Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

Safety Inspections and Gas Monitoring in Hazardous Mining Areas Shortly After Blasting Using Autonomous UAVs

Samuel Nordström[1], Nikolaos Stathoulopoulos[1], Niklas Dahlquist[1], Björn Lindqvist[1], Ilias Tevetzidis[1], Christoforos Kanellakis[1], and George Nikolakopoulos[1]

[1]Lulea tekniska universitet Institutionen for system- och rymdteknik

March 12, 2024

Abstract

This article presents the first ever fully autonomous UAV (Unmanned Aerial Vehicle) mission to perform gas measurements after a real blast in an underground mine. The demonstration mission was deployed around 40 minutes after the blast took place, and as such realistic gas levels were measured. We also present multiple field robotics experiments in different mines detailing the development process. The presented novel autonomy stack, denoted as the Routine Inspection Autonomy (RIA) framework, combines a risk-aware 3D path planning D + * , with 3D LiDAR-based global relocalization on a known map, and it is integrated on a custom hardware and a sensing stack with an onboard gas sensing device. In the presented framework, the autonomous UAV can be deployed in incredibly harsh conditions (dust, significant deformations of the map) shortly after blasting to perform inspections of lingering gases that present a significant safety risk to workers. We also present a *change detection* framework that can extract and visualize the areas that were changed in the blasting procedure, a critical parameter for planning the extraction of materials, and for updating existing mine maps. As will be demonstrated, the RIA stack can enable robust autonomy in harsh conditions, and provides reliable and safe navigation behavior for autonomous Routine Inspection missions.

