2.5D elevation map for efficient query of terrain geometry. The elevation map and segmented point clouds are used to assess the risk of traverse from various perspectives including collision, tip-over, traction loss, and negative obstacles. Individual risk analysis is fused into a single risk value estimate with a confidence value, and sent to the planning module ([Figure 25](#)). We define this risk value in terms of a CVaR (conditional value-at-risk) metric ([Figure 24](#)), which quantifies the severity of the risk of a given path given all uncertainties ζ from the traversability risk analysis according to a desired threshold of probability α . This threshold can be changed by mission-level decision making in order to vary the level of acceptable risk during the mission. This risk metric is particularly useful as it captures tail events with low probability of occurrence, which may have high consequences on the success of the mission and should be taken into account. By approximating all uncertainties with a Gaussian distribution, the CVaR is efficiently evaluated to account for different types of terrain ([Fan et al., 2021](#)).

Semantic traversability factors: In addition to geometric traversability analyses, we can also identify certain terrain features semantically and incorporate them into the risk map. Features such as water are sometimes identifiable in LiDAR point clouds due to differences in reflectivity. Other features such as stairs can be identified using computer vision methods, which locate stairs semantically and also identify stair slope, angle, and height from known geometric priors. This information is useful for identifying nongeometric mobility risks, as well as notifying our planners to approach certain hazards differently (e.g., walking down stairs).

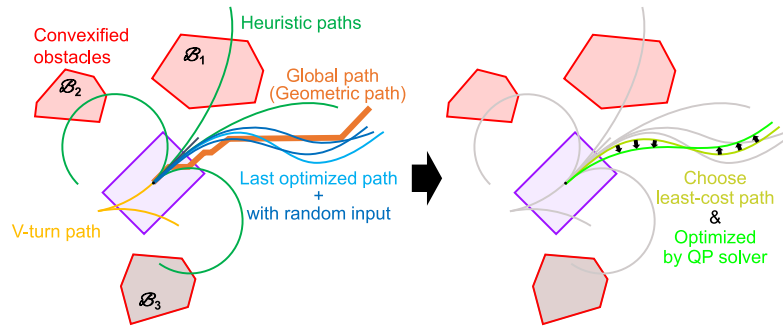


Figure 26. Schematic of our risk-aware kinodynamic trajectory optimization approach, which combines a trajectory library-based search with convex optimization. Global path is determined by A* over a 2D grid. Heuristic paths (including v-turns, random input, etc.) are not necessarily feasible but are used to initialize the local trajectory optimization. Obstacles, constraints, and costs are convexified about the best initial path and used to iteratively solve convex quadratic programming (QP) optimization problems in real time.

8.3. Uncertainty-Aware Traversability Analysis and Motion Planning

Efficient risk-aware kinodynamic planning: Using the computed CVaR metric values on the map, we must search for a path which minimizes these values. This is done in a two-stage hierarchical fashion. The first stage operates on longer distances (40 m) and takes into account positional risk. Using an A* algorithm over a 2D grid, this first stage yields a global geometric plan that minimizes the risk of this path. The second stage operates on shorter distances (8 m) and searches for a kinodynamically feasible trajectory that minimizes CVaR, while maintaining satisfaction of various constraints including obstacles, dynamics, orientation, and control effort. This kinodynamic planner operates in an MPC fashion, and is based on a combination of stochastic trajectory optimization and gradient-based convex optimization techniques (Fan et al., 2021). It runs efficiently in real time at 20–50 Hz, requiring roughly $\sim 20\%$ of one CPU core. As shown in Figure 26, this accounts for a rich set of heuristics, making it robust to local minima. Once a trajectory is optimized, it is sent to an underlying tracking controller for execution on the platform.

Recovery behaviors: In the real world, failures are unavoidable. While our risk-aware traversability and planning architecture is meant to reduce the probability of failures, they do happen from time to time, as a consequence of failure in localization, undetected edge cases in traversability, hardware failures, or unknown unknowns. As a last line of defense, we design behaviors to recover the system from nonfatal failures. Recovery behaviors are autonomously executed when we locally detect that a commanded motion is not being followed, or no valid and safe path is found to move away from the current position. These behaviors include clearing/resetting the local traversability map, increasing the allowable threshold of risk (to try to escape an untraversable area), and moving the robot in an open-loop fashion towards the direction of maximum known free space. In most cases these recovery behaviors are sufficient to recover the robot from a stuck condition, as long as the robot has not suffered a catastrophic failure.

Learning and adaptation: Over the course of a mission we often see changes in vehicle dynamics or environmental factors. For example, decaying battery life, mechanical wear, or mud/water can all affect the vehicle’s intrinsic dynamics. Additionally, changes in surfaces in the environment can strongly affect vehicle motion. To adapt to these changes we employ learning-based methods using Gaussian processes which adapt critical vehicle parameters and dynamics models based on the past history of performance (Fan et al., 2020a,b). By accounting for both the epistemic and aleatoric uncertainties of these statistical models, these methods ensure safety and robustness even in light of changing dynamics models.

Ongoing work: Our ongoing work lies in increasing the ability of NeBula to account for and handle uncertainties in both perception and planning. One major thrust of ongoing work involves




Figure 27. Belief cloud: We aggregate uncertainties in sensor noise and localization into a belief point cloud. (Ellipses describe uncertainties associated with points at the center of each ellipse.) This can be used for accurate and efficient uncertainty-aware traversability analysis and planning.

perception-aware planning. By incorporating sensor models into our traversability maps, we can move towards perception-aware behaviors which maximize sensor coverage, optimally reduce uncertainty, and automatically generate active learning behaviors. This thrust is particularly important in perceptually degraded and complex environments filled with occlusion and traversability hazards. A second direction of ongoing work lies in incorporating localization uncertainty into our traversability risk mapping in a more theoretically satisfying way, while remaining computationally tractable. We call this approach “mapping using belief clouds,” and the idea is to propagate the uncertainty of the robot pose while taking sensor measurements into an aggregated point cloud. The result is a belief cloud, i.e., a point cloud which encodes uncertainty information directly and efficiently. Our ongoing work involves extracting traversability metrics and risks from these uncertainty-aware mapping data (see Figure 27).

9. Uncertainty-Aware Global Planning

Autonomous global planning for environment exploration and coverage is a core part of the NeBula architecture (see Figure 7). NeBula formulates the autonomous exploration in unknown environments under motion and sensing uncertainty by a partially observable Markov decision process (POMDP), one of the most general models for sequential decision making. This formulation allows NeBula to jointly consider sequential outcomes of perception and control at the planning phase in order to achieve higher levels of resiliency during the mission operation. In this section, we discuss our POMDP-based global planning, referred to as PLGRIM (probabilistic local and global reasoning on information roadmaps). For more details, please see Kim et al. (2021), Bouman* et al. (2020) and CoSTAR Team (2020a).

9.1. Problem Formulation

POMDP formulation: A POMDP is described as a tuple $\langle \mathbb{S}, \mathbb{A}, \mathbb{Z}, T, O, R \rangle$, where \mathbb{S} is the set of states of the robot and world, and \mathbb{A} and \mathbb{Z} are the set of robot actions and observations, respectively (Kaelbling et al., 1998; Pineau et al., 2003). At every time step, the agent performs an action $a \in \mathbb{A}$ and receives an observation $z \in \mathbb{Z}$ resulting from the robot’s perceptual interaction with the environment. The motion model $T(s, a, s') = P(s' | s, a)$ defines the probability of being at state s' after taking an action a at state s . The observation model $O(s, a, z) = P(z | s, a)$ is the probability of receiving observation z after taking action a at state s . The reward function $R(s, a)$ returns the expected utility for executing action a at state s . In addition, a belief state $b_t \in \mathbb{B}$ at time t is introduced to denote a posterior distribution over states conditioned on the initial belief b_0

and past action-observation sequence, i.e., $b_t = P(s | b_0, a_{0:t-1}, z_{1:t})$. The optimal policy $\pi^* : \mathbb{B} \rightarrow \mathbb{A}$ of a POMDP for a finite receding horizon is defined as follows:

$$\pi_{t:t+T}^*(b) = \operatorname{argmax}_{\pi \in \Pi_{t:t+T}} \mathbb{E} \sum_{t'=t}^{t+T} \gamma^{t'-t} r(b_{t'}, \pi(b_{t'})), \quad (1)$$

where $\gamma \in (0, 1]$ is a discount factor for the future rewards, and $r(b, a) = \int_s R(s, a) b(s) ds$ denotes a belief reward which is the expected reward of taking action a at belief b . T is a finite planning horizon for a planning episode at time t . Given the policy for last planning episode, only a part of the optimal policy, $\pi_{t:t+\Delta t}^*$ for $\Delta t \in (0, T]$, will be executed at run time. A new planning episode will start at time $t + \Delta t$ given the updated belief $b_{t+\Delta t}$. The computational complexity of a POMDP grows exponentially with the planning horizon (Pineau et al., 2003), and we tackle this challenge with hierarchical belief space representation and planning as to be detailed in Section 9.2.

Application to simultaneous mapping and planning (SMAP): To formalize our SMAP problem as a POMDP, we define the state $s = (q, W)$ as a pair of robot q and world state W . We further decompose the world state as $W = (W_{occ}, W_{cov})$ where W_{occ} and W_{cov} describe the occupancy and the coverage states of the world, respectively associated with their uncertainties (e.g., Agha-mohammadi et al., 2019). A reward function for coverage can be defined as a function of information gain I and action cost C as follows:

$$R(s, a) = \operatorname{fn}(I(W_{cov}, z), C(W_{occ}, q, a)), \quad (2)$$

where $I(W_{cov}, z) = H(W_{cov}) - H(W_{cov} | z)$ is quantified as reduction of the entropy H in W_{cov} after observation z , and $C(W_{occ}, q, a)$ is evaluated from actuation efforts and risks to take action a at robot state q on W_{occ} . Minimizing this cost function in Eq. (1) simultaneously solves for the mapping and planning (SMAP), maximizing the coverage for *artifact detection* and minimizing the action risk (e.g., collision chance). This reward function can be generalized to SLAP problems (e.g., Agha-mohammadi et al., 2018 or Kim et al., 2019b) by incorporating information gain based on localization entropy reduction events, such as a loop closure or landmark detection.

9.2. Hierarchical Coverage Planning on Information Roadmaps

In this subsection, we introduce NeBula's solution for uncertainty-aware global coverage planning, PLGRIM (Probabilistic Local and Global Reasoning on Information roadMaps) (Kim et al., 2021). PLGRIM proposes a hierarchical belief representation and belief space planning structure to scale up to spatially large problems while pursuing locally near-optimal performance (see Figure 28). At

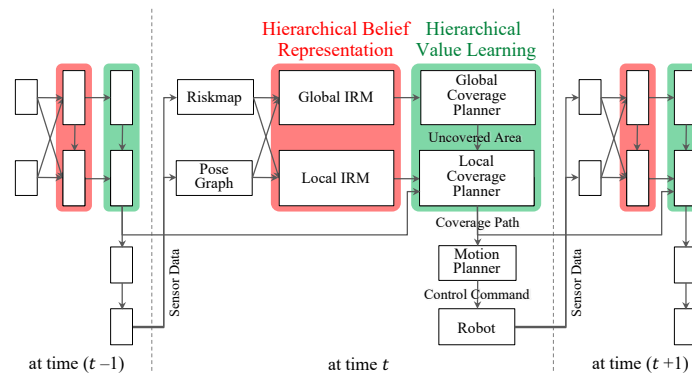


Figure 28. Illustration of PLGRIM framework for large-scale exploration in unknown environments. Over the receding-horizon planning episodes, PLGRIM (i) maintains hierarchical beliefs about the traversal risks and coverage states, and (ii) performs hierarchical value learning to construct an exploration policy.




Figure 29. Hierarchical IRM generated during autonomous exploration of Valentine's cave at Lava Beds National Monument, Tulelake, CA, using a quadraped robot.

each hierarchical level, it maintains a belief about the world and robot states in a compact form, called an information roadmap (IRM), and solves for a POMDP policy to generate a coverage plan over a nonmyopic temporal horizon, in a receding horizon fashion.

Hierarchical POMDP formulation: First, we formulate the receding-horizon SMAP in Eq. (1) into a hierarchical POMDP problem (Kaelbling and Lozano-Pérez, 2011; Kim et al., 2019a). Let us decompose a belief state b into local and global belief states, $b^\ell = P(q, W^\ell)$ and $b^g = P(q, W^g)$, respectively. W^ℓ is a local, rolling-window world representation with high-fidelity information, while W^g is a global, unbounded world representation with approximate information (see Figure 29). With π^ℓ and π^g denoting the local and global policies, respectively, we approximate Eq. (1) as cascaded hierarchical optimization problems as follows:

$$\pi_{t:t+T}(b) \approx \operatorname{argmax}_{\pi^\ell \in \Pi_{t:t+T}^\ell} \mathbb{E} \sum_{t'=t}^{t+T} \gamma^{t'-t} r^\ell(b_{t'}^\ell, \pi^\ell(b_{t'}^\ell; \pi_{t:t+T}^g(b_t^g))), \quad (3)$$

where

$$\pi_{t:t+T}^g(b^g) = \operatorname{argmax}_{\pi^g \in \Pi_{t:t+T}^g} \mathbb{E} \sum_{t'=t}^{t+T} \gamma^{t'-t} r^g(b_{t'}^g, \pi^g(b_{t'}^g)). \quad (4)$$

Here $r^\ell(b^\ell, \pi^\ell(b^\ell))$ and $r^g(b^g, \pi^g(b^g))$ are approximate belief reward functions for the local and global belief spaces, respectively. Note that the co-domain of the global policy $\pi^g(b^g)$ is a parameter space Θ^ℓ of the local policy $\pi^\ell(b^\ell; \theta^\ell)$, $\theta^\ell \in \Theta^\ell$.

Hierarchical belief representation: For a compact and versatile representation of the world, we rely on a graph structure, $G = (N, E)$ with nodes N and edges E , as the data structure to represent the belief about the world state. We refer to this representation as an IRM (Aghamohammadi et al., 2014). We construct and maintain IRMs at two hierarchical levels, namely, local IRM and global IRM (see Figure 29). The local IRM is a dense high-resolution graph that contains high-fidelity information about the occupancy, coverage, and traversal risks, but locally around the




Figure 30. QMDP policy (red arrows displayed above *breadcrumb* nodes) for global coverage planning (GCP). A red sphere indicates the QMDP frontier goal.

robot. In contrast, the global IRM sparsely captures the free-space connectivity. It encodes uncovered area by so-called frontier nodes, which allow for effective representation of large environments, spanning up to several kilometers. In addition to the map uncertainty, IRM can be generalized to incorporate the robot localization uncertainty (e.g., [Kim et al., 2019b](#) or [Agha-mohammadi et al., 2018](#)) in the planning framework when traversing narrow passages and challenging environments where robot location uncertainty can hinder the robot’s ability to navigate the environment.

Hierarchical coverage Planning—GCP: Given local and global IRMs as the hierarchical belief representation of W^l and W^g , respectively, we solve the cascaded hierarchical POMDP problems. At first, a global coverage planner (GCP) solves for the global policy in Eq. (4), providing *global guidance* to a local coverage planner (LCP) of Eq. (3). The global guidance enhances the coverage performance and global completeness. It is especially helpful when the LCP has fully covered the local area and needs global guidance to move to another area. In order for the GCP to scale up to very large problems, we adopt the QMDP approach ([Littman et al., 1995](#)) (see [Figure 30](#)). The main idea is to assume the state becomes fully observable after one step of action under uncertainty, so that the value function for further actions can be evaluated efficiently in a Markov decision process (MDP) setting. This assumption is acceptable for the GCP since its main role is to guide the robot to an uncovered area and let the LCP lead the local coverage behavior. In other words, the GCP’s policy search can terminate at frontier nodes of the global IRM, and thus we can assume no changes in the coverage state during the GCP’s planning episode and adopt QMDP for efficient large-scale planning. The complexity of the GCP is $O(N_{iter}^g |N^g| |N_{nn}^g|^2)$, where N_{iter}^g is the maximum number of iterations for Bellman update, $|N^g|$ is the number of nodes on the global IRM, and $|N_{nn}^g|$ is the number of nearest-neighbor nodes for a node connected by an edge on the global IRM which is bounded to a small finite number. Thus, the complexity of the GCP grows only linearly with $|N^g|$.

Hierarchical coverage planning—LCP: In the hierarchical optimization framework, the LCP solves Eq. (3), given a parameter input from the GCP. The LCP constructs a high-fidelity policy by considering the information gathering (with visibility check given obstacles), traversal risk (based on proximity to obstacles, terrain roughness, and slopes), and robot’s mobility constraints (such as acceleration limits and nonholonomic constraints of wheeled robots). The LCP has two phases: (i) reach the target area based on the GCP’s guidance and (ii) construct a local coverage path after reaching the target area. In the first case, when the target frontier is outside the local IRM range, the LCP instantiates high-fidelity motion commands to reach the target frontier. In the second case, when the target frontier is within the local IRM range, then the LCP performs the information-gathering coverage optimization, as described in Eq. (3). In order to solve the coverage optimization problem we employ the partially observable Monte Carlo planning (POMCP) algorithm ([Silver and Veness, 2010](#); [Kim et al., 2019b](#)). POMCP is a widely adopted POMDP solver that leverages the Monte Carlo sampling technique to alleviate both the curses of dimensionality and




Figure 31. Illustrative example of coverage path planning on the local IRM with Monte Carlo tree search. The field of view of the robot's coverage sensor is represented by a blue circle. Macro actions (six steps on local IRM in this example) associated with the two tree branches, paths A and B, are shown. Note that the final world coverage states in both branches are identical. Path A is evaluated to be more rewarding than B since fewer actions were required to cover the same area.

history (see, e.g., Figure 31). Given a generative model (or a black box simulator) for discrete action and observation spaces, POMCP can learn the near-optimal value function of the reachable belief space with adequate exploration-exploitation trade-off. We limit the complexity of the LCP process to $O(N_{iter}^\ell N_{depth}^\ell)$ of calling the generative model, where N_{iter}^ℓ is the number of iterations for episodic forward simulation and N_{depth}^ℓ is the depth of planning horizon for each iteration. Since it is a local rolling-window planner, there is no increased complexity with the total size of the environment.

9.3. Experimental Evaluation

In order to evaluate our proposed framework, we perform high-fidelity simulation studies with a four-wheeled vehicle (Husky robot) and real-world experiments with a quadruped (Boston Dynamics Spot robot). Both robots are equipped with custom sensing and computing systems, enabling high levels of autonomy and communication capabilities (Otsu et al., 2020). The entire autonomy stack runs in real time on an Intel Core i7 processor with 32 GB of RAM. The stack relies on a multi-sensor fusion framework. The core of this framework is 3D point-cloud data provided by LiDAR range sensors mounted on the robots (Ebadi et al., 2020). We refer to our autonomy stack-integrated Spot as Au-Spot (Bouman* et al., 2020).

Baseline algorithms: We compare our PLGRIM framework against a local coverage planner baseline (next-best-view method) and a global coverage planner baseline (frontier-based method).

1. *Next-Best-View (NBV)*: NBV is a widely adopted local coverage planner that returns a path to the best next view point to move to. It uses an information gain-based reward function as ours but limits the policy search space to a set of shortest paths to sampled view points around the robot. While NBV is able to utilize local high-fidelity information, it suffers from spatially limited world belief and sparse policy space.
2. *Hierarchical Frontier-based Exploration (HFE)*: Frontier-based exploration is a prevalent global coverage planning approach that interleaves moving to a frontier node and creating new frontiers until there are no more frontiers left (e.g., Umari and Mukhopadhyay, 2017). It optimizes for the global completeness of environment exploration but often suffers from local suboptimality due to its large scale of the policy space and myopic one-step look-ahead decision making. The performance of frontier-based methods can be enhanced by modulating the spatial scope of frontier selection, but it still suffers from downsampling artifacts and a sparse policy space composed of large action steps.

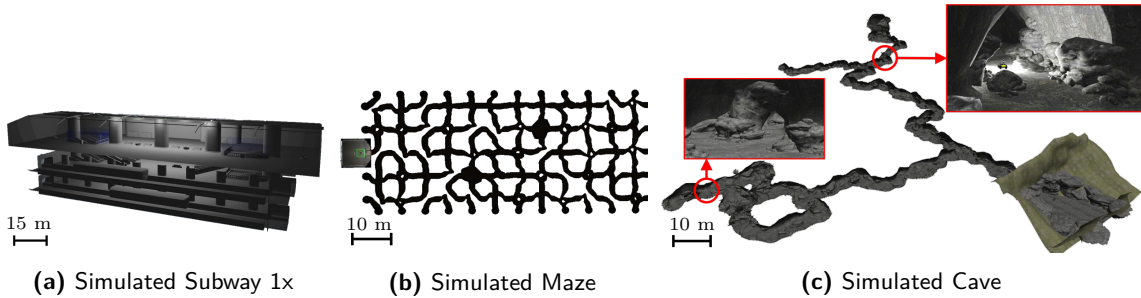


Figure 32. Simulated environments for performance validation: (a) subway station, (b) maze (top-down view), and (c) cave.

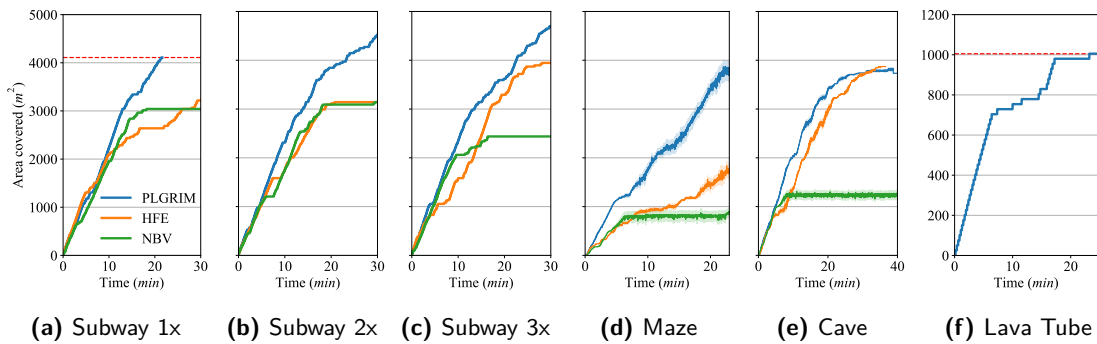


Figure 33. Exploration by PLGRIM and baseline methods in (a)–(c) simulated subways of increasing size, simulated (d) maze and (e) cave, and (f) real-world lava tube. For (d) and (e), the covered area is the average of two runs. Red dashed lines indicate 100% coverage of the environments, wherever applicable.

9.3.1. Simulation Evaluation

We demonstrate PLGRIM’s performance, as well as that of the baseline algorithms, in simulated subway, maze, and cave environments. Figure 32 visualizes these environments. In our comparisons, in order to achieve reasonable performance with the baseline methods in complex simulated environments, we allow baseline methods to leverage our local and global IRM structures as the underlying search space.

Simulated subway station: The subway station consists of large interconnected, polygonal rooms with smooth floors, devoid of obstacles. There are three varying sized subway environments, whose scales are denoted by $1\times$, $2\times$, and $3\times$. Figure 33(a)–(c) shows the scalable performance of PLGRIM against the baselines. In a relatively small environment without complex features (Subway $1\times$), NBV performance is competitive as it evaluates high-resolution paths based on information gain. However, as the environment scale grows, its myopic planning easily gets *stuck* and the robot’s coverage rate drops significantly. HFE shows inconsistent performance in the subway environments. The accumulation of locally suboptimal decisions, due to its sparse environment representation, leads to the construction of a globally inefficient IRM structure. As a result, the robot must perform time-consuming detours in order to *pick up* leftover frontiers.

Simulated maze and cave: The maze and cave are both unstructured environments with complex terrain (rocks, steep slopes, etc.) and topology (narrow passages, sharp bends, dead ends, open spaces, etc.). The coverage rates for each algorithm are displayed in Figure 33(d)–(e). PLGRIM outperforms the baseline methods in these environments. By constructing long-horizon coverage paths over a high-resolution world belief representation, PLGRIM enables the robot to safely explore through hazardous terrain. Simultaneously, it maintains an understanding of the global world, which is used when deciding where to explore next after exhausting all local information. In the cave,




Figure 34. The local IRM (yellow, brown, and white nodes represent uncovered, covered, and unknown areas, respectively) is shown overlaid on the Riskmap. A yellow arrow indicates the robot's location. LCP plans a path (red) that fully covers the local area (snapshot A). When $p(W^\ell)$ updates, the path is adjusted to extend towards the large uncovered swath while maintaining smoothness with the previous path. Another $p(W^\ell)$ update reveals that the path has entered a hazardous area—the wall of lava tube (snapshot B). As a demonstration of LCP's resiliency, the path shifts away from the hazardous area, and the robot is redirected towards the center of the tube (snapshot C). One minute later, the robot encounters a fork in the cave. The LCP path curves slightly toward the fork apex (for maximal information gain) before entering the wider, less-risky channel (snapshot D).

NBV's reliance on a deterministic path, without consideration of probabilistic risk, causes the robot to drive into a pile of rocks and become inoperable. NBV exhibits similarly poor performance in the maze. However, in this case, NBV's myopic planning is particularly ineffectual when faced with navigating a topologically complex space, and the robot ultimately gets *stuck*. As was the case in the subway, HFE suffers in the topologically complex maze due to the accumulation of suboptimal local decisions. In particular, frontiers are sometimes not detected in the sharp bends of the maze, leaving the robot with an empty local policy space. As a result, the robot cannot progress and spends considerable time backtracking along the IRM to distant frontiers.

9.3.2. Real-World Evaluation

We extensively validated PLGRIM on physical robots in real-world environments. In particular, PLGRIM was run on a quadruped robot in Valentine lava tube, located in Lava Beds National Monument, Tulelake, CA. The cave consists of a main tube, which branches into smaller, auxiliary tubes. The floor is characterized by ropy masses of cooled lava. Large boulders, from ceiling breakdown, are scattered throughout the tube. Figures 34 and 35 discuss how PLGRIM is able to overcome the challenges posed by large-scale environments with complex terrain and efficiently guide the robot's exploration. Figure 33(f) shows the area covered over time.

10. Multi-Robot Networking

Multi-robot systems offer advanced capabilities to enable complex and time-constrained missions in large-scale complex environments. Resilient wireless mesh networking solutions are a prerequisite for reliable and efficient multi-robot missions. NeBula is inherently a “networked” solution (Figure 7). While it can be implemented on a single autonomous robot, it also allows for faster and more efficient missions with multiple potentially heterogeneous robots (see Figure 5). NeBula's goal in the SubT challenge is to map an unknown subterranean environment, locate artifacts, and communicate that information to the base station via a wireless mesh network for submission to the DARPA Command Post. Inter-robot wireless communication in subterranean environments is particularly challenging and uncertain in the reliability, capacity, and availability of communication links because of (i) limited line-of-sight opportunities, (ii) the complicated nature of the interaction of radio signals

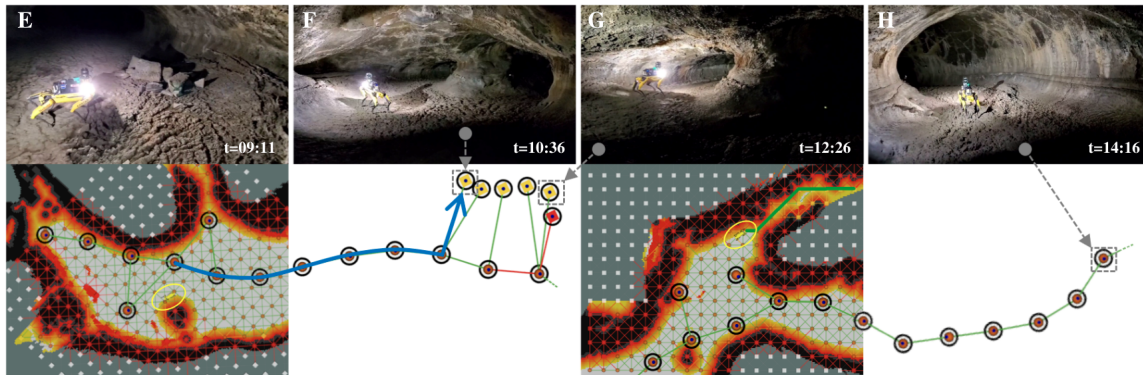


Figure 35. Portions of the global IRM constructed in the lava tube are visualized—yellow nodes represent frontiers, brown nodes represent breadcrumbs. Gray arrows associate a frontier with a snapshot of the robot exploring that frontier. GCP plans a path (blue) along the global IRM to a target frontier after the local area is fully covered (snapshot E). The robot explores the area around the frontier (snapshot F), and then explores a neighboring frontier at the opening of a narrow channel to its right. LCP plans a path (green) into the channel (snapshot G). Later, after all local areas have been explored, the robot is guided back towards the mouth of the cave along the breadcrumb nodes (snapshot H).

with the environment (e.g., reflecting, scattering, multipath), and (iii) the unknown nature of the environment. In this section, we will go over NeBula’s Collaborative High-bandwidth Operations with Radio Droppables (CHORD) communication system for comm-degraded subterranean environments. The objective of CHORD is to maintain high-bandwidth links to multiple robots for efficient commanding, autonomous operation, and data gathering in complex unknown environments. For more details on the development of NeBula’s networking solutions, see [Otsu et al. \(2020\)](#) and [Ginting et al. \(2021\)](#).

10.1. CHORD System Design

Architecture and ConOps: NeBula’s operations in the SubT challenge consist of three general types of agents: (1) static agents (e.g., base station), (2) mobile agents (e.g., robots), and (3) deployable static agents (e.g., communication relay nodes). Each agent communicates by means of a wireless mesh network using commercial-off-the-shelf radios. We use a hybrid of ROS 1 and ROS 2 for the communication middleware ([Ginting et al., 2020, 2021](#)). We use ROS 1 for intra-robot communications, and ROS 2 for inter-robot communications. The mesh network can be extended, as shown in [Figure 36](#), into the subterranean environment by deploying communication nodes ([Section 12.2](#)) from robots to build a backbone wireless mesh network. The decision of where to drop communication nodes is based on the 3D map, data route, signal-to-noise ratio, and the estimated available bandwidth between each radio ([Vaquero et al., 2020](#)). The exact coverage area of each node is dependent on many factors including surface materials, roughness, and environment shape. Dropping communication nodes in range of another node with a route to the base station reduces the uncertainty of getting data back from robots near that node. Robots may also be used to extend the mesh network when they are in communication range of another asset (robot or communication node) with a route to the base station. When a robot needs to send data to the base station, but has no communication route to the base station, it can either return to the communication range of an asset that does have a route to the base station (usually by backtracking to a node on the backbone network) or communicate those data to another robot (referred to as a data mule) that is returning.

Intra-robot communication: Each robot consists of a combination of computers and sensors (see [Section 12](#)) connected by Gigabit Ethernet. Where possible (current generation Husky hardware) this network is separate from the radio network and connects only through a single computer which is connected to a radio. That computer runs the ROS 1 core, the ROS 1-2 bridge, and is

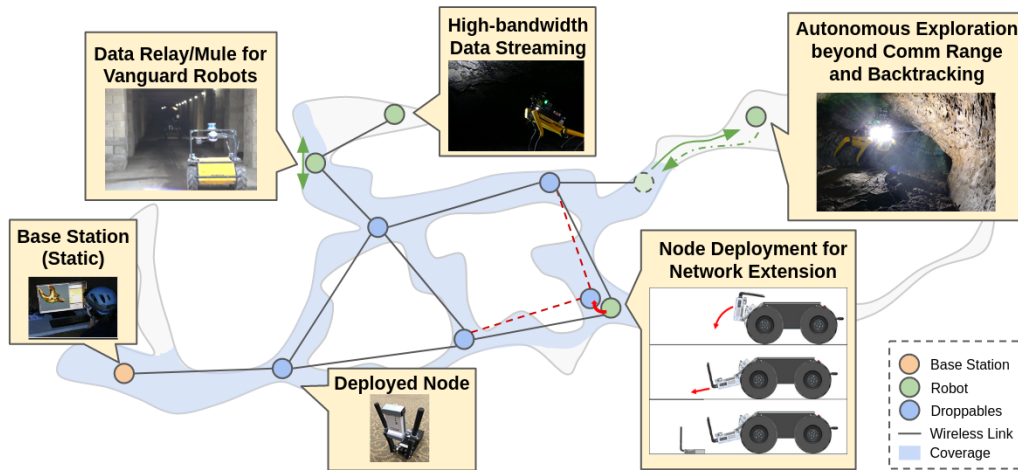


Figure 36. This figure shows how the mesh network is extended into the subterranean environment by deployed communication nodes which form a backbone network. Robots inside the coverage area of the backbone network can extend the network to robots in their communication range. Robots without a route through the network to the base station must be fully autonomous and return to an area with communication coverage to return data. Alternatively (not pictured), multiple robots without a wireless route to the base station may share data and have one act as a data mule to carry the data back to communication range.

responsible for handling inter-agent communication. This further isolates intra-robot communication from inter-agent communication and prevents inadvertent radio traffic. Intra-robot communications are monitored using ROS 1 topic statistics.

Inter-agent communication: Even with relatively high power/bandwidth radios, bandwidth is still a shared, limited, and temporary resource. Efforts must be taken to manage bandwidth usage and be robust to communication loss when robots operate outside the range of the radios. CHORD uses ROS 2 over the wireless mesh network for inter-agent communication. The advanced quality of service (QoS) features of ROS 2 are used to guarantee delivery of important priority data while maintaining network stability over low-bandwidth links. This configuration enables traffic prioritization and resource control. We configured two categories of QoS for inter-robot topics with different mission requirements. Topics that require full message history transfer for post-processing or that deliver mission-critical information have higher priority and are configured so that the messages are reliably delivered even though the network may be down for periods of time. Estimated link capacities, data routes, and data traffic are also monitored to ensure stability.

10.2. Evolution of CHORD

Tunnel Circuit: During the Tunnel competition we used ROS 1 in combination with a custom cross-master messaging mechanism (multimaster-JPL) in combination with radios from Silvus Technologies and Persistent Systems (Otsu et al., 2020). While our communication system and radios supported our six-robot operation well during the tunnel competition, we observed some communication issues. We found that careful attention was needed to avoid ROS 1 attempting to share global topics (like TFs and diagnostics) across all robots. In addition, without better QoS controls, robots outside of the communication range of the backbone network would buffer all data and flood the network on their return. Some of our data products were also larger than expected. See Otsu et al. (2020) for more details.

Urban Circuit: Before the Urban competition we switched to using ROS 2 for inter-agent communication in combination with radios from Silvus Technologies. We observed better performance than our previous ROS 1-only data sharing system. First, by using a different protocol for

inter-robot communication, we isolated the ROS 1 networks completely and avoided unintended data flows between agents. The network isolation also helped to diagnose network issues easily as every inter-robot ROS topic passes through the bridge node. Second, we were able to keep network traffic inside our bandwidth budget, which contributed to the stability of the dynamic network. For more details on the results, see [Ginting et al. \(2021\)](#).

11. Mission Planning and Autonomy

Having the capability to autonomously plan, reconfigure, and perform tasks for a multi-robot system is a crucial component of the NeBula autonomy framework (see [Figure 7](#)), enabling exploration of large, complex, and unknown environments. Especially in the context of the SubT challenge, when there is none or unreliable intermittent communication between robots and a single human supervisor, autonomy is crucial to achieve mission objectives, under time and resource constraints. In this section, we present NeBula’s mission autonomy components, while integrating and allowing a single human supervisor to oversee and interact with a team of more than five heterogeneous robots at the same time, under range-limited and unreliable communication in challenging environments. For technical details, please see [Otsu et al. \(2020\)](#), [Vaquero et al. \(2020\)](#), and [Kaufmann et al. \(2021\)](#).

11.1. Architecture

[Figure 37](#) illustrates the components of the mission autonomy architecture and their interface to components. In the following paragraphs, we describe key components of the system’s mission planning and autonomy.

- *Mission Executive*: The executive is responsible for stepping through the mission flow as defined and specified in the mission file (see next point). The executive triggers the robot’s autonomy behaviors based on its mission state.
- *Mission File*: The mission file defines a set of mission autonomy behaviors and their execution flow during a mission.
- *Autonomy Behaviors*: Mission autonomy behaviors implement the logic to interact with world beliefs and send commands to other modules or platforms. See [Section 11.2](#) for specific behavior implementations.

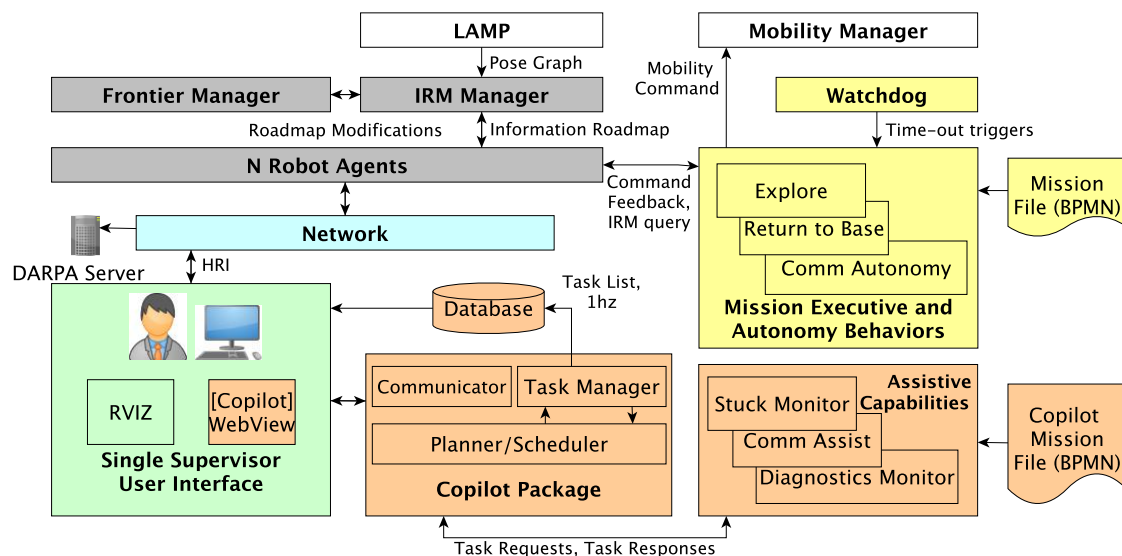


Figure 37. Mission autonomy architecture and components (robot side, yellow; base side, orange).