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

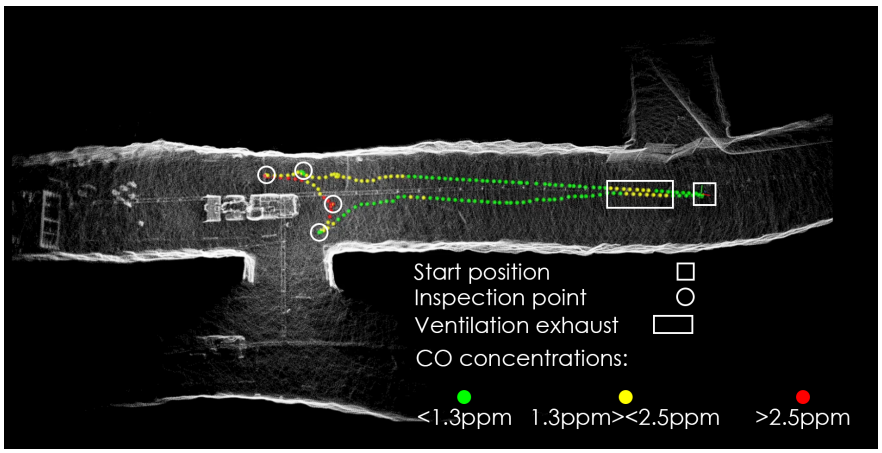

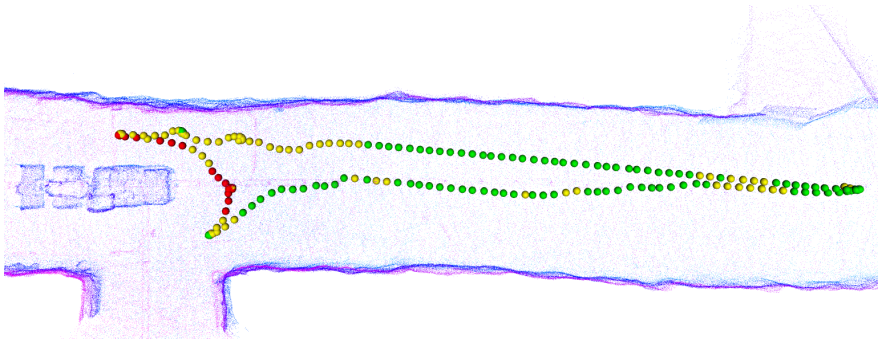

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

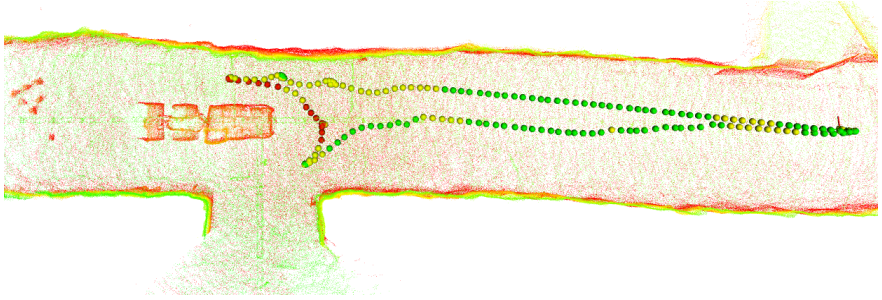

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

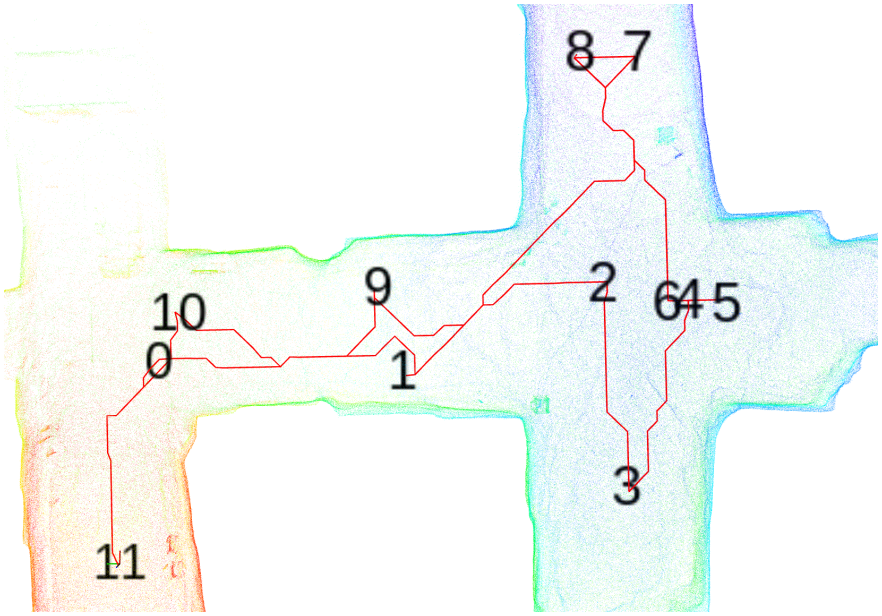

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

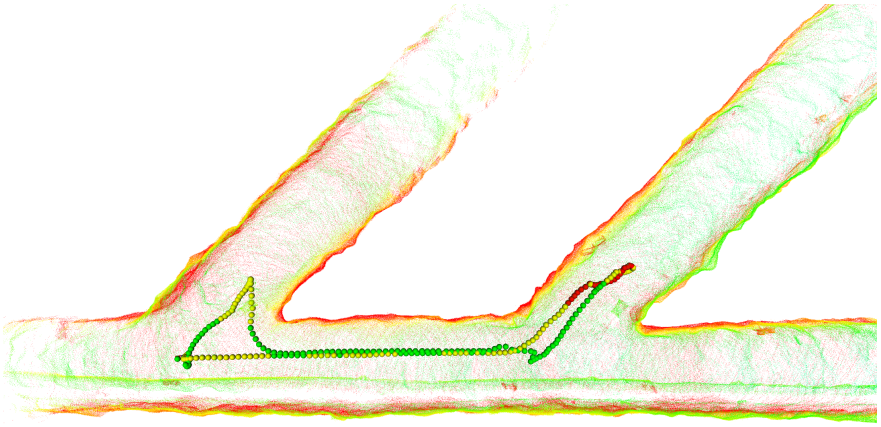

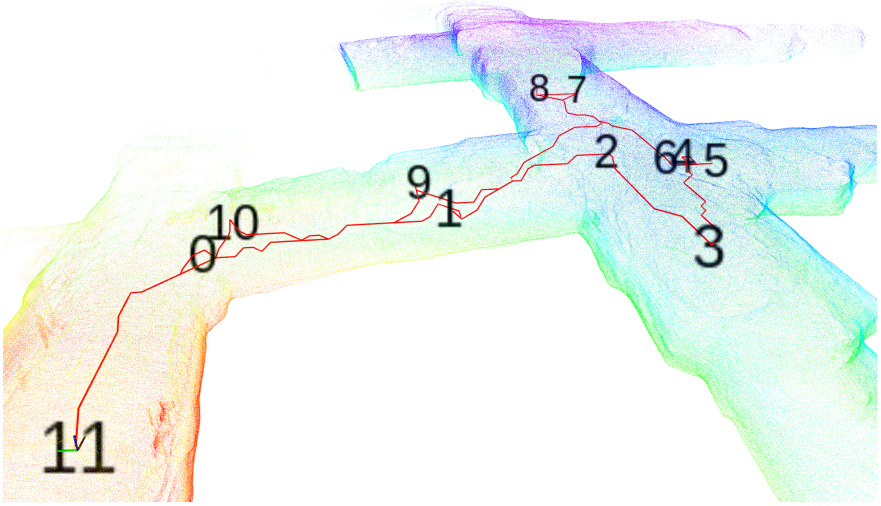

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

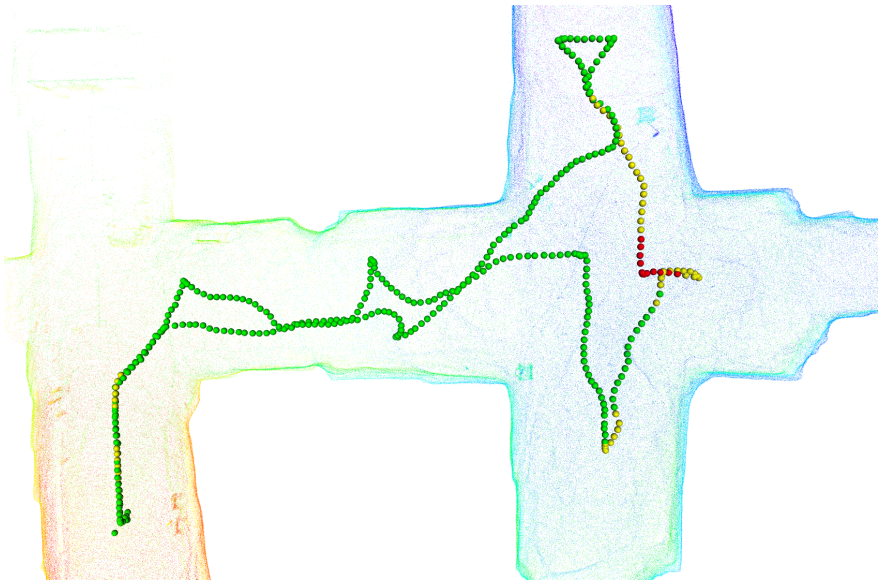

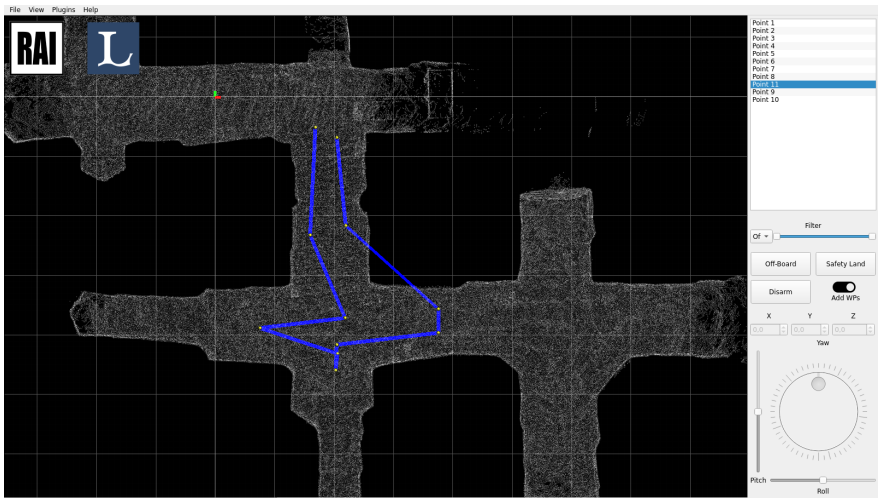