- *IRM Manager*: The IRM is a key data structure to exchange information between humans and machines. Certain robot autonomy behaviors and assistive capabilities utilize the latest belief states in the IRM within autonomous decision-making and planning processes.
- *Mission Watchdog*: The mission watchdog monitors mission progress and communication links. It is an external mechanism that ensures that the robots can transmit their world belief back to the base station regardless of the mission executive states.
- *Copilot and Assistive Behaviors*: The mission autonomy assistant, Copilot for short, encompasses a series of monitoring and assistive capabilities (e.g., system health monitor, comm node drop assistance) that perform autonomous tasks and keep the human in the loop if possible and needed. The details of Copilot, including its assistive capabilities, planning and scheduling, and human interaction via user interfaces, are explained in [Section 11.3](#).
- *User Interface*: The user interface consists of carefully designed web-based components and RViz displays for increased situational awareness. It provides an efficient interface to send various levels of commands with minimal control.

11.2. Autonomy

Exploring unknown and complex environments with a team of multiple robots comes with several challenges. The difficulty of operating a single robot under limited available communications bandwidth, let alone a team, motivates the need for full autonomy during different phases of an exploration mission when communications are not available. In situations where robots are outside of each other's communication range, robots must remain fully autonomous and independently reason about their environment to determine their next task. A subset of this problem arises when robots are within each other's communication range and the robots must devise a coordinated plan. This section describes how we utilize various levels of mission autonomy.

Planning and scheduling: Mission planning and task scheduling constitute the highest layer of the planning modules and are integral components in achieving full autonomy outside and within a communication range ([Figure 4](#)). The mission planner maintains a world belief, as described in [Section 3](#), which comprises the state of the robot team (e.g., robot health, robot location, detected artifacts on each robot), the state of the world (e.g., geometric and semantic maps), the state of the mission (e.g., remaining mission time, margin to desired mission output), and the state of communication (e.g., network connectivity, location of comm nodes). As the world belief increases and improves, the mission planner dynamically retasks robots to new goals or deploys new robots in order to achieve the mission goals defined in the mission specification. Ongoing work focuses on using semantic information about the world (e.g., stair wells, door frames, intersections, room volume) in the mission planner to provide additional belief of where critical information may lie to help improve mission success. The task scheduler works in conjunction with the mission planner and robot autonomy behaviors to actively schedule and assign agents for each task given various constraints. This is one of the mission autonomy features that allows the system to actively deal with temporal uncertainties, dependencies, dynamic resource constraints, and varying risks and mission strategies. The scheduling component is modular and can be interchanged with a variety of existing planning and scheduling frameworks. Currently, a planning problem is formulated using the Planning Domain Definition Language ([Fox and Long, 2002](#)) and is solved using OPTIC ([Benton et al., 2012](#)). The solver updates the task schedule at a fixed cadence.

Executive: The executive is a task manager that ensures each scheduled task is dispatched to each respective agent at the correct time. It tracks the state of all tasks and requests the task scheduler to reschedule when tasks fail, or relaxes temporal constraints when the current schedule is infeasible.

Mission specification: A mission file is used to describe a mission and combines several high-level robot autonomy behaviors to create the flow for more complex scenarios. For this, we use the Traceable Robotic Activity Composer and Executive (TRACE) proposed in ([de la Croix et al., 2017](#)) and the Business Process Modelling Notation (BPMN). [Figure 38](#) shows one example "Exploration"

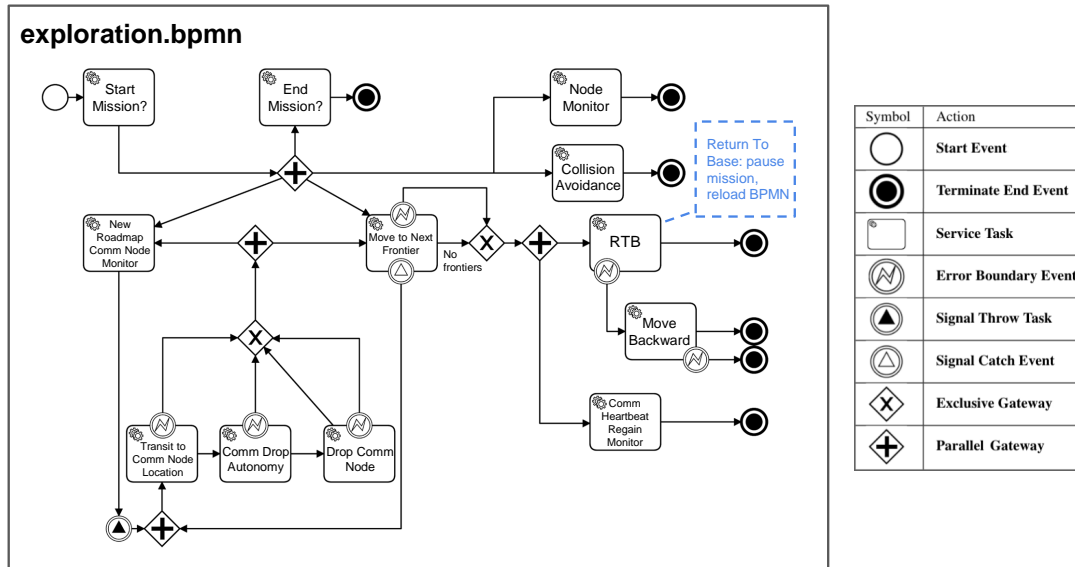


Figure 38. Annotated Exploration mission file used during the self-organized cave circuit.

mission file that was used in one of our Cave exploration scenarios. The mission starts off in the upper-left corner with a prompt for the user to begin the mission which ensures the robots can start moving autonomously in a safe manner. Each rectangular box depicts a service task—they represent robot behaviors in our architecture. Once the mission is started, the mission flow moves towards the parallel gateway. This gateway allows multiple flows to branch off of a single flow. In the case of the exploration BPMN, this allows five behaviors to run in parallel.

Robot autonomy behaviors: Behaviors used in the example mission file (Figure 38) represent a subset of NeBula’s robot autonomy behaviors; some are described here:

- **Move to Next Frontier** - Receives the current frontier as a goal from the global planning module and commands the agent to move to it.
- **New Roadmap Comm Node Monitor** - Monitors the IRM for new mesh network extension requests and interrupts the **Move to Next Frontier** behavior temporarily to initiate **Comm Drop Autonomy**.
- **Comm Drop Autonomy** - Autonomous selection of optimal target location to drop a communication node to maximize communication coverage while minimizing the risk of violating safety and operational constraints for the robots traversing the local environment (Vaquero et al., 2020).
- **Transit to Comm Node Location** - Commands the agent to move towards the dropping location through the IRM.
- **Drop Comm Node** - Instructs the dropper firmware to drop a comm node immediately.
- **Collision Avoidance** - Prevents inter-robot collision by monitoring inter-robot distances. When robots are too close, this behavior performs a prioritized motion planning to resolve the situation.
- **Comms Heartbeat Regain Monitor** - Continuously monitors for heartbeat messages from the base station and terminates successfully when a consistent stream of messages is detected.
- **RTB (Return to Base)** - Locates the agent within the current IRM and finds the shortest path to the base IRM node. Then sends this path to the mobility manager and terminates successfully when the agent has reached the base.
- **Stairs Helper** - Detects stairs and assists during a stair climbing procedure.

11.3. Assistive Autonomy and Human-Robot Interaction

In situations where there is sufficient communication bandwidth to interface with a team of multiple robots, having the capability to provide situational awareness of all robot activities and the mission progress to the operator becomes very useful. Assistive autonomy at the base station facilitates this human-robot interaction especially under limitations like the single-supervisor requirement of the SubT challenge, or available cognitive workload.

Autonomy Copilot MIKE: The Multi-robot Interaction assistant for unKnown cave Environments (Copilot) is introduced in Kaufmann et al. (2021) and supports the single human supervisor during the setup and mission phases of complex multi-robot operations. Copilot treats the human supervisor as one node (or a member) of the multi-robot system and it actively schedules and reschedules the tasks for all members while considering world beliefs, resources, and other constraints (e.g., human cognitive workload, available comm nodes, etc.). Some tasks can be automatically executed or resolved depending on the allowed autonomy level. Currently, Copilot comprises assistive capabilities, robot behaviors, assistive task scheduling, and user interfaces (Figure 37) used to autonomously control the robots and guide the human supervisor's decision making while keeping them in the loop if a communication link is established.

Assistive capabilities: Assistive capabilities are distributed across the robots and the base station to reduce the supervisor's workload. These capabilities assist to (1) detect system anomalies, (2) create tasks for the Copilot core module, (3) monitor the resolution of these tasks, and (4) resolve tasks autonomously or inform the supervisor that actions are needed. The assistive capabilities run independently without knowing the complete system status, which allows a modular extension of the autonomy functionality. Assistive capabilities address human limitations and autonomously intervene during mission execution as needed and when possible. Tasks created by these capabilities can inform the supervisor of the status of communication links between the robots and the base (e.g., communication loss, communication node drop suggestions), a robot's sensor suite status or mobility status (e.g., stuck, tilted, undesirable oscillations), and the system's autonomy software status (e.g., process health).

Assistive task scheduling: The assistive task scheduler shares the implementation of the task scheduler mentioned prior and considers the operator(s) as an agent (or resource) to which tasks can be assigned. As more tasks become automated and can be assigned to the Copilot, the operator becomes more readily available to assist with mission-critical tasks.

User interface: The user interface consists of several web-based components and a configurable RViz view showing a 3D or 2.5D map representation of the explored environment (Otsu et al., 2020). Figure 39 illustrates the single screen interface which includes the Copilot and robot status components. The *Copilot component* is designed to increase situational awareness, inform the supervisor of mission tasks, and dispatch urgent system notifications. The limited interaction between the system and the single human supervisor increases the supervisor's trust in the system's performance, as they are kept in the loop and may still intervene if and when a communication link is established and if the mission strategy requires a change (e.g., taking a more risky, more conservative posture). The user interface provides an additional *robot status component* for each robot to directly interact with agents or conduct troubleshooting. Figure 39 gives an example of this component which is used when communication links exist; the human supervisor needs to shift their attention to a single robot to directly interact with it. This component comprises elements to reflect real-time housekeeping data, such as current position, sensor health status, estimated battery remaining, and remaining comm nodes. It also integrates controls for changing a robot's role to either act as a vanguard explorer or provide support to the leading agents in the mission. Lastly, the component integrates the current BPMN mission diagram and highlights an agents' active autonomy behavior to reflect the robot internal state of the mission executive.

Human trust in autonomy: Copilot has been deployed in multiple field tests and mission simulations in diverse environments providing us with useful feedback regarding the interface design and multi-robot operations. Previous to the implementation of Copilot, we learned that our human

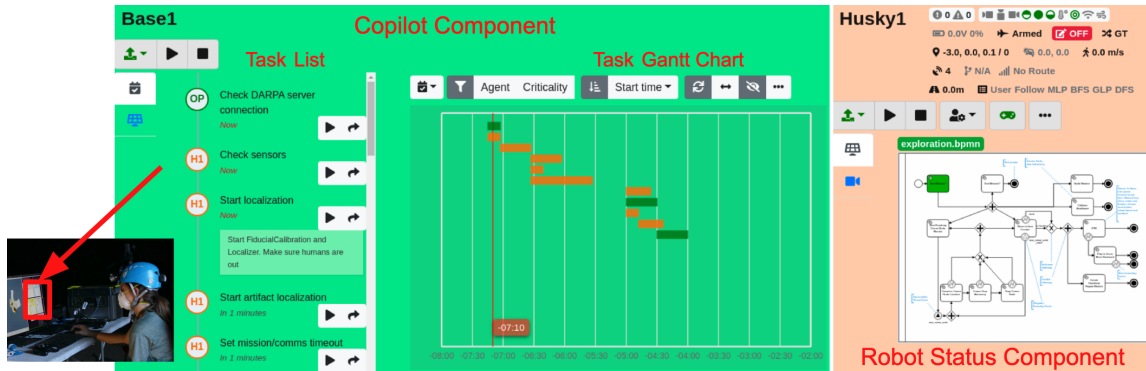


Figure 39. Base station single screen setup and human supervisor interacting with the user interface (left). Autonomy assistant Copilot MIKE's interface component with task list and Gantt chart mission overview (middle). Robot status component displaying the robot's state in the exploration mission (compare Figure 38), real-time information regarding sensor hardware, current pose, and other housekeeping data that represent current beliefs (right).

supervisors were overwhelmed with remembering tasks instead of strategically overseeing the robots' activities and overall mission progress. Now our human supervisors let the Copilot handle several decision-making processes and must trust its autonomy capabilities (Kaufmann et al., 2021), even if a supervisor could make their own decisions.

12. Mobility Systems and Hardware Integration

As described in Section 3, to simultaneously address various challenges associated with exploring unknown challenging terrains, we rely on a team of heterogeneous robots with complementary capabilities in mobility, sensing, computing, and endurance. These assets are deployed in the mission as the robots learn about the terrain specifications. Figure 40 shows four classes of our robots: (i) wheeled rovers, to cover general and relatively smooth surfaces and mild obstacles; (ii) legged robots, to cover more challenging and uneven terrains where surmounting obstacles or staircases are required (Bouman* et al., 2020; CoSTAR Team, 2020a; Miller et al., 2020; Jenelten et al., 2020); (iii) tracked robots, to complement legged platforms in handling different surface material; and (iv) aerial and hybrid locomotion, to enable traversing vertical shafts, and areas that are not accessible by ground robots.

12.1. Ground Robots

Our ground robots are able to carry heavy payloads. Hence they are equipped with high levels of sensory and processing capabilities enabling complex autonomous behaviors and artifact detection. Their battery capacity allows them to have longer operational time than flying vehicles. Below, we discuss the NeBula Autonomy Payload on these ground robots.

Architecture: The NeBula payload (Figure 41) consists of the NeBula Sensor Package (NSP), NeBula Power and Computing Core (NPCC), NeBula Diagnostics Board (NDB), and NeBula Communications Deployment System (NCDS). Its electronics and software architecture is modular, to facilitate adaptation to varying mechanical and power constraints of each platform in our heterogeneous robotic fleet.

Design principles: The key design principles are as follows:

- *Durability:* Shock proofing of the system increases longevity and self-recovery chances when exploring challenging terrains.
- *Lightweight materials:* Payload reduction maximizes robot agility and battery life.







Topology	Wheeled Vehicles Husky 	Legged Platform Spot 	Hybrid Vehicles Rollercopter, Drivcopter 
High-level Computers	SLAM: Multiple, High-power CPU Object Recognition: NVIDIA GPU	SLAM: Mid-power CPU Object Recognition: NVIDIA GPU	SLAM: Lightweight SB-CPU's Object Recognition: NVIDIA GPU
Low Level Microcontrollers	Safety features: Atmel MCU Comm Dropper: STM MCU	Safety features: Atmel MCU Comm Dropper: STM MCU	Safety features: Atmel MCU
Sensors	LiDAR, depth and thermal cameras, gas and signal sensors, external IMU.	LiDAR, depth and thermal cameras, gas and signal sensors, external IMU.	LiDAR, depth and color cameras, IMU.
Comm Dropper	Yes	Yes	No
Mass	50kg	40kg	1-5kg
Battery	Lead acid	Lithium ion	Lithium polymer
Operational time	3 hours	1.5 hours	<60 minutes
Topology	High-speed Data Mules Balto 	Tracked Vehicles Telemax 	Aerial Vehicles Quadrotor, Coax-octocopter 
High-level Computers	SLAM: Single mid-power CPU Object Recognition: NVIDIA GPU	SLAM: Single high-power CPU Object Recognition: NVIDIA GPU	SLAM and object detection: Lightweight low-power CPU
Low Level Microcontrollers	Safety features: Atmel MCU	Safety features: Atmel MCU Comm Dropper: STM MCU	Safety features: Atmel MCU
Sensors	LiDAR, depth and thermal cameras, gas and signal sensors, external IMU.	LiDAR, depth and thermal cameras, gas and signal sensors, external IMU.	Depth and color cameras, IMU, 2D lidars
Comm Dropper	No	Yes	No
Mass	10kg	75kg	0.5-3kg
Battery	Lithium Polymer	Lead acid	Lithium Polymer
Operational time	2 hours	3 hours	<20 minutes

Figure 40. Hardware specifications for some of the CoSTAR robots.

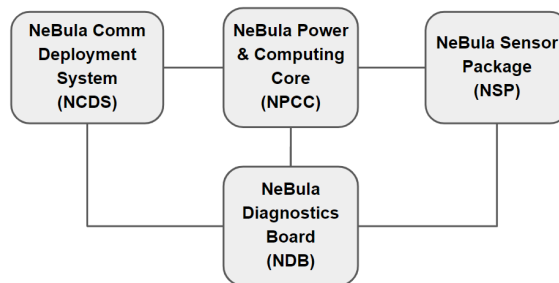


Figure 41. High-level overview of the “NeBula Autonomy Payload” architecture.

- *Modularity:* A critical feature when developing a heterogeneous fleet of robots with various capabilities is modularity. For example, processors and sensors are adaptable across different mobility systems with differing mass/size constraints.
- *Portability:* Ease of transportation minimizes damage to valuable hardware.
- *HW-autonomy co-design:* Adaptable design process enables reconfiguring sensors to adapt to autonomy evolution.



Figure 42. NSP equipped on Husky (left) and Spot (right).

NeBula Sensor Package (NSP): The NSP empowers the NeBula autonomy solution on the robots by gathering sensory information in real time from the environment. NSP is heterogeneous; on each robot NSP consists of a subset of the following sensors: LiDARs, monocular, stereo, and thermal cameras, external IMUs, encoders, contact sensors, ultra-high-lumen LEDs, radars, gas sensors, and UWB and wireless signal detectors. NSP is protected by custom superstructures with impact protection. A combination of hard resin urethane, semi-rigid carbon-infused nylon, and aluminum was used. The NSP interfaces with the NPCC via high-bandwidth USB and Ethernet for data and custom serial messages and push-pull connectors for high-amp power. One version is depicted in [Figure 42](#), which shows NSP payloads utilizing an array of Velodyne Puck VLP-16, Intel Realsense d435i, and FLIR Boson 640, among several others; the sensors are interchangeable and reconfigurable to accommodate different sensor arrangements, manufacturers, and sensing modalities.

NeBula Power and Computing Core (NPCC): The NPCC is an auxiliary payload which provides power to all NeBula sensors and computers used for autonomy. Aluminum enclosures provide protection to the internal electronics in the event of atypical loads and impacts such as falls and collisions. It is designed with considerations for thermal cooling and haptics due to extensive cycling of the connector interface panel. The modular, auxiliary-mount design was tweaked to accommodate for the reduced flight weight of the drones. The NPCC is powered from an external lithium high-capacity battery to provide isolation and extended battery life for the internal battery across the heterogeneous fleet. The payload uses two high-power computers for sensing, autonomy, and semantic scene understanding and also hosts the low-level microprocessor. On some robots, the NPCC is equipped with a GPU-based system-on-module with a custom interface to the power modules, cameras, and sensors to accommodate machine learning and semantic understanding functionalities. The various configurations of the NPCC can be seen in [Figure 43](#).

NeBula Diagnostics Board (NDB): The NDB implements the system diagnostics, which monitors the vital power elements of the robot such as battery voltage, input current, and individual regulator voltages. When the robot boots initially, all voltage regulators are powered up in a staggered sequence to limit the inrush loads to the NDB. After each voltage regulator is enabled, the processor checks that the voltage is within the expected range and reports errors if any are found. In addition, the current monitoring checks for high current draw when each of the regulators is enabled to detect possible short circuits in the various robot sensors and mechanisms. The power module has an input protection circuit to protect against voltage transients, reverse polarity, and undervoltage. A custom ROS message combines all these measurements and is constantly publishing the hardware power status to a specific robot diagnostic topic using the rosserial interface.




Figure 43. NeBula Power and Computing Core. Wheeled/tracked (left), legged (middle), and aerial vehicles (right).



Figure 44. NCDS equipped on Husky (left) and Spot (right).

NeBula Comm Deployment System (NCDS): As discussed in our ConOps in [Section 3](#), we construct and expand a wireless mesh network near the environment entrance, to extend the reach of the base station. In order to do so, ground robots are equipped with a comm deployment system (NCDS), which allows them to carry and deploy communication radios (comm nodes) and static assets during the mission. The radios are autonomously deployed using the NCDS which mounts at the rear of the robots seen in [Figure 44](#). The comm nodes are encased in aluminum until release, and the NCDS electronics are mounted locally on the mechanism and sealed for ingress protection. The core protection and brake-lock/release mechanism was modularized across the fleet though the deployment capacity was reduced for the legged and tracked vehicles due to available mounting points, payload constraints, and sensor occlusion. It is driven by a geared brushless motor and activated via a custom embedded system communicating over rosserial. A ROS message provides the status of each radio (loaded, dropped). The NCDS circuit board interfaces with the NPCC via serial communication. It monitors the hard stop switches used for calibration and the switches responsible for detecting if radios are loaded/released. A high-level representation is depicted in [Figure 45](#).

The NCDS software architecture relies on a finite state machine (FSM), with the following representative activities:

- *Start up:* Checks for battery power. Enables motor power and establishes links with motors.
- *Calibration:* Servos is actuated until limit switches are contacted. Encoder positional values are then stored in local memory.
- *Load/Unload:* Manual switches allows the user to load and swap radios.
- *Deploy Radio:* Perform radio deployment. Use sensors to detect if deployment was successful.
- *Shutdown:* Disables motor power.

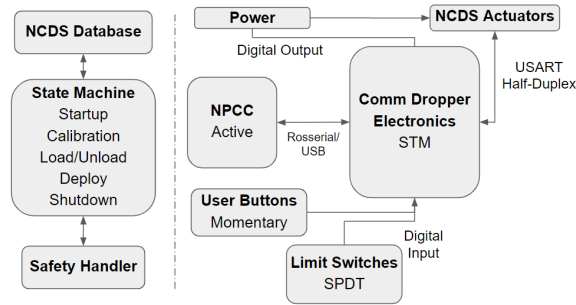


Figure 45. NCDS software architecture (left) and NCDS hardware architecture (right).



Figure 46. Static asset (comm node and UWB).

12.2. Static Assets

Static assets refer to hardware solutions that are not capable of changing their position after deployment. Our static assets are composed of comm nodes and UWB modules as pictured in [Figure 46](#). The comm nodes construct NeBula’s mesh network to enable more efficient inter-robot information exchange, while the UWB provides auxiliary landmarks and provides ranging measurement to assist the SLAM and global localization modules, as explained in [Section 6](#). Static nodes can be powered prior to the start of a mission or prior deployment. They are carried by mobile robots and deployed by the NCDS.

12.3. Flying Vehicles

NeBula has been implemented on several flying robots of varying sizes. Flying robots are responsible for exploring vertical shafts and areas not accessible with ground robots, or relaying data from the ground robots to the base station by quickly flying to regions with strong communication connectivity, e.g., for data muling (see [Section 10](#)). Their processors and sensing capability are much more constrained than our ground robots. The concept behind the drone development is to keep a balance between the payload, the size of the drone, and the endurance. Thus, an iterative design has been performed in order to conclude in an optimal hardware configuration ([Jung et al., 2021](#)). [Figure 47](#) shows several NeBula-powered aerial vehicles:

A specific example is shown in [Figure 48](#), a custom drone extensively utilized in this project that carefully balances speed, weight, autonomy capability, and flight time. The vehicle’s weight is 1.5 kg and provides 12 minutes of flight time. A 2D rotating RPLiDAR A2 is mounted on top of the vehicle, providing range measurements at 10 Hz and a monocular visual sensor at 30 FPS. The velocity estimation is based on the PX4Flow optical flow sensor at 20 Hz, while the height measurements are provided by the single beam LiDAR-lite v3 at 100 Hz, both installed on the




Figure 47. Flight time of similarly sized vehicles with different autonomy capabilities. The general trend for aerial robots is that long flight times are possible at the expense of large vehicles with less access and agility.




Figure 48. Drone equipped with a forward looking camera, LED light bars, optical flow, 2D LiDAR, and single beam LiDAR looking upward and downward.

bottom of the vehicle pointing down as indicated in [Figure 48](#). Furthermore, the aerial platform is equipped with two 10 W LED light bars in both front arms for providing additional illumination for the forward looking camera and four low-power LEDs pointing down for providing additional illumination for the optical flow sensor.

12.4. Hybrid

To extend the range of the flying drones, our team has been investigating and designing new hybrid ground/aerial vehicles, referred to as “rollocopters.” Rollocopters are designed to mainly roll on the ground; when rolling is not possible, they can fly, negotiate a nonrollable terrain element, and land on the other side, and continue rolling. The combined rolling/flying behavior can extend the operational lifetime severalfold. Furthermore, the structure used for rolling the robot is designed to provide impact resilience while flying which provides further robustness. Our hybrid platforms consist of an adapted version of the above-mentioned NeBula payload along with the pixhawk firmware ([Meier et al., 2015](#)). [Fan et al. \(2019\)](#) describes the details of integration of the NeBula autonomy stack on the rollocopter platform.

Design evolution: [Figure 50](#) shows different variations of the rollocopter platform. The Hytaq version ([Kalantari and Spenko, 2014](#)) consists of a rotating cage, which encloses the drone to enable rolling. [Figure 49](#) shows a timelapse for early hybrid mobility tests on reinforced versions of this platform at JPL’s Mars Yard ([CoSTAR Team, 2018a](#)). Its modular, multi-agent version, called Shapeshifter ([Agha-mohammadi et al., 2020](#); [Tagliabue et al., 2020](#)), self-assembles this shell using permanent electromagnets. The passive-wheeled-rollocopter (PW-rollocopter) ([Fan et al., 2019](#); [Lew et al., 2019](#)) uses two independent passive wheels to enable large sensory FOV and autonomy. To reduce propeller-generated dust, DrivoCopter ([Kalantari et al., 2020](#)) and BAXTER ([Choi et al.,](#)



Figure 49. Hybrid platform flying over an obstacle at the JPL Mars Yard (CoSTAR Team, 2018a).



Figure 50. Hybrid platform hardware evolution (top row, left to right) pre-challenge Hytaq design (Kalantari and Spenko, 2014), modular hybrid vehicle (Agha-mohammadi et al., 2020), (bottom row, left to right) autonomous rollocopter (Fan et al., 2019; Lew et al., 2019), DrivoCopter (Kalantari et al., 2020), BAXTER (Choi et al., 2020), spherical rollocopter (Sabet et al., 2020, 2019).

2020) implement dedicated wheeled actuators, at the expense of a slight reduction in flight time. The latest work is focused on omnidirectional spherical rollocopters that provide maximum agility and maneuverability while having the minimum friction (Sabet et al., 2019, 2020).

13. Experiments

In this section, we present Team CoSTAR’s multi-year effort to validate NeBula technologies on physical and virtual systems. We first briefly discuss our simulation-based validation approach, and then discuss the performance of NeBula on heterogeneous physical robot teams in various challenging real-world environments. The videos in CoSTAR Team (2020a,b,c,d, 2021) depict some highlights of these runs.

13.1. Simulation Results

Simulator configuration: We use simulations both for component development and integration testing. NeBula relies on multiple simulator configurations at different fidelity levels to enable faster and more focused development. Examples include a high-fidelity Gazebo simulator for testing ground vehicle behaviors, flight stack software-in-the-loop simulations for analyzing flight performance, a docker-based multi-robot networking simulator, and a low-fidelity dynamics simulator for Monte Carlo analysis. Our simulator setup is portable, and used in local and cloud environments depending on the computational resources it requires and the number of agents that are deployed during the test.




Figure 51. Robot and environment simulation models.

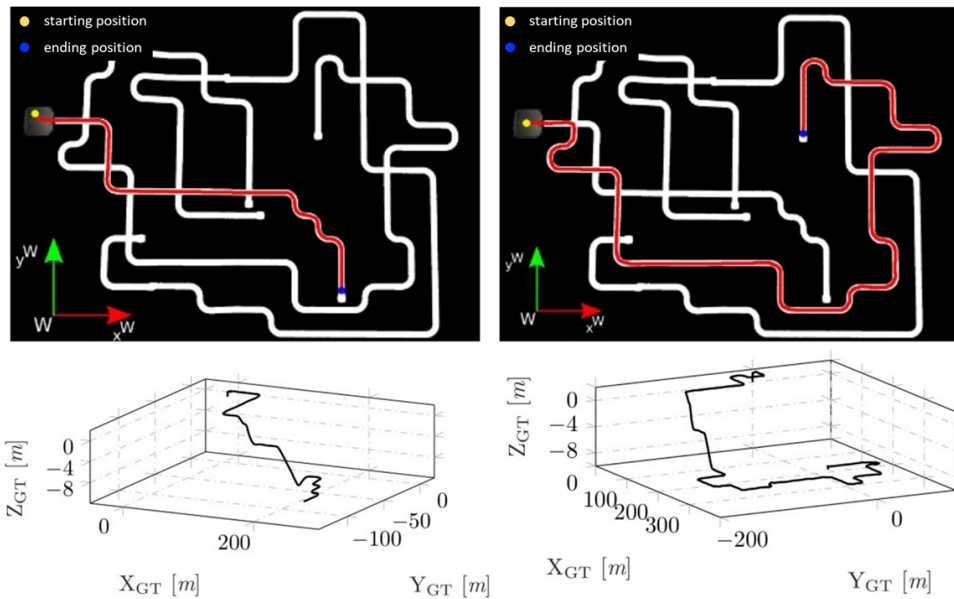


Figure 52. Drone navigation performance comparison in low-light tunnel environment with perception-aware nonlinear MPC, and heuristics-based junction sequence following.

Robot and environment models: Figure 51 shows selected examples of our simulated robot and environment models. In addition to many existing open-source models including the ones available at SubT Tech Repo (DARPA, 2018b), we built our custom robot and environment models, some of which are submitted to and available on the same repository. For example, Figure 52 shows a comparative study of drone navigation algorithms using the robot and environment models that are publicly available on the repository.

Self-organized simulation events: We regularly organize simulation-based events (“virtual” demos emulating the virtual DARPA SubT Challenge; see Chase and Kitchen, 2020) to track performance statistics over time. In each virtual demo, we evaluate the latest system performance based on several evaluation metrics. We developed a set of automated analysis tools to evaluate statistics on exploration, localization, mapping, artifact scoring, and human intervention. Simulation enables us to measure the statistics from large-scale tests which cannot be obtained easily with

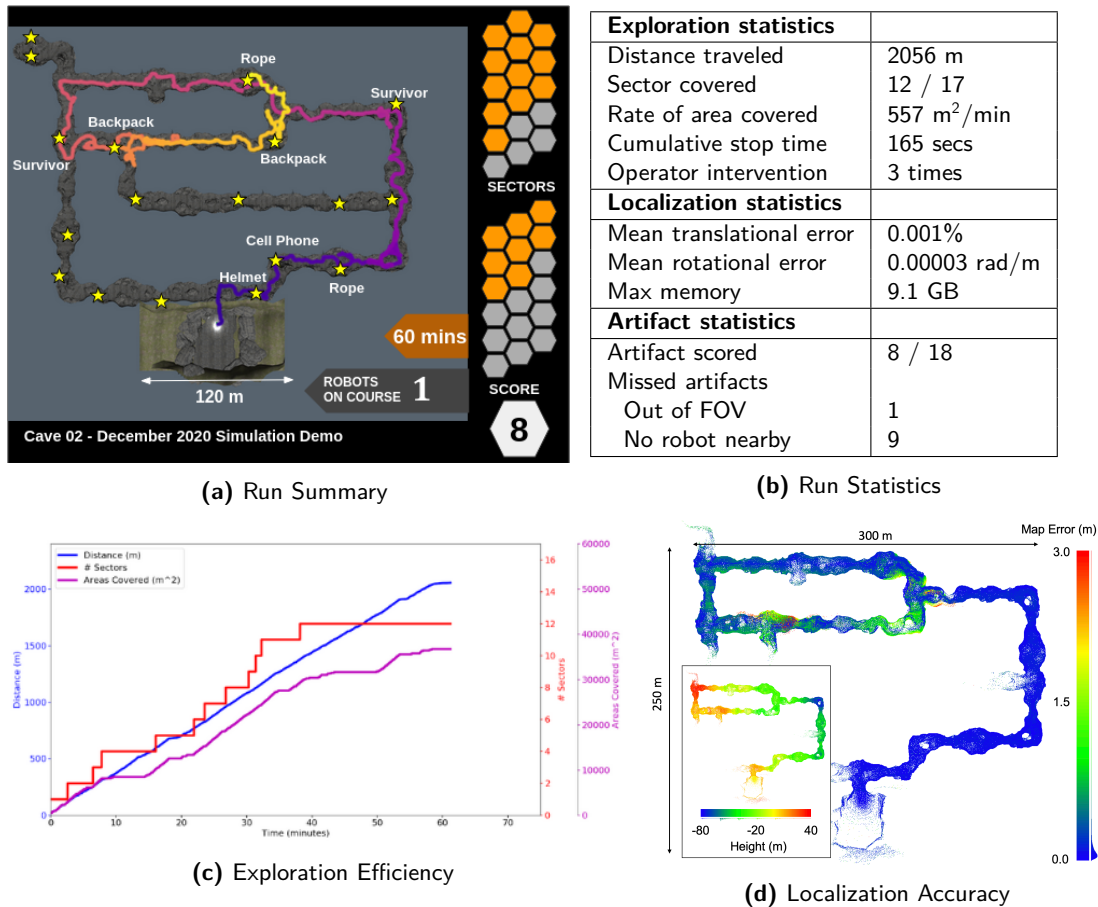


Figure 53. Sample result from Team CoSTAR's regular virtual demos.

hardware experiments. Figure 53 shows exploration and operation statistics of a single robot simulation in the Gazebo simulator. The competition statistics (Figure 53(a)) provide high-level evaluation with the number of sectors covered, number of artifacts scored, and robot trajectories in the course. The exploration statistics (Figure 53(b), (c)) evaluate how quickly and efficiently the autonomous exploration behavior covered the large-scale environment. The human supervisor intervention is monitored to measure the reliability of the autonomous system (Figure 53(b)). The localization and mapping performance (Figure 53(b), (d)) is evaluated against the ground truth dataset generated from the simulator. Finally, the artifact scoring performance is quantified with detection and localization evaluation (Figure 53(b)).

13.2. Field Tests and Demonstrations

Our system has been rigorously field tested in over 100 field tests in 17 different off-site locations, from lava tubes to mines 240 m underground as well as numerous on-site locations including JPL's Mars Yard. This testing regime is a fundamental part of field hardening our system to perform robustly in real-world challenging environments. A snapshot of some of our field-test locations is shown in Figure 54. In this section, we discuss highlights of our system's performance at these self-organized field tests and demonstrations.

Mine and tunnel field tests—ground vehicles: Field tests in tunnel environments as well as various underground mines stressed a variety of system capabilities, from dusty silver mines with narrow passages, to muddy coal mines with multiple decision points and massive scale. Figure 55

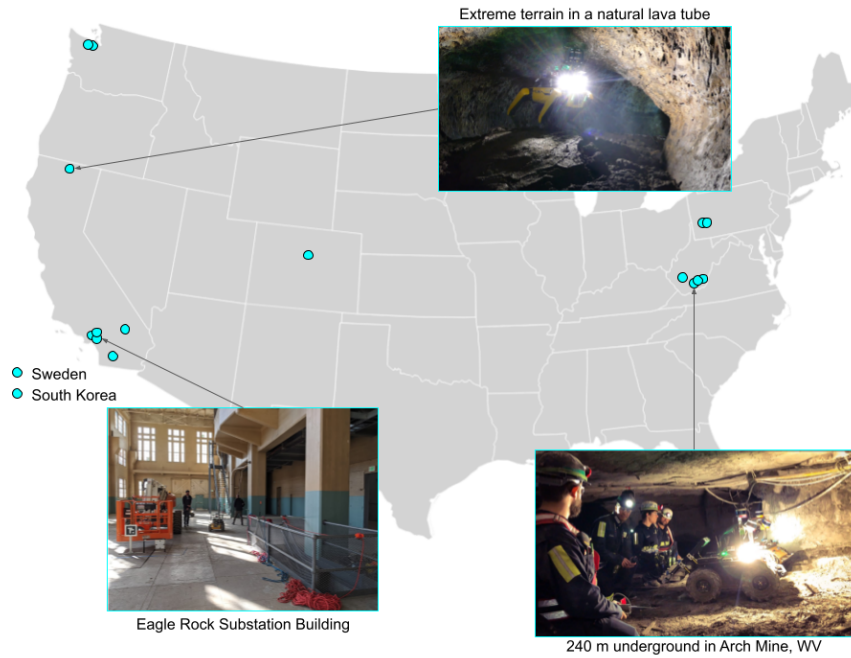


Figure 54. Locations and snapshots from field test activities. Each blue circle is a different field-test location.