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

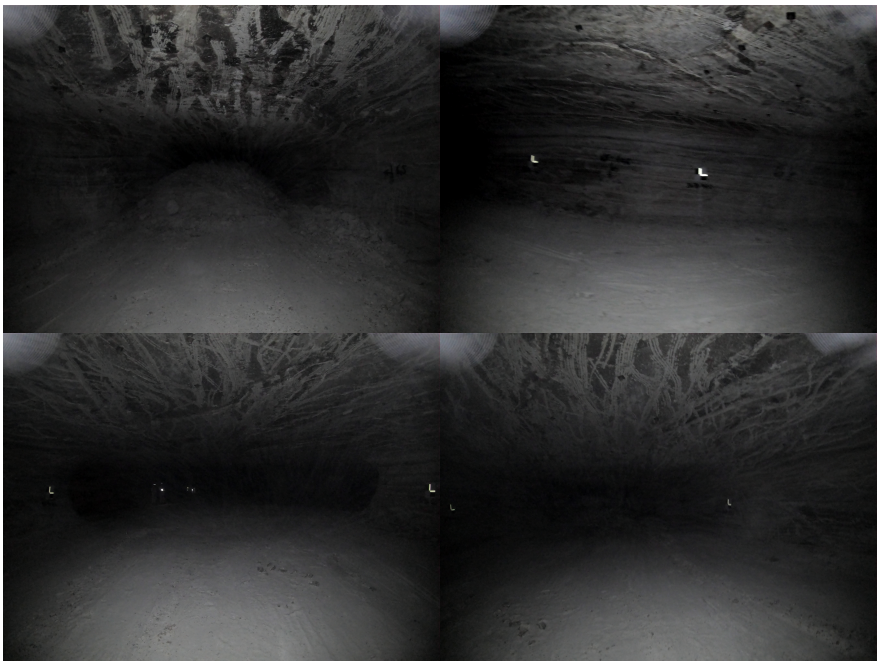

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

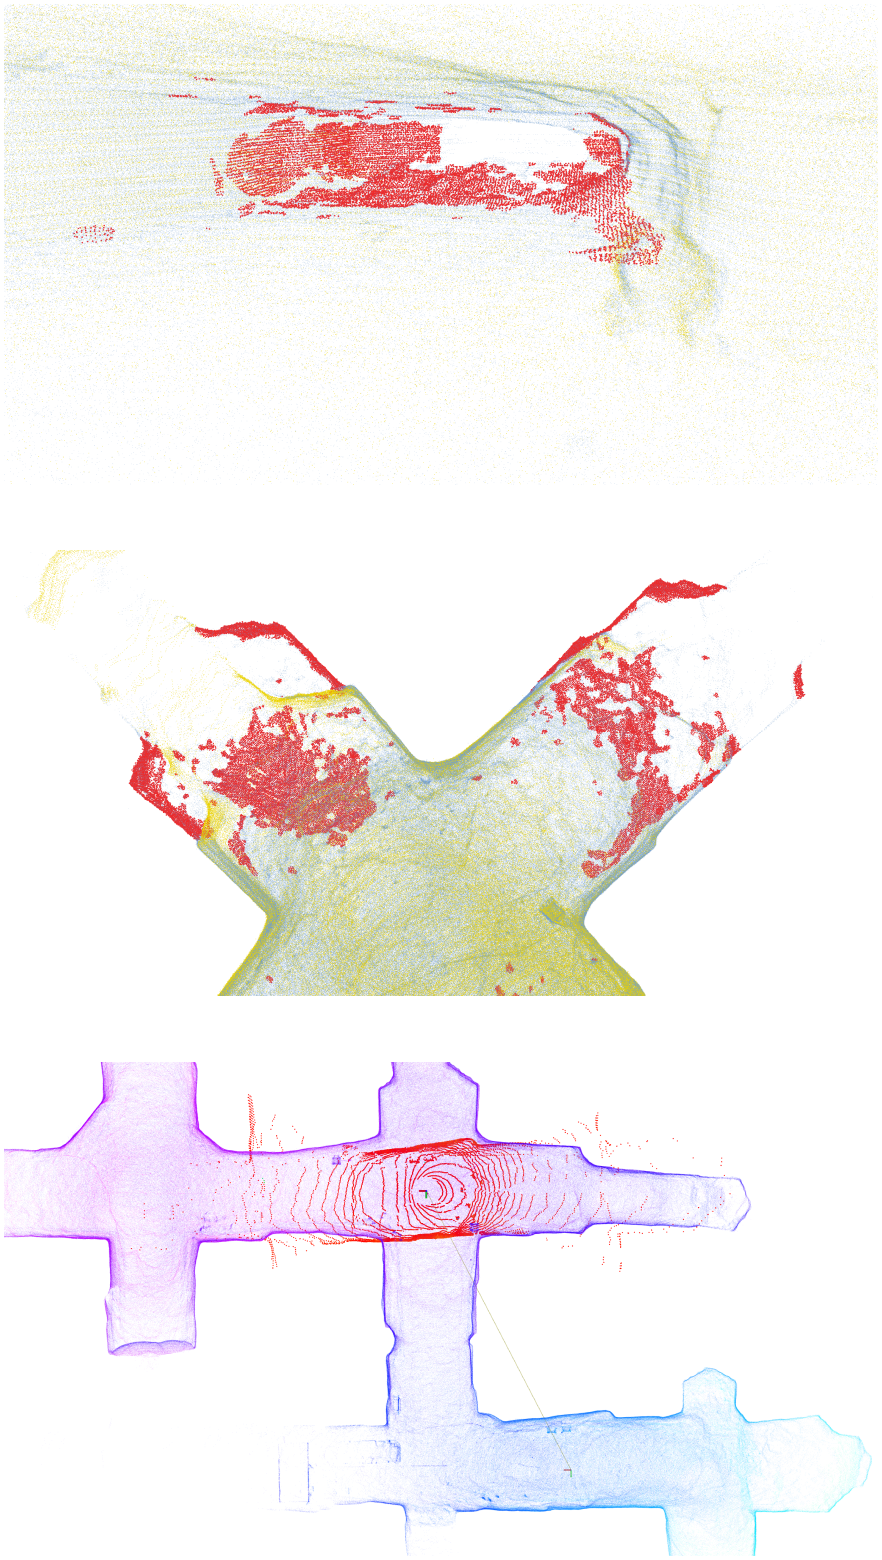

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

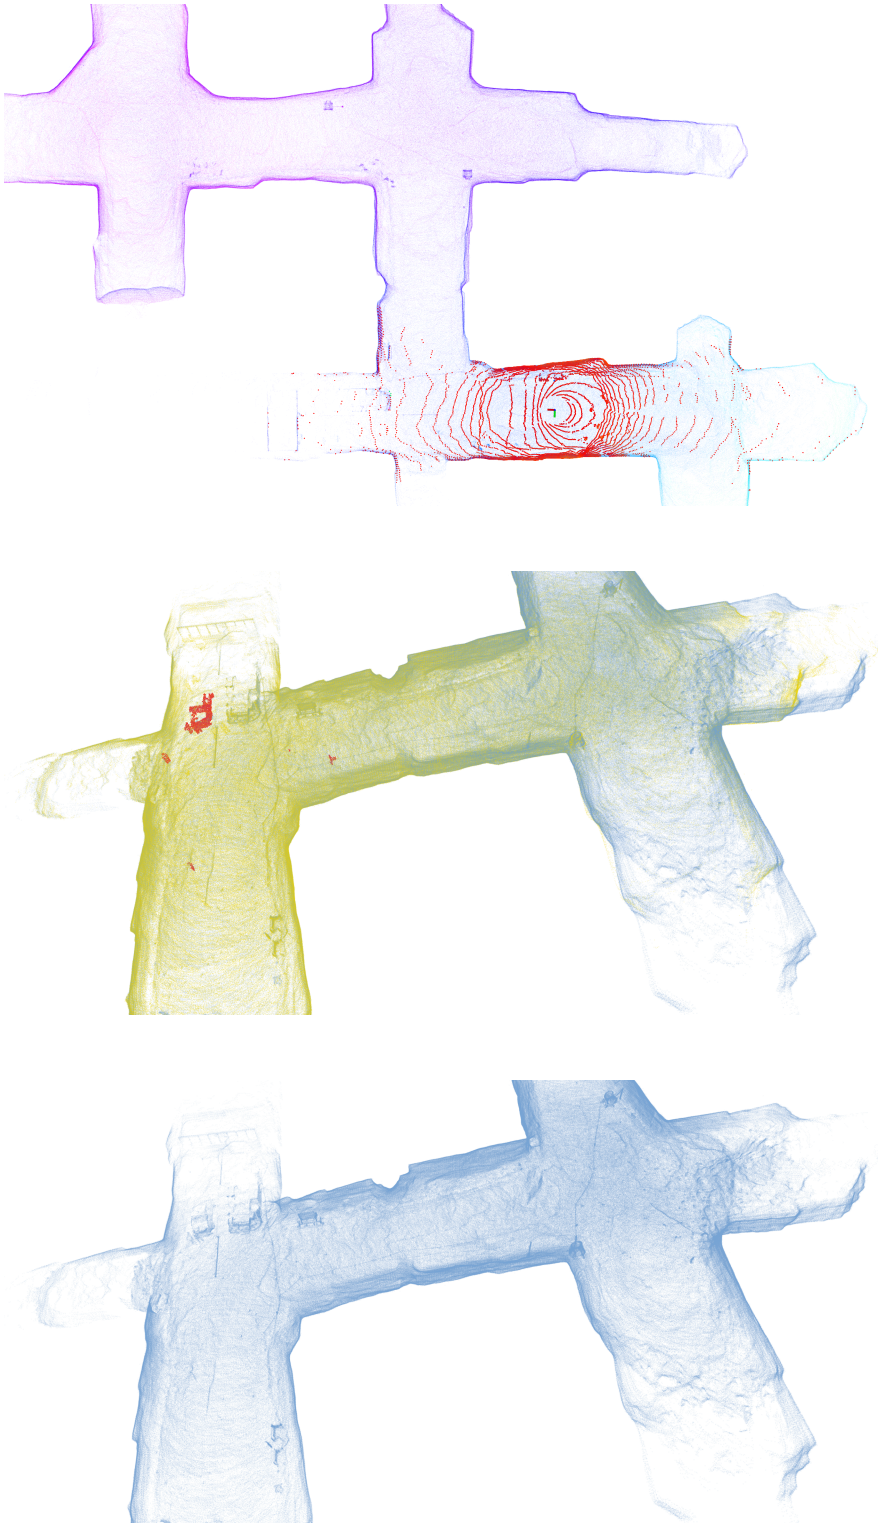

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

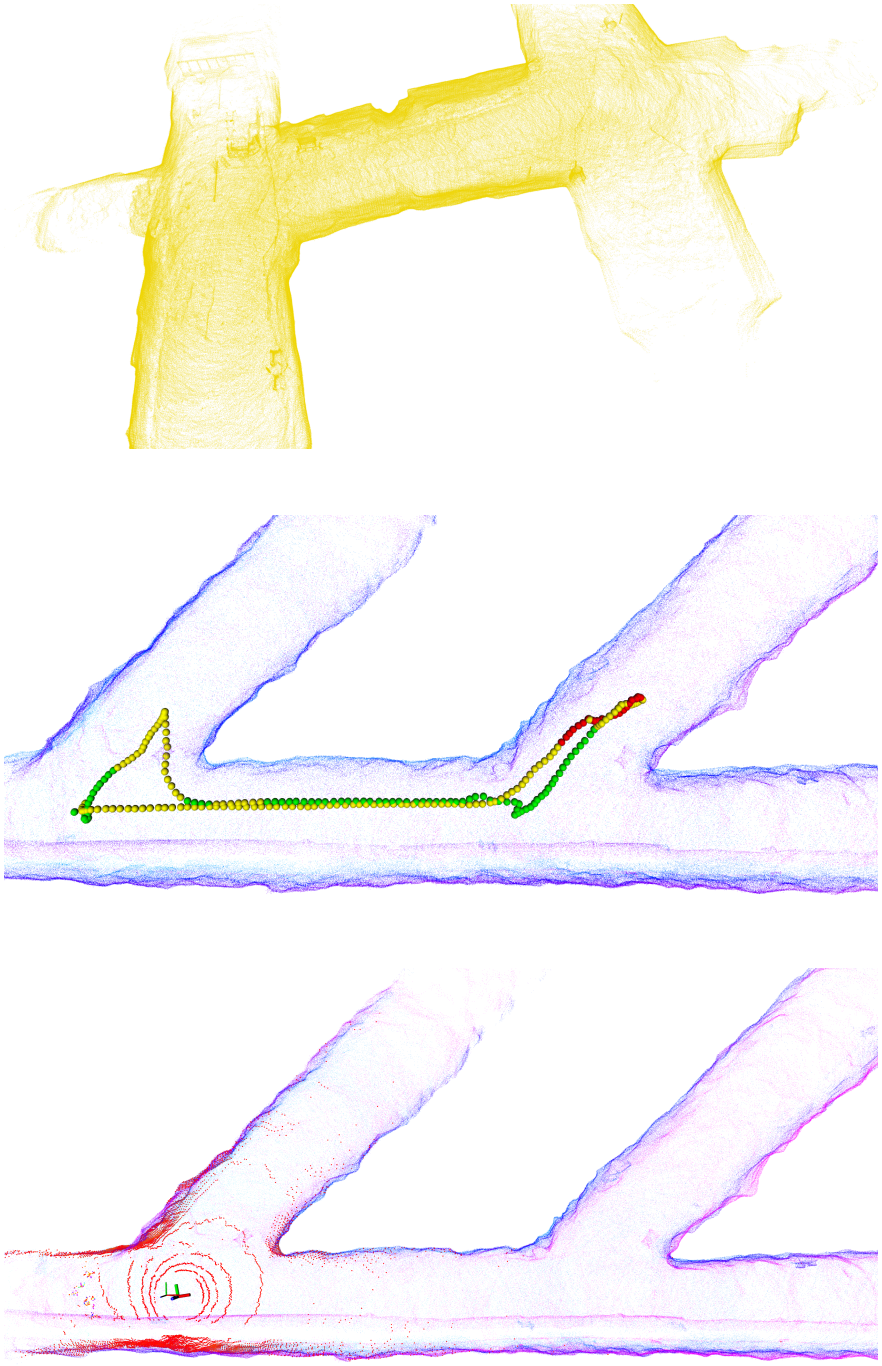

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

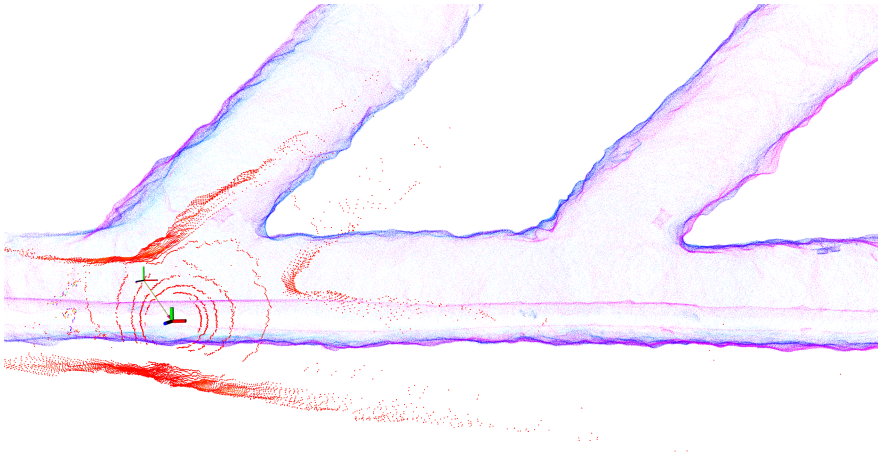

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

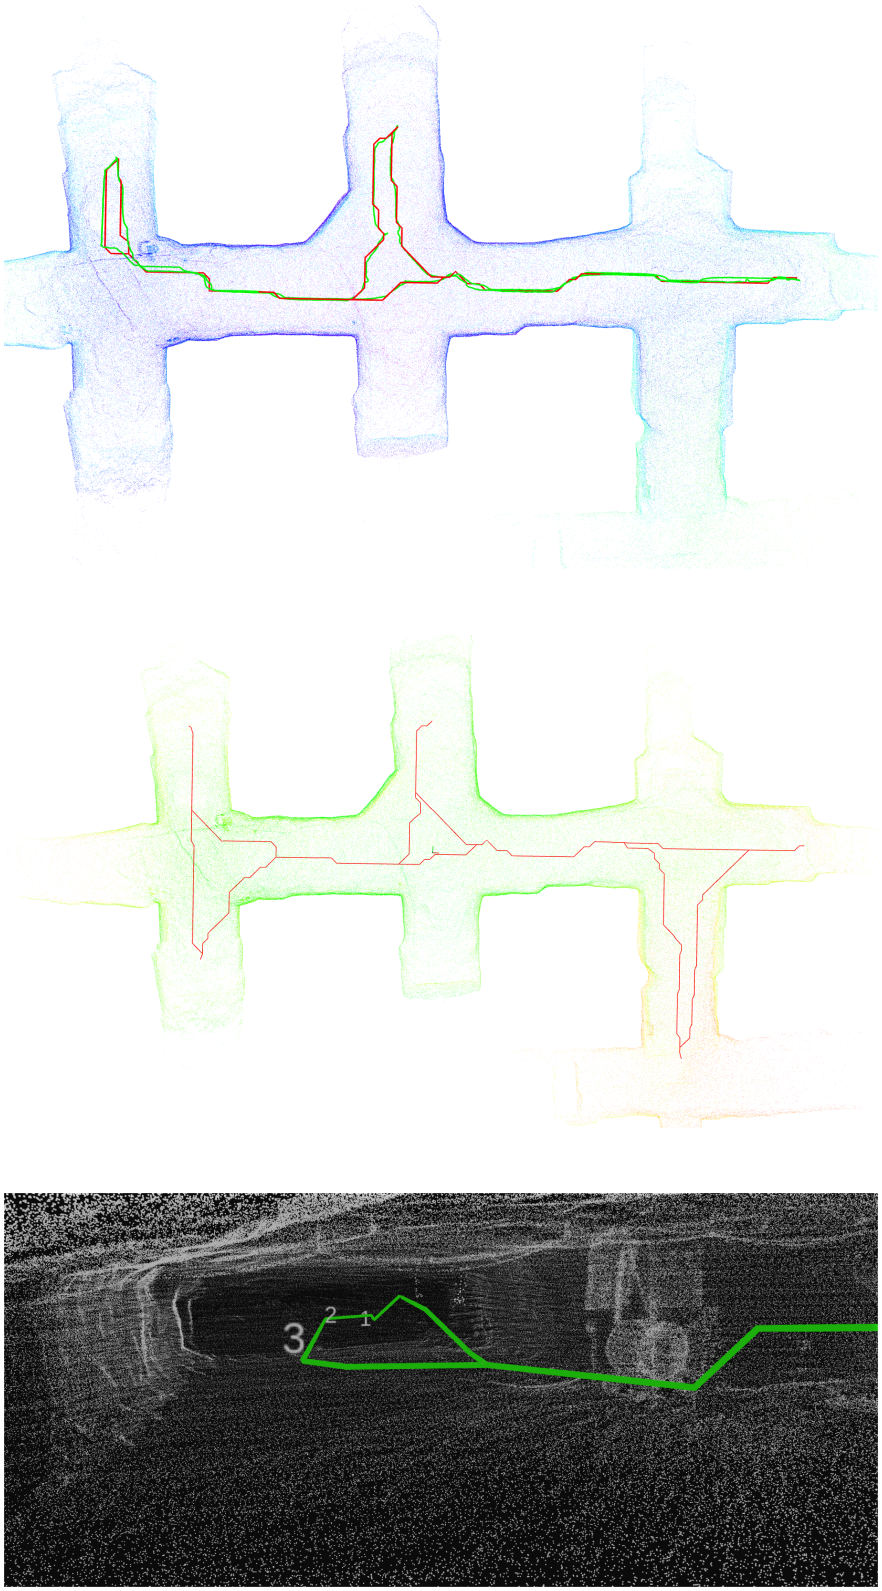