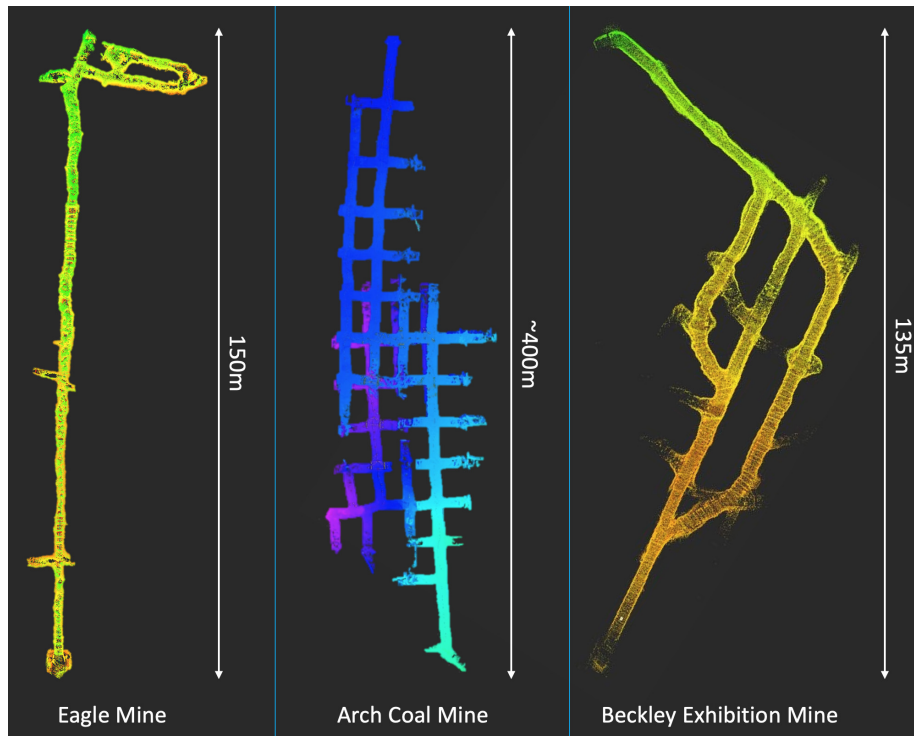


Figure 55. SLAM point-cloud maps from field tests in different mine locations. Left: A silver mine in southern California. Center: A portion of a large coal mine in West Virginia, 800 ft underground. Right: A near-surface historic coal mine in West Virginia.




Figure 56. Autonomous flights in Mjölkuddsberget Mine, Sweden (left), and Hwasoon Mine, Korea (right).

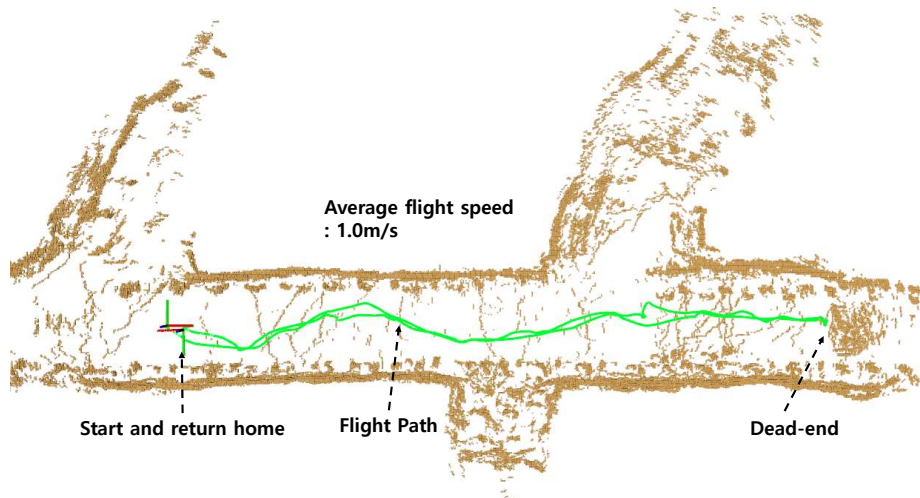


Figure 57. Autonomous flight test at Beckley Mine, Morgantown, WV. In this example, the drone flies approximately 50 m over 45 s. While flying, the drone meets the dead end and returns to the communication range to deliver the map and artifact data before explore other branches of the mine.

presents three maps from these tests. The Arch Mine tests (Figure 55, center) were in a portion of an active coal mine, and included extreme traversability hazards (see Figure 22), complex topology, and severely degraded communications. The multi-robot map is the result of a 30-minute operation, with a total of $400 \times 100 \text{ m}^2$ of area covered by the two-robot team. Eagle Mine and Beckley Exhibition Mine are maps from a single robot.

Tunnel field tests—airial vehicles: Aerial vehicles are deployed to perform fast navigation and exploration, while providing a rough topological map and collecting data from areas inaccessible to ground robots. Figure 56 shows two of the several hundred successful flight experiments conducted in tunnels located in the United States, South Korea, and Sweden. One specific example of drone flight is shown in Figure 57 (CoSTAR Team, 2018b). In this example the drone shows successful flight autonomy capability with local planning method in a confined and cluttered environment. The reliability of the developed system is best exemplified with the well over a thousand logged minutes of safe autonomous flight in complex and perceptually challenging tunnel environments.

Urban field tests—ground vehicles: Urban field testing provides more abundance of possible locations, but also a large variety of challenging conditions. We tested our system, in narrow cubicle farms of multi-level office buildings, with regions of self-similar corridors (Figure 58, top left); in multi-level parking structures with wide open spaces (Figure 58, bottom left), and in large, open industrial buildings with high ceilings, narrow stairways, and narrow doorways (Figure 58, right).

Urban field tests—hybrid vehicles: Hybrid ground/aerial mobility shows its benefit in the urban environment, where there are plenty of easy-to-roll flat surfaces as well as vertical openings






Figure 58. SLAM point-cloud maps from field tests in different locations. Topleft: An office building at JPL. Bottom left: A multi-level parking structure at JPL. Right: Eagle Rock substation and surrounds (pictured at top right).



(a) Time sequence of planned hybrid transition from rolling to flying and back to rolling. Colored blocks are occupied voxels colored by z height. Blue/Red path indicates hybrid planned path where blue is rolling and red is flying. Pink sphere is goal waypoint. Motion primitives from local planner are shown. Note that as the vehicle moves forward, an obstacle is revealed and a small hop over it is planned and executed.



(b) Power-energy consumption comparison with hybrid (green), rolling-only (blue), and flying-only (red) mobility modes. The hybrid mode flies over three obstacles, while rolling-only mode was tested without any obstacles.

Figure 59. Rollocopter mobility mode comparison.

to navigate in multi-level environments. Autonomous hybrid mobility testing results (Fan et al., 2019) show the benefit of hybrid mobility compared to pure ground or flying vehicles. The hybrid mode allows the robot to fly over obstacles that block the way for ground robots (Figure 59(a)). In comparison to the pure aerial vehicles, the hybrid mode shows better energy consumption profile by rolling on easy flat terrain (Figure 59(b)).

13.3. DARPA Organized Events

This section outlines results from testing NeBula in the series of DARPA facilitated test events as part of the Subterranean Challenge (Orekhov and Chung, 2021). The timeline of these events is summarized in Figure 60, with substantial developments over each of the 6-month intervals between




Figure 60. Overview of the timeline of DARPA organized test activities, with the expectation of increasing capability from one event to the next.




Figure 61. Four heterogeneous mobility platforms deployed at the STIX event. Clockwise from top left: Rollocopter (with Husky in the course), Husky, Flipper, and Scout.

tests by all teams, including (but not limited to) [Tranzatto et al. \(2021\)](#); [Issacs et al. \(2021\)](#); [Hudson et al. \(2021\)](#); [Rouček et al. \(2021\)](#); [Scherer et al. \(2021\)](#); [Ohradzansky et al. \(2021\)](#); [Lu et al. \(2021\)](#).

13.3.1. SubT Integration Exercise (STIX)

Environment and robots: STIX was held in April 2019 at the Edgar Experimental Mine, Idaho Springs, CO. The participating teams were offered two practice sessions and one simulated scored run in an environment representative of the Tunnel Circuit. For Team CoSTAR, one of the primary objectives was to quantify the capability and limitation of various mobility systems through the evaluation in a real mine environment. To that end, we deployed four different types of mobility platforms shown in [Figure 61](#), including a Clearpath Husky wheeled vehicle, Flipper tracked vehicle, Scout quadcopter, and Rollocopter aerial/ground hybrid vehicle.

Performance: Each platform demonstrated autonomous exploration capabilities in the mine. Husky autonomously traversed 200 m until it reached the first fork. During the drive, Husky detected five artifacts and reported back to the base station: two of them were correct in type and locations. Soon after Husky turned into a side passage, the communication to the base station was lost, which triggered the return to base behavior as designed. The robot successfully returned to the

base by following the breadcrumb nodes in IRM. Rollocopter performed four 3-minute runs during the practice sessions, each of which consisted of autonomous take-off, exploration, and landing operations. The robot exhibited robust navigation capabilities, traveling over 100 m in cumulative distance. Dust was a major issue for drone flight in the narrow passageways, causing vision-based state estimation failures. Our approach of using heterogeneous odometry sources enabled us to be resilient to these failures to a large degree (Santamaria-Navarro et al., 2019). Dust largely obscured all cameras after a few minutes into each run. With no measurement updates from odometry sources, we relied on IMU-only inertial odometry to safely land the vehicle. These results led to future work on improving state estimation resilience in the presence of dust and variable lighting. It also helped us to improve the camera placement design for the various scales of environments. The Flipper was able to navigate over the train tracks in the tunnel because of its tracked design but it was much slower than the Husky. This is because of greater contact area of the tracks that causes slower turning behavior compared to differential drive (wheeled) robots for the same commanded track/wheel speeds (Dixit and Burdick, 2020). This is why we decided to incorporate a hybrid vehicle that had both wheels and tracks (Telex) in the next circuit. We extensively tested Telex in the Arch mine (Thakker et al., 2020) and were able to navigate over different types of terrain (in tracked mode) while maintaining speed when the robot was on flat terrain (in wheel mode).

13.3.2. Tunnel Circuit

Environment: The Tunnel Circuit took place in August 2019 at the NIOSH mine complexes in Pittsburgh, PA. There were two courses, Experimental (EX) and Safety Research (SR), focusing on different aspects of the tunnel environmental challenges. The EX course is composed of long straight corridors with featureless flat walls. The SR course has a grid-like structure that provides many decision points and loop closures.

Robots: Team CoSTAR staged seven robots at the starting gate (see Figure 62). The robot roster includes four Clearpath Husky robots (Husky 1–4), one Telerob Telex track/wheel hybrid robot (Telex 1), one high-speed RC car (Xmaxx 1), and one aerial drone (Scout 1). Our strategy was to adaptively deploy the heterogeneous robot team based on the complexity and challenges of the unseen course elements. The information from vanguard robots told that the environment is




Figure 62. Robot team deployed at the Tunnel Circuit event. Top left: Four Huskies at the staging area. Top right: Drone deployed at the Experimental Mine. Bottom left: Husky deploying a communication node. Bottom right: RC car deployed at the Experimental Mine.




Figure 63. Driving distance statistics during the Tunnel and Urban Circuits. Prefixes refer to different robot types: D, Drone, H, Husky (wheeled robot), S, Spot (legged robot), X, XMax (1/5 scale RC car).




Figure 64. SLAM point-cloud map from a single Husky robot run in the Safety Research run of the Tunnel Circuit. The robot starts from the bottom left, and completes the top loop in a clockwise direction.

benign and mostly accessible by the ground vehicles, leading the decision to rely more on wheeled platforms that show higher endurance.

Performance: In each run, we sent three or four robots to the course and achieved more than 2 km combined traversal (Figure 63). The video in CoSTAR Team (2020c) depicts some highlights of these runs. The longest single-robot drive was 1.26 km by Husky 4 at the SR mine on Day 3, including long periods of no-communication autonomous exploration and successful returning to the communication range at the end of the mission. Figure 64 shows the map created by Husky 4 in this run, with an error under 1% of the distance traveled. Four communication nodes are deployed from the robots to build a backbone network, covering the areas near the mine entrance and extending the reach of base station for faster data retrieval. We detected 16 artifacts during our four runs of the Tunnel Circuit, leading to a circuit score of 11 and the second-place finish among 11 teams (Table 7).

Challenges: During the Tunnel Circuit event, we faced many real-world challenges which contributed to improve our system toward the next Circuit events. The flat featureless walls of

Table 7. Number of Scored Points for Each Competition Run

Circuit	Tunnel				Urban			
	SR R1	SR R2	EX R1	EX R2	A R1	A R2	B R1	B R2
Run								
Run scores	4	7	2	4	5	7	4	9
Course scores	7		4		7		9	
Circuit scores	11				16			
Circuit rank	2nd				1st			



Figure 65. Robot team deployed at the Urban Circuit event. Left: Robot fleet and communication node deployment. Right: A snapshot of Spot stair climbing (DARPA TV).

the EX course affected our localization performance which was purely LiDAR based as of the Tunnel Circuit. This motivated the development of multi-modal methods (Section 5 and Section 6). Multi-robot operations in a comm-degraded environment also posed challenges to our networking system. Based on the analysis of the Tunnel experience, we carefully redesigned the inter-robot networking protocol and deployed it in the Urban Circuit (Section 10). The drone traveled 35 m in 43 seconds before it experienced critical state estimation error due to poor lighting. To eliminate this single-point failure in the future flights, we put more efforts on multi-modal sensing and parallel estimation (Section 5).

13.3.3. Urban Circuit

Environment: The Urban Circuit took place in February 2020 at the Satsop Business Park in Elma, WA. The unfinished power plant in the park was chosen for the place of the second circuit event where the robots were exposed to challenges from urban structures. The two courses, Alpha and Beta, cover two floors of the power plant with size around $90 \times 90 \times 15 \text{ m}^3$. There are many small rooms and narrow corridors divided by thick walls that prevent direct wireless communications.

Robots: Team CoSTAR staged seven robots and deployed four of them during this event (Figure 65), including the newly introduced Boston Dynamics Spot quadruped robots (Spot 1 and 2). The other two robots are Clearpath Husky (Husky 1 and 4) with major upgrades in onboard electronics and sensor stack. All robots are deployed in each run, acting in the vanguard and supporting roles based on the assignment, their location, and time of the mission. One Spot was dedicated to climbing stairs to explore the floor at a different level.

Performance: We detected 25 artifacts during our four runs of the Urban Circuit, leading to a circuit score of 16 and the first-place finish among 10 teams (Table 7). Figure 66 shows a multi-robot map generated at the Beta course, including details on scored artifacts at multiple levels and at the furthest extent of exploration. The four-robot team traversed a combined total of 2.3 km (Figure 63), including a large closed loop by the two Husky robots around the central round structure. The video in CoSTAR Team (2020a) depicts Spot robots exploring the multi-level courses autonomously.

Figure 67 shows the breakdown of scoring in the two highest scoring runs of the Urban competition, summarizing the reasons for not scoring the remaining artifacts. In addition to the

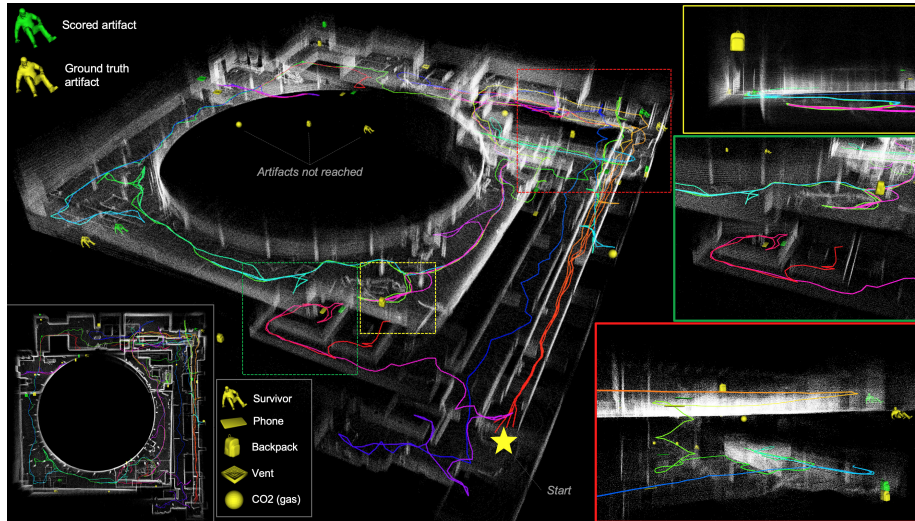


Figure 66. Multi-robot map, trajectories, scored artifacts, and ground truth artifacts. All robots start from the star on the bottom right with trajectory colors cycling from red through yellow to blue and back to red based on distance traveled. Red inset shows multi-level exploration and scoring (upper-level survivor and lower-level backpack) by Spot 1 with a stair descent (green part of trajectory). Green inset shows the furthest extent of lower-level exploration where a cell phone was scored by Spot 1. Yellow inset shows an instance of an unscored artifact, where it is placed out of view of ground robots, requiring flight. Other unscored artifacts are seen in the center of the map where no robot arrived near them.




Figure 67. Scoring breakdown for the two highest-scoring Urban Circuit runs.

artifacts missed due to the robot not reaching the artifact locations (Figure 66), we missed artifacts due to the limited field of view of the sensory suite on the robots (Figure 66) and other challenges in the artifact detection pipeline.

Challenges: Overall, Team CoSTAR’s system showed significant improvement from the Tunnel Circuit. Nonetheless, the circuit identified numerous ongoing challenges calling for further development. These development areas include improved depth and breadth of coverage, enhanced sensor fields of view for artifact detection, aerial robot developments, autonomous fault recovery, and careful attention to comm node placement/configuration to maximize wireless link quality for fast data retrieval.

13.4. Self-Organized Cave Circuit

Environment overview: We self-organized a Cave Circuit event at the Valentine Cave of the Lava Beds National Monument, Tulelake, CA. The cave, a class of lava tubes, was formed from volcanic flows, and an active research target for the future lunar/Martian cave exploration (Blank et al., 2018,




Figure 68. A point-cloud map from the SLAM module and traversability maps during autonomous exploration in the Valentine Cave at Lava Beds National Monument, Tulelake, CA.

2020). The length of the cave is approximately 300 m with an elevation change of 20 m. The tortuous and steep terrain in the cave limited the range of operations of wheeled platforms; hence tests were focused on the Spot quadrupeds (drone testing was not possible due to regulatory challenges).

Circuit event organization: The test was run as close to the competition setup as possible. The global frame and artifact locations were determined based on a pre-existing, high-precision Faro scan of the cave. A fiducial gate was surveyed into the cave frame and used to calibrate the robots. All artifact classes other than CO₂ and Vent were placed throughout the cave, with a prioritization on ropes and helmets (the cave-specific artifacts).

Results: The cave demonstration was remarkable in that entire runs were fully autonomous with zero human intervention: the operator only started the mission. Figure 68 shows a map from one such run. Spot traversed 400 m on average over four runs. Each mission was ended not by a time limit, but by the full accessible environment being covered, with side passages limited by traversability hazards (very low ceilings, cliffs). The videos of the cave demonstration are available at [CoSTAR Team \(2020b,d, 2021\)](#).

14. Lessons Learned

At the time of submission, NeBula's uncertainty-aware and perception-aware principles are implemented in a number of modules throughout our autonomy solution. The module interactions have been tightly co-designed and joint probability distributions across some of the modules are incorporated to bolster the overall system's resiliency. This uncertainty-aware architecture is verified through intensive field testing campaigns using heterogeneous platforms with different variations of mobility, sensing, and computing capabilities. We observed that the joint perception-planning approach brought resiliency to the system behaviors in challenging real-world environments, where uncertainty is ubiquitous. This section summarizes the lessons learned from developing and fielding the NeBula autonomy solution.