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

Posted on 12 Mar 2024 — The copyright holder is the author/funder. All rights reserved. No reuse without permission. — https://doi.org/10.22541/au.171027386.64290736/v1 — This is a preprint and has not been peer-reviewed. Data may be preliminary.

Safety Inspections and Gas Monitoring in Hazardous Mining Areas Shortly After Blasting Using Autonomous UAVs

Samuel Nordström*
Robotics & AI Group
Department of Computer,
Electrical and Space Engineering,
Luleå University of Technology,
Luleå SE-97187, Sweden
samkar@ltu.se

Nikolaos Stathoulopoulos
Robotics & AI Group
Department of Computer,
Electrical and Space Engineering,
Luleå University of Technology,
Luleå SE-97187, Sweden
niksta@ltu.se

Niklas Dahlquist
Robotics & AI Group
Department of Computer,
Electrical and Space Engineering,
Luleå University of Technology,
Luleå SE-97187, Sweden
nikdah@ltu.se

Björn Lindqvist
Robotics & AI Group
Department of Computer,
Electrical and Space Engineering,
Luleå University of Technology,
Luleå SE-97187, Sweden
bjolin@ltu.se

Ilias Tevetzidis
Robotics & AI Group
Department of Computer,
Electrical and Space Engineering,
Luleå University of Technology,
Luleå SE-97187, Sweden
ilitev@ltu.se

Christoforos Kanellakis
Robotics & AI Group
Department of Computer,
Electrical and Space Engineering,
Luleå University of Technology,
Luleå SE-97187, Sweden
chrkan@ltu.se

George Nikolakopulos
Robotics & AI Group
Department of Computer,
Electrical and Space Engineering,
Luleå University of Technology,
Luleå SE-97187, Sweden
geonik@ltu.se

Abstract

This article presents the first ever fully autonomous UAV (Unmanned Aerial Vehicle) mission to perform gas measurements after a real blast in an underground mine. The demonstration

*This work has been partially funded by the European Unions Horizon 2020 Research and Innovation Programme under the Grant Agreement No. 101003591 NEX-GEN SIMS, and partially by the Sustainable Underground Mining (SUM) Academy Programme under project SP14. **Acknowledgments** Thanks to the K+S Group, Epirock Rock Drills AB, and LKAB (Luossavaara-Kiirunavaara AB) for the suport and access to mine environments and opportunity to perform the field evaluations presented in this article

mission was deployed around 40 minutes after the blast took place, and as such realistic gas levels were measured. We also present multiple field robotics experiments in different mines detailing the development process. The presented novel autonomy stack, denoted as the Routine Inspection Autonomy (RIA) framework, combines a risk-aware 3D path planning D^*_+, with 3D LiDAR-based global relocalization on a known map, and it is integrated on a custom hardware and a sensing stack with an onboard gas sensing device. In the presented framework, the autonomous UAV can be deployed in incredibly harsh conditions (dust, significant deformations of the map) shortly after blasting to perform inspections of lingering gases that present a significant safety risk to workers. We also present a *change detection* framework that can extract and visualize the areas that were changed in the blasting procedure, a critical parameter for planning the extraction of materials, and for updating existing mine maps. As will be demonstrated, the RIA stack can enable robust autonomy in harsh conditions, and provides reliable and safe navigation behavior for autonomous Routine Inspection missions.

Keywords— Field Robotics, Mining Robotics, Unmanned Areal Vehicles, Gas Monitoring, Change Detection

1 Introduction

Robots are increasingly used more and more in real-world application areas and for different tasks (Zeng, 2022), and as Technology Readiness-Levels increase the interest in the adoption of robotic technologies is increasing in the mining industry as well. More specifically, mining operations can present significant safety risks to operators, and as such modern mines place a large focus on increasing the levels of safety in relation to the active day-to-day operations in the mine. Two such risks are direct products of the blasting operation to expand the mining tunnel: 1) the potential cave-in or partial tunnel collapse after a blast indicated by the state of the mining face, hanging walls, and surrounding tunnel, and 2) the risk of lingering gases not only from encountered gas pockets but mainly from gases produced by the explosives themselves. As such, after each blasting operation manual inspections must be performed to ensure structural integrity and low levels of gases in the areas surrounding the blast before the excavation process can begin. The idea is to remove all human operators from the hazardous area until it is declared safe through the use of inspection robots, and concepts surrounding improved safety by using robots to inspect mining areas have been around for some time (Basha and Patel, 2022; Zimroz et al., 2019; Reddy et al., 2015) for various applications. Utilizing robots in mines is not only a way of improving safety, but it can also save money by automating and increasing the efficiency of the overall routine inspection tasks to, for example, reduce idling times during the blasting cycle.

Unmanned Aerial Vehicles (UAVs), and rotorcraft in particular, are, due to their ability to operate in 3D while ignoring the ground terrain, showing great promise in mining inspection applications as discussed in Shahmoradi et al. (2020). In that survey, teleoperated (manual) UAVs are shown to inspect rock mass characteristics (Turner et al., 2018), inspection of the distribution of rock fragmentation (or muck piles) after blasting (Bamford et al., 2017), and for providing visual inspection of unreachable areas (Freire and Cota, 2017). The survey mentions that for current teleoperations, two of the main challenges for UAVs in underground mines are the communication and connectivity for out-of-line-of-sight operations (Forooshani et al., 2013; Ranjan et al., 2014) and the inability for drone pilots to access nearby safe areas to operate from when it comes to inspections of high-risk areas. This problem can be mitigated completely by autonomous robots that do not rely on any communication with the operator during mission execution, and whose autonomy and sensing stack enables it to operate without any installed infrastructure in the mine. The quest remains in designing such systems that can reliably and safely execute large-scale autonomous inspection missions in the harsh mining environments, and integrating them with the mining workflow and the operators for performing actually useful tasks in the mine.

This manuscript presents developments in that direction, aimed specifically at safety inspection missions after blasting, which is one of the critical inspection use-cases. We present the RIA (Routine Inspection Autonomy) stack, that is utilizing a combination of risk-aware 3D path planning and 3D LiDAR relocalization modules. This is supported by an easy-to-use user interface which enables an operator to easily click high-level waypoints to specify the inspection mission, and as such the framework operates under the assumption that there exists an apriori map of the area that one can define areas to be inspected in. We also present a multi-sensing approach to safety inspections where we can

provide a visual inspection of the state of the drift where the blast took place, onboard gas sensing in fused with the onboard localization data, and a point cloud-based framework for *change detection* that can segment out the areas of the mine map that has changed between routine inspections (before and after blasting) in order to update mine maps and plan the excavation process.

The background of this work lies in the Horizon 2020 project Next Generation Carbon Neutral Pilots for Smart Intelligent Mining Systems (NEXGEN SIMS) Horizon (2020) and the Sustainable Underground Mining (SUM) Academy Programme, and the field trials demonstrated were only made possibly by our mining industry collaborators at the K+S Group, at LKAB (Luossavaara-Kiirunavaara AB), and Epirock Rock Drills AB. The main field demonstration included in this work took place around 40 minutes after a large-scale blasting operation at a depth of 700 m in a potash mine operated by the K+S Group in Germany, and mimics the real field conditions where such technologies could be adopted perfectly. As such, this manuscript provides, in addition to the proposed autonomy kit, also a high-impact technology demonstrator with great potential to reduce the risk to human operators during post-blast inspections. Supporting the large-scale demonstration, we include additional smaller scale experiments at other relevant underground mining sites showing the progression of the demonstrator and its use in wider range of scenarios.

2 Related works

2.1 Mining Robotics & Autonomous Gas Monitoring

Using autonomous robots for gas monitoring and other inspection tasks in the mining industry has been an area of interest in the last years. Source localization or *sniffing* algorithms have been trialed to search indoor areas for high gas concentrations (Arain et al., 2021), while initial trials for using autonomous ground robots to detect gases, cracks, etc. in mines and tunnels have been attempted (Kim and Choi, 2023; Protopapadakis et al., 2016) in smaller-scale missions. While UAV systems have lower battery life and less carrying capacity than a wheeled ground robot, they can enter areas regardless of the ground terrain, which in this work will be demonstrated by navigating into mining drifts that were just blasted to perform inspections despite the rubble on the ground. Detecting gases with an UAV has its negatives, the propellers will mix the gases around the UAV and change the pressure (which can affect measurements), but it is still possible to measure gases as shown in Rossi and Brunelli (2015). This related work only used basic autonomy in a lab setting to prove the concept of measuring gases with a drone.