14.1. Heterogeneous System Design and Integration

Designing a system with heterogeneous robots is a nontrivial task. Accomplishing the objectives of complex real-world operations (such as the one in the SubT) with constraints on time, resources/cost, robot size, weight, power, etc., can be too difficult or impossible for a single robot. This calls for multi-robot solutions. An important lesson learned and a major future direction is, “It is critical to have principled tools that can translate the design choices to overall mission metrics.” These design choices range from the selection of type and number of robots to the selection of sensors, instruments, and algorithms, to the type of wires and connectors, and batteries. As some concrete examples that our team has encountered: “A slight change in the weight (and hence operation time) of one single robot” or “selection of a specific wire that can slightly affect the data traffic” can have a significant effect on the overall aggregated performance of the robot team and mission metrics. There are hundreds if not thousands of such design parameters and choices (both in HW and SW), which calls for systematic tools to quickly abstract them and translate them to high-level mission metrics.

An important observation is that for the fast-paced development of a very large autonomy architecture for a multi-robot system, the architecture should be highly modular and adaptive. Since the architecture must incorporate robots’ heterogeneous capabilities (mobility, sensing, and compute), appropriate abstractions need to be implemented to operate robots in a unified framework. We have performed many iterations of the architecture redesign as our entire system capabilities grow. A key lesson learned is that integration testing at a regular interval is essential to maintain the stability of such a growing system. Team CoSTAR has organized monthly virtual and physical integration demonstrations with competition-like set-ups. Regular testing enforces the new module to comply with the system interaction rules, ensuring functionality despite constant augmentations by multiple subteams throughout the development. Regular performance tracking also allows us to understand the current technology state and highlights the next domain to be studied.

14.2. Resilient State Estimation

In challenging real-world robotic operations, perceptual degradation is common, and state estimation can suffer from deterioration of sensor measurements and estimation quality. In addition, physical systems are exposed to the risk of sensor failures that are made prevalent by frequent field testing. These failures are difficult to model in many cases, yet systems must be able to react to these uncertain events to ensure system stability is not lost. One key lesson from our experience is to let the system predict a failure and quickly adapt to it. We observed that adding redundancy (e.g., HeRO at Section 5) in a principled way shows a lot of potential to improve the resiliency of an odometry system by incorporating multi-modal sensors and algorithms. HeRO’s built-in health check mechanism allows for preliminary detection of various types of failures and adapting the system to mitigate the effect. This proactive approach has been effective in our field tests and circuit events, providing the foundation to support high-level autonomous behaviors.

Another critical observation is on “how and when to trust the sensor models.” Following NeBula’s philosophy of uncertainty-aware predictive planning, the planner needs to be able to predict the joint distributions over the system’s future actions and estimated states. To predict the future performance, we divide unknowns into “known unknowns” and “unknown unknowns.” Known unknowns refer to uncertainties that can be reasonably modeled using probability distribution functions. Unknown unknowns refer to unmodeled uncertainties resulting from unexpected events in the system operation. To be resilient and robust to these uncertainties, we have observed that a cascaded framework (like HeRO) of loosely coupled and tightly coupled layers is promising. The cascaded framework first copes with unknown unknowns, by relying on a loosely coupled layer and checking the consensus across different sensory modalities to detect anomalies and reject faulty estimation channels. Once left with inliers, it handles “known unknowns” by tightly fusing the estimation channels using their predicted probabilistic models.

14.3. Large-Scale Positioning and Mapping

The SLAM design and deployment efforts presented in this paper highlight the importance of using complementary sensor information for localization and mapping. In complex environments, various sensors have their own advantages and disadvantages. For example, wheel-encoder-based odometry or visual odometry provides a useful source of information in long corridors where LiDAR scan matching is subject to drift; IMU data allow for accurate short-term relative rotation estimation while history-based scan matching can resolve longer-term displacement estimation. This need to fuse heterogeneous sensors also emphasizes the importance of a flexible SLAM framework where one can easily define and fuse sensor data when available. In this sense, the use of a factor graph framework reduced the integration efforts and allowed us to design a unified SLAM back end that can be easily reconfigured. The importance of developing a reconfigurable framework is further exacerbated by the desire of running our SLAM solution on a heterogeneous set of platforms (e.g., robots with different sensor suites and different computational capabilities), which demanded our framework to be sensor agnostic (e.g., adjust to different sensors with minimal parameter tuning) and reconfigurable (e.g., enable and disable sensors to fuse via configuration files).

While having a reliable odometry solution is critical in large-scale mapping, even the most accurate odometry systems accumulate error over long distances across extreme terrains. This remarks the importance of loop closures to keep the localization errors bounded. However, when it comes to complex, large-scale environments, with perceptually degraded conditions, a key lesson learned is that loop closure needs to be incorporated in a resilient manner into the overall framework. Computational constraints will limit the ability to search the map history and find correct loop closures in a reasonable time. Even after loop closures are found, it is crucial to maintain multiple hypotheses or at least remove incorrect loop closures resulting from perceptual aliasing. In this paper, we observed that *graph-based outlier rejection* (e.g., a variation of [Mangelson et al. \(2018\)](#)) is helpful in filtering some of the incorrect hypotheses. This family of methods relies on approximate max clique solvers for consistency maximization, which becomes computationally expensive due to the size and density of the SLAM graphs in competition settings. Unfortunately, they often fail to select inlier loop closures, inducing “jumps” in the robot trajectory when the solver is stuck in suboptimal maximal cliques. Future work includes improving loop closure detection (e.g., prioritizing the most informative loop closures, using different sensors for place recognition) and adopting more recent methods for outlier rejection based on graduated nonconvexity.

Computational aspects can be also improved by sharing the workload across robots. Centralized multi-robot SLAM methods require increasing computational resources for larger teams of robots. Adopting a distributed SLAM system can improve scalability and reduce communication bandwidth. In general, distributing the overall computation across various robots (based on their capabilities) can better scale to large teams of robots.

An important lesson learned is that augmenting geometric information with semantics can increase the resiliency of the SLAM system. For instance, incorporating semantic information such as intersections, stairs, and man-made objects in the graph can increase the robustness of the loop closures. Further, identifying a set of stairs in the environment provides a readily usable prior on the geometry of the stairs; similarly, identifying doors provides actionable information for navigation. Metric-semantic mapping is an active research area and a tight integration of geometric, semantic (and physical) reasoning has the potential to improve the robustness and accuracy of the map built by the robots.

Finally, while this section is concerned with localizing the robots and building a map, the quality of the map reconstruction is highly affected by the trajectories covered by the robots. Following NeBula's uncertainty-aware planning perspective, performing active loop closures to reduce the uncertainty in the robot location has a major impact on the SLAM accuracy. Active loop closures can be enabled by guiding robots to rendezvous points to create inter-robot measurements or by visiting parts of the environment that have unambiguous detection signatures. Towards this goal, quantifying and actively reducing uncertainty is crucial and is a fundamental trait of the NeBula

framework. Active localization and mapping go beyond loop closures, and tightly couple perception, action, and communication.

Uncertainty-aware traversability analysis: Navigating over perceptually degraded challenging terrains requires a radical departure from systems designed to operate in known environments with clear landmarks and easily distinguishable obstacles. Designed to detect complex geometric hazards at various scales, traversability analysis highly depends on the quality of local mapping and state estimation. As such, an important lesson has been that accounting for the “perception uncertainty” in planning is a key to building a reliable traversability system operating on challenging terrains and perceptual conditions. We learned that multi-fidelity mapping approaches (similar to the one presented in Section 8) improve the balance between computational constraints and accuracy in the presence of degraded state estimation and sensor measurement uncertainty. Uncertainty-aware estimation trusts and accumulates measurements in the world belief based on measurements’ accuracy and quality. Using this world belief, the traversability layer precisely quantifies the perception-aware traversability risks and costs when negotiating challenging terrains. Finally, a critical lesson in traversability algorithm design and development is constant field testing: The presented uncertainty-aware approach has been field hardened in over 100 field test sites with diverse traversability-stressing elements.

14.4. Scalable Belief-Space Global Planning

Global planning for area coverage and exploration behaviors is one of the modules where the awareness of uncertainty plays a critical role in achieving high performance in the real world. We remark several key dilemmas encountered while developing and testing the global planner, related to uncertainty representation and uncertainty-aware decision making: (i) scalability vs. information fidelity in world belief representation, (ii) planning horizon vs. planning time for each receding-horizon planning episode, and (iii) the consistency of plans for smooth motions vs. resiliency to sudden changes in world belief over time. We learned that hierarchical approaches (similar to the ones described in Section 9) have the potential to address the above challenges. Our method, PLGRIM, utilizes (i) hierarchical IRMs, (ii) longer-horizon, higher-resolution POMDP solvers with manageable computation load, and (iii) receding horizon planning with resiliency logic. However, those are still open problems. An important future direction is “quick online adaptation” of parameters that balance the environment scales and complexity with the computing system and sensor limitations. Another important open problem is related to encoding and capturing a more accurate and reliable representation of the high-dimensional world belief.

14.5. Semantic Understanding and Artifact Detection

Detecting objects and understanding the semantics of the environment in perceptually degraded subterranean settings is a challenging task. In the DARPA Subterranean Challenge, artifacts have multiple signatures, ranging from visual, thermal, auditory to gas-based, and radio-based artifacts. These various signatures and payload constraints (size, weight, power, etc.) on different heterogeneous platforms call for a method that assesses the value of adding various sensors and sensing modalities to each vehicle. Further, the choice and configuration of each sensor on each robot have a significant impact on the scoring performance (see Section 7). For example, for artifacts that can be detected visually, the salient camera parameters include field of view, reliability, image quality, resolution, and frame rate. Hence an important lesson has been that the artifact collection system needs a tight co-design of software and hardware. This includes the whole pipeline from the sensors, to the cables, to the processors and algorithms. For example, a change in a USB hub or cable can significantly impact the choices of other hardware and algorithms by removing certain data choking elements in the pipeline. A related unexpected failure mode (which occurred in one of our competitions) at the system level was sensor data transfer and communication, caused by both USB driver and networking bandwidth limitations. Correcting this problem required a thorough analysis

of the hardware configuration, including the development and testing of custom driver software by our team.

The second major lesson was related to data. While the existing standard object detection datasets are quite large and diverse, learning to detect objects in the perceptually degraded environments (e.g., variable lighting and with obscurants), such as the ones in the DARPA Subterranean Challenge-like environments, are still out of distribution in relation to mainstream datasets.

Finally after these trade-offs, there are still many limitations in the perception side of the artifact detection architecture. So the third, and the most important, lesson learned is that the planning and perception for artifact detection need to be tightly co-designed. In other words, active perception is required where the planner needs to take actions that compensate for perception shortcomings. Examples of such actions are (i) executing local search maneuvers where the robot sweeps the larger parts of the scene with its sensors to compensate for its limited field of view, and (ii) actively getting closer to certain targets or making measurements from various angles to increase the detection confidence. This mode might also include deciding which robot and which sensor should gather more information to increase the confidence. It also includes (iii) changing perception pipeline parameters such as the camera resolution, input rate, etc., to handle computational constraints by focusing and limiting the attention of the system on the important parts of the input channel. Active perception for semantic understanding is a highly open area, with a lot of future work in these domains.

14.6. Bandwidth-Aware Communication System Design

Communications between computational units (e.g., ROS nodes) should ideally take into account the predicted available bandwidth in the link whether that link is on the same computer, via a high-speed Ethernet, a larger wired network, or radio/wireless. We took a number of successful steps in this direction: (i) separating inter-agent (ROS 2) and intra-agent (ROS 1) communications, (ii) using different QoS settings for different classes of topics, (iii) monitoring inter-robot communications and estimated bandwidth, and (iv) monitoring intra-robot communications. As future multi-agent projects move away from ROS 1 and towards ROS 2 (or another middleware/communication solution), the same principles apply. We see that there is a need for continued improvements and autonomy in routing and QoS systems to optimally use communication resources in these future systems.

Another important lesson learned is that, due to the communication-degraded nature of the subterranean environments, the planner needs to support and improve the communication and multi-robot networking performance by dropping communication nodes at the optimal locations and by actively carrying (muling) data using mobile robots between the various nodes. Hence, aligned with NeBula's philosophy, (1) taking networking uncertainties into account, (2) evaluating and predicting the potential value of the communicated information, and (3) co-designing the planning system and networking system are the critical observations to increase the robustness and performance of the overall multi-robot robot networking system in communication-degraded environments.

14.7. Supervised Autonomy to Full Autonomy

Achieving mission success depends heavily on the mission-level autonomy from both single- and multi-robot perspectives. This is especially the case in contexts such as the SubT Challenge where large robot teams must intelligently coordinate themselves under strict communication, time, and resource constraints to explore and map kilometer-scale environments and find objects of interest.

Closer to the course entrance a communication link with the human supervisor can be established. Thus the system must be capable of autonomously balancing between deeper exploration, operational risk, the value of the collected onboard information, and the cost of bringing the systems closer to the surface for the data retrieval. A key lesson learned is that even when robots are close to the course entrance and a communication link with the human supervisor is established, with a large heterogeneous multi-robot team, a single human supervisor becomes a bottleneck in the control loop due to excessive cognitive workloads of the supervisor. Hence, for successful operations, NeBula's

mission executive and Copilot assistant are designed to support such tasks by adjusting the system autonomy levels, gradually delegating human tasks where the human is viewed as a resource, and providing suitable instructions and feedback to reduce the cognitive load of the single supervisor. Integral to the effort is maintaining a world belief that can be easily interpreted by the system or supervisor, and trigger behaviors or sequences of actions that bring the system closer towards achieving mission completion.

However, there is a lot of future research in this area. Our ongoing efforts have been focusing on some of these areas, including efficient ways of specifying more complex tasks and missions, balancing human-machine task distribution, planning using semantic information, and scheduling under uncertainty in task execution and future event occurrence.

14.8. Simulator-Based Development

The development of our system was greatly accelerated by the use of different computer simulation environments, discussed herein. The challenges inherent to real-world robotics in uncertain and fully autonomous settings pose tremendous risks to naive robots and algorithms in the early stages of development. Taking development into a simulation and out of hardware has afforded us numerous opportunities to pursue and achieve high-risk-high-gain algorithmic strategies. Many safety-critical features are difficult to test with physical hardware. By carefully modeling different fleet and hardware configurations, we were able to develop, test, and eventually deploy our heterogeneous multi-robot team in the real world, with optimal networking, hardware, and software payloads. Although simulators offer an enormous benefit to robotic systems development, there are significant differences between simulations and real-world systems which must be accounted for. It is important to understand simulator limitations and always perform proper validation and testing on physical hardware platforms.

15. Conclusion

Team CoSTAR's approach to the DARPA Subterranean Challenge lies in our autonomy solution, NeBula (Networked Belief-aware Perceptual Autonomy), which emphasizes resilience and intelligent decision making through uncertainty awareness. NeBula has led to second and first place in DARPA Subterranean Challenge's Tunnel (in 2019) and Urban (in 2020) competitions, respectively. When dealing with exploration and operation in unknown environments, uncertainty is inherent to all decisions. As the main principle, NeBula focuses on quantifying and taking advantage of uncertainty at multiple levels of our autonomy stack, including state estimation, mapping, traversability, planning, communications, and other state domains. We combine these technologies in a synergistic way, which examines the interaction between interrelated modules.

During this competition we have demonstrated and hope to continue to demonstrate autonomous exploration in extreme environments on multiple platforms (including wheeled, legged, and aerial). Solving this problem remains of paramount interest in a wide range of applications, especially when it comes to missions to explore unknown planetary bodies beyond our home planet. We believe the uncertainty-aware and platform-agnostic nature of most NeBula components is a step towards resilient and safe robotic autonomy solutions in unknown and extreme environments with both single- and multi-robot systems.

Acknowledgments

The work is partially supported by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004), and Defense Advanced Research Projects Agency (DARPA).¹

We also would like to thank our collaborators and sponsors including Boston Dynamics, JPL Office of Strategic Investments, Intel, Clearpath Robotics, Velodyne, Telerob USA, Markforged, Silvus Technologies, ARCS California State University Northridge, Vectornav, HM Engineering, ARCH Mine, Mine Safety and Health Administration (MSHA), National Institute for Occupational Safety and Health, Polytechnique Montreal, and West Virginia University.

We would like to thank the rest of Team CoSTAR, our collaborators, and advisors for their support, fruitful discussions, and their contributions to the project. In particular, we would like to thank Giulio Autelitano, Nicholas Ohi, Mamoru Sobue, Yuki Kubo, Kenny Chen, Jennifer Blank, Kevin Lu, Sahand Sabet, Sandro Berchier, Yasin Almalioglu, Micah Corah, Meriem Ben Miled, Joseph Bowkett, Navid Nasiri, Tomoki Emmei, Rianna M Jitosh, Gita Temelkova, Dave Gallagher, Leon Alkalai, Richard Volpe, Lorraine Fesq, Robert Stirbl, Randy Odle, Cagri Kilic, Leonardo Forero, Daniel Tikhomirov, Juan Nieto, Kyle Strickland, Chris Denniston, Antoni Rosinol, Alexander Nikitas Dimopoulos, Harrison Delecki, Nicolas Marchal, Oriana Claudia Peltzer, Joshua Ott, Nicholas Palomo, Adam Jaroh, Heiko Helble, Aaron Ames, Issa Nesnas, Fernando Mier-hicks, Shreyansh Daftry, Nicholas Peng, Mitra Azar, Corey Mack, Darmindra Arumugam, Jack Bush, Terry Suh, Mike Paton, Jared Beard, Jennifer Nguyen, Evangelos Theodorou, Michelle Tan, Andrew Haeffner, Pradyumna Vyshnav, Marcus Abate, Alexandra Stutt, Kartik Patath, Arghya Chatterjee, Shehryar Kattak, Chris Heckman, Kostas Alexis, Zack Jackowski, Michael Perry, Marco da Silva, Nhut Tan Ho, Barry Ridge, Brett Kennedy, Larry Matthies, Richard French, Jeffery Hall, Andila Wijekulasuriya, Paul Backes, Marco Tempest, Joe Bautista, Kayla Mesina, Erica Bettencourt, Amiel Gitai Hartman, Jeff Delaune, Daniel Pastor, Yu Gu, Jason Gross, Jacopo Villa, Soojean Han, David Chan, Rashied Amini, Jake Ketchum, Andrew Bieler, Slawek Kurdziel, Steve Zhao, Reza Radmanesh, Jose Mendez, Roberto Mendez, Christopher Patterson, Jack Dunkle, Jong Tai Jang, Petter Nilsson, Filip Claesson, Emil Fresk, Ransalu Senanayake, Mohammadjavad Khojasteh, Andrew Singletary, Alexandra Bodrova, and Tesla Wells.

The content is not endorsed by and does not necessarily reflect the position or policy of the government or our sponsors.

16. Glossary: Acronyms

APE Absolute Position Error

AUV Autonomous Underwater Vehicle

BPMN Business Process Modelling Notation

BRM Belief Roadmap

CHORD Collaborative High-bandwidth Operations with Radio Droppables

CIO Contact-Inertial Odometry

Comm Communication

ConOps Concept of Operations

CoSTAR Collaborative SubTterranean Autonomous Robots

CNN Convolutional Neural Network

CVaR Conditional Value-at-Risk

DARPA Defense Advanced Research Projects Agency

DLT Direct Linear Transform

EdgeTPU Edge Tensor Processing Unit

FGA Flat Ground Assumption

FIRM Feedback-based Information Roadmap

FOV Field of View

FPS Frames per Second

GCP Global Coverage Planning

GESTALT Grid-based Estimation of Surface Traversability Applied to Local Terrain

GICP Generalized Iterative Closest Point
















GO-LLC Geometric-Only LiDAR Loop Closures

GNSS Global Navigation Satellite System
GPS Global Positioning System
GPU Graphical Processing Unit
GTSAM Georgia Tech Smoothing and Mapping
HeRO Heterogeneous and Resilient Odometry Estimator
HFE Hierarchical Frontier-based Exploration
IMU Inertial Measurement Unit
iSAM2 incremental Smoothing and Mapping
IRM Information Roadmap
JPL Jet Propulsion Laboratory
LAMP Large-scale Autonomous Mapping and Positioning
LCP Local Coverage Planner
LIO LiDAR-Inertial Odometry
LION LiDAR-Inertial Observability-aware Navigation for vision-denied environments
LO Lidar Odometry
LOCUS Lidar Odometry for Consistent operations in Uncertain Settings
MAV Micro Aerial Vehicle
MDC Motion Distortion Correction
MER Mars Exploration Rover
MIKE Multi-robot Interaction assistant for unKnown cave Environments (aka “Copilot,” “Copilot MIKE”)
MPC Model Predictive Control
MSL Mars Science Laboratory
MSR Mars Sample Return
NBV Next-Best-View
NeBula Networked Belief-aware Perceptual Autonomy
NCDS NeBula Communications Deployment System
NDB NeBula Diagnostics Board
NPCC NeBula Power and Computing Core
NSP NeBula Sensor Package
PLGRIM Probabilistic Local and Global Reasoning on Information roadMaps
POMCP Partially Observable Monte Carlo Planning
POMDP Partially Observable Markov Decision Process
QoS Quality of Service
RA-LLC Range-Aided LiDAR Loop Closures
RC Remote Control
RF Radio Frequency
RHC Receding Horizon Control
RIO RaDAR-Inertial Odometry
RRT Rapidly-exploring Random Trees
QP Quadratic Programming
ROAMS Rover Analysis Modeling and Simulation
ROS Robot Operating System
RSSI Received signal strength indication
SLAM Simultaneous Localization and Mapping
SLAP Simultaneous Localization and Planning
SMAP Simultaneous Mapping and Planning
SOG Sum of Gaussians
STEP Stochastic Traversability Evaluation and Planning
STIX SubT Integration Exercise
TIO Thermal-Inertial Odometry
TRACE Traceable Robotic Activity Composer and Executive

UAV Unmanned Aerial Vehicle
UWB Ultrawideband
VIO Visual-Inertial Odometry
VO Visual Odometry
WIO Wheel-Inertial Odometry

ORCID

Ali Agha  <https://orcid.org/0000-0001-5509-1841>
 Kyohei Otsu  <https://orcid.org/0000-0002-9363-8817>
 Benjamin Morrell  <https://orcid.org/0000-0002-9768-3615>
 Rohan Thakker  <https://orcid.org/0000-0002-2672-2103>
 Angel Santamaria-Navarro  <https://orcid.org/0000-0002-6328-1423>
 Sung-Kyun Kim  <https://orcid.org/0000-0003-2460-1153>
 Amanda Bouman  <https://orcid.org/0000-0002-4215-2913>
 Xianmei Lei  <https://orcid.org/0000-0003-4971-6977>
 Jeffrey Edlund  <https://orcid.org/0000-0003-3092-4493>
 Muhammad Fadhil Ginting  <https://orcid.org/0000-0003-4030-5151>
 Kamak Ebadi  <https://orcid.org/0000-0002-4094-7077>
 Matthew Anderson  <https://orcid.org/0000-0001-8884-3448>
 Edward Terry  <https://orcid.org/0000-0002-7495-0506>
 Michael Wolf  <https://orcid.org/0000-0003-4534-0036>
 Andrea Tagliabue  <https://orcid.org/0000-0003-3376-4114>
 Tiago Stegun Vaquero  <https://orcid.org/0000-0002-4681-4363>
 Matteo Palieri  <https://orcid.org/0000-0001-8218-3346>
 Scott Tepsuporn  <https://orcid.org/0000-0003-2028-9377>
 Yun Chang  <https://orcid.org/0000-0002-2829-5256>
 Arash Kalantari  <https://orcid.org/0000-0003-2184-8548>
 Brett Lopez  <https://orcid.org/0000-0001-5366-911X>
 Nobuhiro Funabiki  <https://orcid.org/0000-0002-4713-9689>
 Thomas Touma  <https://orcid.org/0000-0002-4921-2210>
 Jesus Tordesillas  <https://orcid.org/0000-0001-6848-4070>
 Nikhilesh Alatur  <https://orcid.org/0000-0003-4783-788X>
 Jeremy Nash  <https://orcid.org/0000-0002-1854-8504>
 Sunggoo Jung  <https://orcid.org/0000-0001-7672-6885>
 Hanseob Lee  <https://orcid.org/0000-0003-4186-9418>
 Christoforos Kanellakis  <https://orcid.org/0000-0001-8870-6718>
 John Mayo  <https://orcid.org/0000-0002-7175-6599>
 Scott Harper  <https://orcid.org/0000-0002-1232-0621>
 Marcel Kaufmann  <https://orcid.org/0000-0001-5132-592X>
 Anushri Dixit  <https://orcid.org/0000-0002-9698-2189>
 Gustavo J. Correa  <https://orcid.org/0000-0003-4112-9495>
 Carlyn Lee  <https://orcid.org/0000-0001-6206-071X>
 Jay Gao  <https://orcid.org/0000-0002-0713-7865>
 Gene Merewether  <https://orcid.org/0000-0003-3046-9400>
 Gautam Salhotra  <https://orcid.org/0000-0003-2201-9303>
 Maira Saboia Da Silva  <https://orcid.org/0000-0002-6277-1703>
 Benjamin Ramtoula  <https://orcid.org/0000-0002-9185-7995>
 Taeyeon Kim  <https://orcid.org/0000-0003-3139-1799>
 Tara Bartlett  <https://orcid.org/0000-0002-4336-6219>
 Alex Stephens  <https://orcid.org/0000-0001-6209-0548>

Eric Heiden  <https://orcid.org/0000-0002-2031-8564>
 Thomas Lew  <https://orcid.org/0000-0002-8209-3604>
 Abhishek Cauligi  <https://orcid.org/0000-0003-1262-0859>
 Tristan Heywood  <https://orcid.org/0000-0002-9047-2292>
 Henry A. Leopold  <https://orcid.org/0000-0002-2685-5635>
 Hyungho Chris Choi  <https://orcid.org/0000-0002-4758-2007>
 Shreyansh Daftry  <https://orcid.org/0000-0002-0157-5944>
 Olivier Toupet  <https://orcid.org/0000-0003-4993-7390>
 Inhwan Wee  <https://orcid.org/0000-0001-5294-1150>
 Abhishek Thakur  <https://orcid.org/0000-0003-4149-1312>
 Micah Feras  <https://orcid.org/0000-0002-9324-7875>
 Giovanni Beltrame  <https://orcid.org/0000-0001-9755-8630>
 George Nikolakopoulos  <https://orcid.org/0000-0003-0126-1897>
 Luca Carlone  <https://orcid.org/0000-0003-1884-5397>
 Joel Burdick  <https://orcid.org/0000-0002-3091-540X>

References

- Agha-mohammadi, A., Chakravorty, S., and Amato, N. (2014). FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *International Journal of Robotics Research*, 33(2):268–304.
- Agha-mohammadi, A., Agarwal, S., Kim, S.-K., Chakravorty, S., and Amato, N. M. (2018). SLAP: Simultaneous localization and planning under uncertainty via dynamic replanning in belief space. *IEEE Transactions on Robotics*, 34(5):1195–1214.
- Agha-mohammadi, A., Heiden, E., Hausman, K., and Sukhatme, G. (2019). Confidence-rich grid mapping. *International Journal of Robotics Research*, 38(12–13):1352–1374.
- Agha-mohammadi, A., Tagliabue, A., Schneider, S., Morrell, B., Pavone, M., Hofgartner, J., Nesnas, I. A., Amini, R. B., Kalantari, A., Babuscia, A., et al. (2020). The shapeshifter: A morphing, multi-agent, multi-modal robotic platform for the exploration of Titan (preprint version). *arXiv preprint arXiv:2003.08293*.
- Agarwal, V. (2019). Ground up design of a multi-modal object detection system. Master’s thesis, Carnegie Mellon University, Pittsburgh, PA.
- Ahmad, S., Sunberg, Z. N., and Humbert, J. S. (2021). APF-PF: Probabilistic depth perception for 3D reactive obstacle avoidance. In *2021 American Control Conference (ACC)*, pages 32–39.
- Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., and Sivic, J. (2016). NetVLAD: CNN architecture for weakly supervised place recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5297–5307.
- Autopilot, P. (2021). PX4FLOW Smart Camera. <https://docs.px4.io/master/en/sensor/px4flow.html>.
- Bajracharya, M., Maimone, M. W., and Helmick, D. (2008). Autonomy for Mars rovers: Past, present, and future. *Computer*, 41(12):44–50.
- Balta, H., Bedkowski, J., Govindaraj, S., Majek, K., Musialik, P., Serrano, D., Alexis, K., Siegwart, R., and De Cubber, G. (2017). Integrated data management for science & a fleet of search-and-rescue robots. *Journal of Field Robotics*, 34(3):539–582.
- Benton, J., Coles, A., and Coles, A. (2012). Temporal planning with preferences and time-dependent continuous costs. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 22.
- Blank, J. G., Roush, T. L., Stoker, C. L., Colaprete, A., Datta, S., Wong, U., Deans, M., Fong, T., Léveillé, R., Northup, D. E., et al. (2018). Planetary caves as astrobiology targets. In *AGU Fall Meeting 2018*. AGU.
- Blank, J. G., Agha-mohammadi, A., Bell, Jr., E. R., Crown, D. A., Morrell, B., Patterson, C. J. S., Uckert, K., and Whelley, P. L. (2020). Volcanic caves as priority sites for astrobiology science. White paper submitted to the *Decadal Survey on Planetary Science & Astrobiology*.
- Bochkovskiy, A., Wang, C.-Y., and Liao, H.-Y. M. (2020). YOLOv4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*.
- Bonnabel, S., Barczyk, M., and Goulette, F. (2016). On the covariance of ICP-based scan-matching techniques. In *2016 American Control Conference (ACC)*, pages 5498–5503.

- Bouman, A., Ginting, M., Alatur, N., Palieri, M., Fan, D. D., Touma, T., Pailevanian, T., Kim, S., Otsu, K., Burdick, J., and Agha-Mohammadi, A. (2020). Autonomous Spot: Long-range autonomous exploration of extreme environments with legged locomotion. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Chase, R., and Kitchen, S. (2020). How we won the DARPA SubT Challenge: Urban circuit virtual track. <https://spectrum.ieee.org/darpa-subt-challenge-urban-circuit-virtual>.
- Choi, H. C., Wee, I., Corah, M., Sabet, S., Kim, T., Touma, T., Shim, D. H., and Agha-mohammadi, A. (2020). BAXTER: Bi-modal Aerial-Terrestrial Hybrid Vehicle for Long-endurance Versatile Mobility. *International Symposium on Experimental Robotics*.
- Choudhary, S., Carlone, L., Nieto, C., Rogers, J., Christensen, H. I., and Dellaert, F. (2017). Distributed mapping with privacy and communication constraints: Lightweight algorithms and object-based models. *International Journal of Robotics Research*, 36(12):1286–1311.
- CoSTAR Team (2018a). Rollocopter early versions. <https://youtu.be/XHltzQv3cTE>.
- CoSTAR Team (2018b). Scout tunnel experiment. https://www.youtube.com/watch?v=MdYXu_ImXik.
- CoSTAR Team (2020a). Autonomous Spot demonstration in urban environment. <https://youtu.be/VitjPLRdY8g>.
- CoSTAR Team (2020b). Mars Dogs: AI powered autonomous robot concept for Mars. <https://youtu.be/A8YR2J6KTxU>.
- CoSTAR Team (2020c). NeBula Autonomy Highlights in Team CoSTAR's Multi-robot Operations in the DARPA Subterranean Challenge Tunnel and Urban Circuit. <https://youtu.be/O4SZxZmQh-Q>.
- CoSTAR Team (2020d). Robotic exploration over extreme terrain in Martian-analog lava tubes. https://youtu.be/Apy-PiZK_d0.
- CoSTAR Team (2021). Search for life: NASA JPL explores Martian-like caves. <https://youtu.be/qTWDZr4U8>.
- Cushing, G. (2012). Candidate cave entrances on Mars. *Journal of Cave and Karst Studies*, 74:33–47.
- DARPA (2018a). DARPA Subterranean Challenge. <https://www.subtchallenge.com>.
- DARPA (2018b). SubT Tech Repo. <https://www.subtchallenge.world/models>.
- DARPA (2021). DARPA Subterranean Challenge: Scoring Metrics. https://www.subtchallenge.com/resources/SubT_Challenge_Finals_Rules.pdf.
- de la Croix, J.-P., Lim, G., Vander Hook, J., Rahmani, A., Droge, G., Xydes, A., and Scrapper, Jr., C. (2017). Mission modeling, planning, and execution module for teams of unmanned vehicles. In *Unmanned Systems Technology XIX*, volume 10195, page 101950J. International Society for Optics and Photonics.
- Dellaert, F. (2012). Factor graphs and GTSAM: A hands-on introduction. Technical report, Georgia Institute of Technology.
- Dixit, A., and Burdick, J. (2020). The kinematics of tracked vehicles via the power dissipation method.
- Ebadi, K., Chang, Y., Palieri, M., Stephens, A., Hatteland, A., Heiden, E., Thakur, A., Funabiki, N., Morrell, B., Wood, S., Carlone, L., and Agha-mohammadi, A. (2020). LAMP: Large-Scale Autonomous Mapping and Positioning for exploration of perceptually-degraded subterranean environments. In *2020 International Conference on Robotics and Automation (ICRA)*.
- Ebadi, K., Palieri, M., Wood, S., Curtis, P., and Agha-mohammadi, A. (2021). DARE-SLAM: Degeneracy-Aware and Resilient Loop Closing in Perceptually-Degraded Environments. *Journal of Intelligent and Robotic Systems*, 102(1):1–25.
- Fakoorian, S., Palieri, M., Santamaria-Navarro, A., Guaragnella, C., Simon, D., and Agha-mohammadi, A. (2020). Maximum correntropy Kalman filter for orientation estimation with application to lidar inertial odometry. In *ASME 2020 Dynamic Systems and Control Conference*. doi: 10.1115/DSCC2020-3256.
- Fakoorian, S., Santamaria-Navarro, A., Lopez, B. T., Simon, D., and Agha-mohammadi, A. (2021). Towards Robust State Estimation by Boosting the Maximum Correntropy Criterion Kalman Filter with Adaptive Behaviors. In *IEEE Robotics and Automation Letters*, 6(3): 5469–5476.
- Fan, D. D., Thakker, R., Bartlett, T., Miled, M. B., Kim, L., Theodorou, E., and Agha-mohammadi, A.-a. (2019). Autonomous hybrid ground/aerial mobility in unknown environments. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3070–3077. IEEE.
- Fan, D. D., Agha-mohammadi, A., and Theodorou, E. A. (2020a). Deep learning tubes for tube MPC. *Robotics Science and Systems (RSS)*.
- Fan, D. D., Nguyen, J., Thakker, R., Alatur, N., Agha-mohammadi, A., and Theodorou, E. A. (2020b). Bayesian learning-based adaptive control for safety critical systems. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4093–4099.

- Fan, D. D., Otsu, K., Kubo, Y., Dixit, A., Burdick, J., and Agha-Mohammadi, A.-A. (2021). Step: Stochastic traversability evaluation and planning for safe off-road navigation. *arXiv preprint arXiv:2103.02828*.
- Fedschun, T. (2018). Thailand cave rescue clears hurdle, divers reach passageway that may lead to youth soccer team. <https://www.nytimes.com/2018/07/12/world/asia/thailand-cave-rescue-seals.html>.
- Ford, J., Jones, H., Sharif, K., and Whittaker, R. (2020). Technologies enabling the exploration of lunar pits. *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*.
- Forster, C., Carlone, L., Dellaert, F., and Scaramuzza, D. (2015a). Imu preintegration on manifold for efficient visual-inertial maximum-a-posteriori estimation. *Robotics: Science and Systems (RSS)*.
- Forster, C., Carlone, L., Dellaert, F., and Scaramuzza, D. (2015b). On-manifold preintegration theory for fast and accurate visual-inertial navigation. *IEEE Transactions on Robotics*, pages 1–18.
- Fox, M., and Long, D. (2002). PDDL+: Modelling continuous time-dependent effects.
- Funabiki, N., Morrell, B., Nash, J., and Agha-mohammadi, A. (2020). Range-aided pose-graph-based SLAM: Applications of deployable ranging beacons for unknown environment exploration. *IEEE Robotics and Automation Letters*, 6(1):48–55.
- Furgale, P., Rehder, J., and Siegwart, R. (2013). Unified temporal and spatial calibration for multi-sensor systems. In *IEEE/RSJ International Conference on Intelligent Robots & Systems*.
- Gelfand, N., Ikemoto, L., Rusinkiewicz, S., and Levoy, M. (2003). Geometrically stable sampling for the ICP algorithm. In *Proceedings of the Fourth International Conference on 3-D Digital Imaging and Modeling, 2003 (3DIM 2003)*, pages 260–267.
- Ginting, M. F., Touma, T., Edlund, J. A., Buscicchio, A., Otsu, K., and Agha, A. (2020). Deployable mesh network for enabling reliable communication from within subsurface voids to the planetary surface. In *AGU Fall Meeting*.
- Ginting, M. F., Otsu, K., Edlund, J. A., Gao, J., and Agha-mohammadi, A. (2021). CHORD: Distributed data-sharing via hybrid ROS 1 and 2 for multi-robot exploration of large-scale complex environments. *IEEE Robotics and Automation Letters*, 6(3):5064–5071.
- Heiden, E., Hausman, K., Sukhatme, G. S., and Agha-mohammadi, A.-a. (2017). Planning high-speed safe trajectories in confidence-rich maps. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2880–2886.
- Hess, W., Kohler, D., Rapp, H., and Andor, D. (2016). Real-time loop closure in 2D lidar SLAM. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1271–1278.
- Himmelsbach, M., Hundelshausen, F. V., and Wuensche, H.-J. (2010). Fast segmentation of 3D point clouds for ground vehicles. In *2010 IEEE Intelligent Vehicles Symposium (IV)*, pages 560–565.
- Hines, T., Stepanas, K., Talbot, F., Sa, I., Lewis, J., Hernandez, E., Kottege, N., and Hudson, N. (2021). Virtual surfaces and attitude aware planning and behaviours for negative obstacle navigation. *IEEE Robotics and Automation Letters*, 6(2):4048–4055.
- Hudson, N., Talbot, F., Cox, M., Williams, J., Hines, T., Pitt, A., Wood, B., Frousheger, D., Surdo, K. L., Molnar, T., Steindl, R., Wildie, M., Sa, I., Kottege, N., Stepanas, K., Hernandez, E., Catt, G., Docherty, W., Tidd, B., Tam, B., Murrell, S., Bessell, M., Hanson, L., Tychsen-Smith, L., Suzuki, H., Overs, L., Kendoul, F., Wagner, G., Palmer, D., Milani, P., O'Brien, M., Jiang, S., Chen, S., and Arkin, R. C. (2021). Heterogeneous ground and air platforms, homogeneous sensing: Team CSIRO Data61's approach to the DARPA Subterranean Challenge. *Journal of Field Robotics*, volume 2, pages 595–636.
- Issacs, J., Knoedler, K., Herdering, A., Beylik, M., and Quintero, H. (2021). Teleoperation for urban search and rescue applications. *Submitted to the Journal of Field Robotics*, volume 2, pages 1177–1190.
- Jenelten, F., Miki, T., Vijayan, A. E., Bjelonic, M., and Hutter, M. (2020). Perceptive locomotion in rough terrain—online foothold optimization. *IEEE Robotics and Automation Letters*, 5(4):5370–5376.
- Jung, S., Lee, H., Shim, D. H., and Agha-mohammadi, A.-a. (2021). Collision-free local planner for unknown subterranean navigation. *ETRI Journal*, 43(4):580–593.
- Kaelbling, L., Littman, M., and Cassandra, A. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134.
- Kaelbling, L. P., and Lozano-Pérez, T. (2011). Planning in the know: Hierarchical belief-space task and motion planning. In *Workshop on Mobile Manipulation, IEEE International Conference on Robotics and Automation*.
- Kaess, M., Johannsson, H., Roberts, R., Ila, V., Leonard, J. J., and Dellaert, F. (2012). iSAM2: Incremental smoothing and mapping using the Bayes tree. *International Journal of Robotics Research*, 31(2):216–235.
- Kalantari, A., and Spenko, M. (2014). Modeling and performance assessment of the HyTAQ, a hybrid terrestrial/aerial quadrotor. *IEEE Transactions on Robotics*, 30(5):1278–1285.

- Kalantari, A., Touma, T., Kim, L., Jitsho, R., Strickland, K., Lopez, B. T., and Agha-Mohammadi, A. (2020). Drivocopter: A concept hybrid aerial/ground vehicle for long-endurance mobility. In *2020 IEEE Aerospace Conference*, pages 1–10.
- Kalita, H., Morad, S., Ravindran, A., and Thangavelautham, J. (2018). Path planning and navigation inside off-world lava tubes and caves. In *2018 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 1311–1318.
- Kaufmann, M., Vaquero, T. S., Correa, G. J., Otsu, K., Ginting, M. F., Beltrame, G., and Agha-mohammadi, A. (2021). Copilot MIKE: An autonomous assistant for multi-robot operations in cave exploration. In *2021 IEEE Aerospace Conference*, pages 1–9. IEEE.
- Kesner, S. B., Plante, J.-S., Boston, P. J., Fabian, T., and Dubowsky, S. (2007). Mobility and power feasibility of a microbot team system for extraterrestrial cave exploration. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 4893–4898.
- Khattak, S., Mascarich, F., Dang, T., Papachristos, C., and Alexis, K. (2019). Robust thermal-inertial localization for aerial robots: A case for direct methods. In *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 1061–1068.
- Kim, S.-K., Salzman, O., and Likhachev, M. (2019a). POMHDP: Search-based belief space planning using multiple heuristics. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, volume 29, pages 734–744.
- Kim, S.-K., Thakker, R., and Agha-Mohammadi, A. (2019b). Bi-directional value learning for risk-aware planning under uncertainty. *IEEE Robotics and Automation Letters*, 4(3):2493–2500.
- Kim, S.-K., Bouman, A., Salhotra, G., Fan, D. D., Otsu, K., Burdick, J., and Agha-mohammadi, A. (2021). PLGRIM: Hierarchical value learning for large-scale exploration in unknown environments. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, volume 31.
- Kinsey, J. C., Eustice, R. M., and Whitcomb, L. L. (2006). A survey of underwater vehicle navigation: Recent advances and new challenges. In *IFAC Conference of Manoeuvring and Control of Marine Craft*, volume 88, pages 1–12. Lisbon.
- Kramer, A., Stahoviak, C., Santamaria-Navarro, A., Agha-mohammadi, A., and Heckman, C. (2020). Radar-inertial ego-velocity estimation for visually degraded environments. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5739–5746.
- Lajoie, P.-Y., Ramtoula, B., Chang, Y., Carlone, L., and Beltrame, G. (2020). DOOR-SLAM: Distributed, online, and outlier resilient SLAM for robotic teams. *IEEE Robotics and Automation Letters*, 5(2):1656–1663.
- Léveillé, R. J., and Datta, S. (2010). Lava tubes and basaltic caves as astrobiological targets on Earth and Mars: A review. *Planetary and Space Science*, 58(4):592–598.
- Lew, T., Emmei, T., Fan, D. D., Bartlett, T., Santamaria-Navarro, A., Thakker, R., and Agha-mohammadi, A.-a. (2019). Contact inertial odometry: Collisions are your friends. *Proceedings of the International Symposium on Robotics Research*.
- Littman, M. L., Cassandra, A. R., and Kaelbling, L. P. (1995). Learning policies for partially observable environments: Scaling up. In *Machine Learning Proceedings 1995*, pages 362–370. Elsevier.
- Lu, C.-L., Huang, C.-I., Huang, J.-T., Hsu, C.-C., Chang, P.-K., Ewe, Z. L., Huang, P.-J., Huang, Y.-Y., Huang, S.-C., Li, P.-L., Wang, B.-H., Yim, L.-S., Huang, S.-W., Liu, Z.-Y., Bai, M.-H., and Wang, H.-C. (2021). A heterogeneous unmanned ground vehicle and blimp robot team for search and rescue using data-driven autonomy and communication-aware navigation. *Journal of Field Robotics*, volume 2, pages 557–594.
- Mangelson, J. G., Dominic, D., Eustice, R. M., and Vasudevan, R. (2018). Pairwise consistent measurement set maximization for robust multi-robot map merging. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2916–2923.
- Meier, L., Honegger, D., and Pollefeys, M. (2015). PX4: A node-based multithreaded open source robotics framework for deeply embedded platforms. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6235–6240.
- Millane, A., and Taylor, Z. (2019). evo: Python package for the evaluation of odometry and SLAM. https://github.com/ethz-asl/lidar_align.
- Miller, I. D., Cladera, F., Cowley, A., Shivakumar, S. S., Lee, E. S., Jarin-Lipschitz, L., Bhat, A., Rodrigues, N., Zhou, A., Cohen, A., et al. (2020). Mine tunnel exploration using multiple quadrupedal robots. *IEEE Robotics and Automation Letters*, 5(2):2840–2847.

- Morrell, B., Palieri, M., Funabiki, N., Thakur, A., Agha, A., and Blank, J. G. (2020). Robotic localization and multi-sensor, semantic 3D mapping for exploration of subsurface voids. In *AGU Fall Meeting 2020*. AGU.
- Mur-Artal, R., and Tardós, J. D. (2017). ORB-SLAM2: An open-source SLAM system for monocular, stereo and RGB-D cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262.
- Nelson, E. (2016). Berkley localization and mapping. <https://github.com/erik-nelson/blam>.
- Ohradzansky, M. T., Rush, E. R., Riley, D. G., Mills, A. B., Ahmad, S., McGuire, S., Biggie, H., Harlow, K., Miles, M. J., Frew, E. W., Heckman, C., and Humbert, J. S. (2021). Multi-agent autonomy: Advancements and challenges in subterranean exploration. *Submitted to the Journal of Field Robotics*, volume 2, pages 1068–1104.
- Orekhov, V. L., and Chung, T. H. (2021). The DARPA Subterranean Challenge: A synopsis of the circuits stage. *Submitted to the Journal of Field Robotics*, volume 2, pages 735–747.
- Otsu, K., Tepsuporn, S., Thakker, R., Vaquero, T. S., Edlund, J. A., Walsh, W., Miles, G., Heywood, T., Wolf, M. T., and Agha-mohammadi, A. (2020). Supervised autonomy for communication-degraded subterranean exploration by a robot team. In *IEEE Aerospace Conference*.
- Palieri, M., Morrell, B., Thakur, A., Ebadi, K., Nash, J., Chatterjee, A., Kanellakis, C., Carlone, L., Guaragnella, C., and Agha-mohammadi, A. (2020). LOCUS: A multi-sensor lidar-centric solution for high-precision odometry and 3D mapping in real-time. *IEEE Robotics and Automation Letters*, 6(2):421–428.
- Pineau, J., Gordon, G., and Thrun, S. (2003). Point-based value iteration: An anytime algorithm for POMDPs. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1025–1032, Acapulco, Mexico.
- Qin, Y., Pedram, S. A., Feyzabadi, S., Allan, M., McLeod, A. J., Burdick, J. W., and Azizian, M. (2020). Temporal segmentation of surgical sub-tasks through deep learning with multiple data sources. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 371–377.
- Qualcomm Technologies, Inc. Machine Vision SDK. <https://developer.qualcomm.com/software/machine-vision-sdk>.
- Ramtoula, B., Caccavale, A., Beltrame, G., and Schwager, M. (2020). MSL-RAPTOR: A 6DoF relative pose tracker for onboard robotic perception. *arXiv preprint arXiv:2012.09264*.
- Redmon, J., and Farhadi, A. (2018). YOLOv3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.
- Rehder, J., Nikolic, J., Schneider, T., Hinzmann, T., and Siegwart, R. (2016). Extending kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes. In *Proceedings of the IEEE Conference on Robotics and Automation*.
- Rosinol, A., Abate, M., Chang, Y., and Carlone, L. (2020). Kimera: An open-source library for real-time metric-semantic localization and mapping. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1689–1696.
- Rouček, T., Pecka, M., Čížek, P., Petříček, T., Bayer, J., Šalanský, V., Heřt, D., Petrлік, M., Baca, T., Spurny, V., Krátký, V., Petráček, P., Baril, D., Vaidis, M., Kubelka, V., Pomerleau, F., Faigl, J., Zimmermann, K., Saska, M., Svoboda, T., and Krajník, T. (2021). System for multi-robotic exploration of underground environments CTU-CRAS-NORLAB in the DARPA Subterranean Challenge. *Submitted to the Journal of Field Robotics*.
- Sabet, S., Agha-Mohammadi, A., Tagliabue, A., Elliott, D. S., and Nikraves, P. E. (2019). Rollocopter: An energy-aware hybrid aerial-ground mobility for extreme terrains. In *IEEE Aerospace Conference*, pages 1–8.
- Sabet, S., Poursina, M., Nikraves, P. E., Reverdy, P., and Agha-Mohammadi, A.-A. (2020). Dynamic modeling, energy analysis, and path planning of spherical robots on uneven terrains. *IEEE Robotics and Automation Letters*, 5(4):6049–6056.
- Santamaria-Navarro, A., Thakker, R., Fan, D. D., Morrell, B., and Agha-mohammadi, A. (2019). Towards resilient autonomous navigation of drones. *Proceedings of the International Symposium on Robotics Research*.
- Santana, P. F., Barata, J., and Correia, L. (2007). Sustainable robots for humanitarian demining. *International Journal of Advanced Robotic Systems*, 4(2):23.
- Sasaki, T., Otsu, K., Thakker, R., Haesaert, S., and Agha-mohammadi, A. (2020). Where to map? Iterative rover-copter path planning for mars exploration. *IEEE Robotics and Automation Letters*, 5(2):2123–2130.

- Scherer, S., Agrawal, V., Best, G., Cao, C., Cujic, K., Darnley, R., DeBortoli, R., Dexheimer, E., Drozd, B., Garg, R., Higgins, I., Keller, J., Kohanbash, D., Nogueira, L., Pradhan, R., Tatum, M., K. Viswanathan, V., Willits, S., Zhao, S., Zhu, H., Abad, D., Angert, T., Armstrong, G., Boirum, R., Dongare, A., Dworman, M., Hu, S., Jaekel, J., Ji, R., Lai, A., Hsuan Lee, Y., Luong, A., Mangelson, J., Maier, J., Picard, J., Pluckter, K., Saba, A., Saroya, M., Scheide, E., Shoemaker-Trejo, N., Spisak, J., Teza, J., Yang, F., Wilson, A., Zhang, H., Choset, H., Kaess, M., Rowe, A., Singh, S., Zhang, J., A. Hollinger, G., and Travers, M. (2021). Resilient and modular subterranean exploration with a team of roving and flying robots. *Submitted to the Journal of Field Robotics*, volume 2, pages 678–734.
- Segal, A., Haehnel, D., and Thrun, S. (2009). Generalized-ICP. In *Robotics: Science and Systems*, volume 2, page 435. Seattle, WA.
- Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., and Rus, D. (2020). LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping. In *IEEE/RSJ IROS*.
- Silver, D., and Veness, J. (2010). Monte-Carlo planning in large POMDPs. In *Advances in Neural Information Processing Systems*, pages 2164–2172.
- Stamenković, V., Beegle, L., Zacny, K., Arumugam, D., Baglioni, P., Barba, N., Baross, J., Bell, M., Bhartia, R., Blank, J., et al. (2019). The next frontier for planetary and human exploration. *Nature Astronomy*, 3(2):116–120.
- Tagliabue, A., Schneider, S., Pavone, M., and Agha-mohammadi, A. (2020). Shapeshifter: A multi-agent, multi-modal robotic platform for exploration of titan. In *2020 IEEE Aerospace Conference*, pages 1–13.
- Tagliabue, A., Tordesillas, J., Cai, X., Santamaria-Navarro, A., How, J., Carlone, L., and Agha-mohammadi, A. (2021). LION: Lidar-Inertial Observability-aware Navigator for vision-denied environments. In *17th International Symposium on Experimental Robotics*, pages 380–390.
- Terry, E., Lei, X., Morrell, B., Daftry, S., and Agha-mohammadi, A. (2020). Object and gas source detection with robotic platforms in perceptually-degraded environments. In *RSS Workshop: Robots in the Wild: Challenges in Deploying Robust Autonomy for Robotic Exploration*.
- Thakker, R., Alatur, N., Fan, D. D., Tordesillas, J., Paton, M., Otsu, K., Toupet, O., and Agha-mohammadi, A. (2020). Autonomous off-road navigation over extreme terrains with perceptually-challenging conditions. In *17th International Symposium on Experimental Robotics (ISER)*.
- Titus, T., Wynne, J., Boston, P., De León, P., Demirel-Floyd, C., Jones, H., Sauro, F., Uckert, K., Agha-Mohammadi, A., Alexander, E., Jr., Ashley, J., Azua, A., Chiao, L., Cushing, G., Dedecker, J., Fairén, A., Frumkin, A., De Waele, J., Harris, G., and Romo, R. (2020). Science and technology requirements to explore caves in our solar system. *White Paper*.
- Touma, T., Blank, J. G., Ginting, M. F., Patterson, C., and Agha, A. (2020). Mars dogs: Biomimetic robots for the exploration of Mars, from its rugged surface to its hidden caves. In *AGU Fall Meeting*.
- Tranzatto, M., Mascarich, F., Bernreiter, L., Godinho, C., Camurri, M., Khattak, S. M. K., Dang, T., Reijgwart, V., Loeje, J., Wisth, D., Zimmermann, S., Nguyen, H., Fehr, M., Solanka, L., Buchanan, R., Bjelonic, M., Khedekar, N., Valceschini, M., Jenelten, F., Dharmadhikari, M., Homberger, T., De Petris, P., Wellhausen, L., Kulkarni, M., Miki, T., Hirsch, S., Montenegro, M., Papachristos, C., Tresoldi, F., Carius, J., Valsecchi, G., Lee, J., Meyer, K., Wu, X., Nieto, J., Smith, A., Hutter, M., Siegwart, R., Mueller, M., Fallon, M., and Alexis, K. (2021). CERBERUS: Autonomous legged and aerial robotic exploration in the tunnel and urban circuits of the DARPA Subterranean Challenge. *Submitted to the Journal of Field Robotics*, volume 2, pages 274–324.
- Umari, H., and Mukhopadhyay, S. (2017). Autonomous robotic exploration based on multiple rapidly-exploring randomized trees. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1396–1402.
- Vaquero, T. S., Saboia da Silva, M., Otsu, K., Kaufmann, M., Edlund, J. A., and Agha-mohammadi, A. (2020). Traversability-aware signal coverage planning for communication node deployment in planetary cave exploration. In *Artificial Intelligence, Robotics and Automation in Space*.
- Vejpongsa, T. (2018). Flooding complicates cave search for thai soccer team. <https://apnews.com/article/be65215e881540fc9e63d2edfe98efa1>.
- Xiong, Y., Liu, H., Gupta, S., Akin, B., Bender, G., Kindermans, P.-J., Tan, M., Singh, V., and Chen, B. (2020). MobileDets: Searching for object detection architectures for mobile accelerators. *arXiv preprint arXiv:2004.14525*.
- Ye, H., Chen, Y., and Liu, M. (2019). Tightly coupled 3D lidar inertial odometry and mapping. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3144–3150.

Yurtsever, E., Lambert, J., Carballo, A., and Takeda, K. (2020). A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access*, 8:58443–58469.

Zhang, J., and Singh, S. (2014). LOAM: Lidar odometry and mapping in real-time. In *In Robotics: Science and Systems*, volume 2, page 9.

How to cite this article: Agha, A., Otsu, K., Morrell, B., Fan, D. D., Thakker, R., Santamaria-Navarro, A., Kim, S.-K., Bouman, A., Lei, X., Edlund, J., Ginting, M. F., Ebadi, K., Anderson, M., Pailevanian, T., Terry, E., Wolf, M., Tagliabue, A., Vaquero, T. S., Palieri, M., Tepsuporn, S., Chang, Y., Kalantari, A., Chavez, F., Lopez, B., Funabiki, N., Miles, G., Touma, T., Buscicchio, A., Tordesillas, J., Alatur, N., Nash, J., Walsh, W., Jung, S., Lee, H., Kanellakis, C., Mayo, J., Harper, S., Kaufmann, M., Dixit, A., Correa, G. J., Lee, C., Gao, J., Merewether, G., Maldonado-Contreras, J., Salhotra, G., Da Silva, M. S., Ramtoula, B., Fakoorian, S., Hatteland, A., Kim, T., Bartlett, T., Stephens, A., Kim, L., Bergh, C., Heiden, E., Lew, T., Cauligi, A., Heywood, T., Kramer, A., Leopold, H. A., Melikyan, H., Choi, H. C., Daftry, S., Toupet, O., Wee, I., Thakur, A., Feras, M., Beltrame, G., Nikolakopoulos, G., Shim, D., Carlone, L., & Burdick, J. (2022). NeBula: TEAM CoSTAR's robotic autonomy solution that won phase II of DARPA Subterranean Challenge. *Field Robotics*, 2, 1432–1506.

Publisher's Note: Field Robotics does not accept any legal responsibility for errors, omissions or claims and does not provide any warranty, express or implied, with respect to information published in this article.