Many works are motivated by search-and-rescue scenarios in subterranean areas (Wang et al., 2014; Zhao et al., 2017; Gomathi et al., 2015; Li et al., 2020; Lindqvist et al., 2022b) where gasses are one of the potential dangers that rescue workers need to consider. Similar rescue scenarios were also of great interest in the recent DARPA Subterranean Challenge (DARPA, 2020), where teams of robots were tasked to search vast underground areas to find and localize objects of interest. The SubT Challenge significantly pushed the state-of-the-art of localization and exploratory navigation systems for subterranean environments (Tranzatto et al., 2022; Rouček et al., 2020; Agha et al., 2022; Lindqvist et al., 2022c).

Autonomous UAVs, as opposed to teleoperated, have been deployed in mines previously to inspect subterranean or mining tunnels (Kanellakis et al., 2018; Papachristos et al., 2019). The significant contributions and differences to the state-of-the-art in this manuscript is mainly the scale of the mission and the fact that the mission was performed in the context of a direct industry use-case and in real field conditions right after a real blast. Additionally, our work presents a complete mission execution scenario in the context of how such technologies could be used by the industry in the future, using an advanced autonomy kit capable of handling the harsh environments. This manuscript also includes a multi-modal sensing and inspection capability, where we perform real-time measurements for lingering gases, visual safety inspection of the blasted drifts and surrounding area, as well as 3D reconstruction and change detection using 3D LiDAR point clouds. To the best of the authors' ability, no autonomous robotic gas inspection mission has previously been performed shortly after a blast in an underground mine.

2.2 Path Planning & Navigation

To perform a fully autonomous gas monitoring mission, a complete-coverage path planner such as in Jiang et al. (2022) is one option. Coverage planners generally focus on designing trajectories that fully inspect a certain area based on some mission criteria, such as covering the surface of a structure completely with a camera footprint. In our

mission use-case, such a planner would need to run online as the post-blast map will have significant deformations and blockages, and might produce an unnecessarily long mission as the environment is large. Another option is a source localization (Bhatt et al., 2018) algorithm to guide the UAV towards high gas concentrations. This would be difficult to apply for consistent inspection mission execution in our use-case, as there is no guarantee of a gas-concentration gradient leading the UAV to critical gas pockets in the blast zone. To enable the high-TRL mission demonstrations in this manuscript, we are going down a more straight-forward route of visiting the most important inspection waypoints that can provide the mine operators with the information required, in a compromise between full coverage and the length of the mission. To do so, the most common approach is occupancy map (Hornung et al., 2013a) based path planning, such as sampling based RRT LaValle et al. (2001) (Rapidly-exploring Random Trees), grid search algorithms like A^* (Warren, 1993), more modern versions such as Jump-Point Search (Harabor and Grastien, 2014), or a myriad of other variations (Xiong et al., 2021; Wang et al., 2021; Hornung et al., 2013b; Zhang et al., 2021) to the same problem. Classic map-based waypoint navigation is not a novel or particularly hard task, and has been extensively applied in robotic navigation. What is challenging is to do it in sensor deprecated areas that are dark, dusty, and with potentially large changes in the environment like the ones caused by blasts in mines, and most importantly on degraded or imperfect occupancy maps, that are not fully dense and where there could be many gaps in the maps. One solution to gaps in the map is filtering (Atapour-Abarghouei and Breckon, 2018) but it is hard to filter gaps while keeping actual openings, making the solution dependent on good tuning for the specific use-case and environment. Another problem is generating robot-safe paths that take the dimensions of the robot into account such that following the path does not lead to collisions. This issue has been attempted to be solved via inflation of the occupied space (Li et al., 2018). These challenges can also be solved with path planners that have knowledge of the robots' size and dynamics (Yan F., 2013; Tordesillas et al., 2019; Tordesillas and How, 2021; Dang et al., 2020). Such methods generally work well as they prevent the robot from going too close to any occupied space while also addressing the previously mentioned "gaps in the map" problem. Problems with these methods can arise when you desire to have large safety marginals to promote as safe as possible navigation, but there still are occasions where the robot must violate those outer margins to reach its destination. In this work, we are relying on our own developed algorithm denoted as D_+^* (Karlsson et al., 2022). D_+^* is an extension of D^*-lite (Koenig and Likhachev, 2002), where we add proximity risk costs to occupied voxels and treat unknown space (e.g. not set to free or occupied) as a risk. D_+^* is implemented with the ROS-integrated occupancy mapper OctoMap Hornung et al. (2013a) and can operate on both online and saved offline maps. The risk costs enable using large safety margins that promote safe behavior but allowing the UAV to violate them (the outer ones) when absolutely needed to reach its target. As such, we can say that D_+^* plans the shortest and safest path to the target based on the added risk heuristic. The risk costs also prevent impossible shortcuts or corner cutting due to incomplete maps or gaps in the occupied space. Finally, the tuning needed for D_+^* is minimal, as it is only a question of adjusting the voxel size of the occupancy map and the number of voxels needed for the desired safety marginal.

2.3 Global Localization and Relocalization

Several solutions exist for the problem of Simultaneous Localization And Mapping (SLAM) (Koval et al., 2022), but in the case of multiple autonomous inspection missions, it is important to be able to perform map-based navigation. The first step for map-based navigation is to estimate the robot's global pose inside the given point cloud map, and then continuously estimate the robots' position with respect to the initial coordinates of the map. In the current literature Stathoulopoulos et al. (2024b), there is a limited number available frameworks that support global relocalization within a 3D point cloud map. Koide et al. (2019) provided a series of packages that include global localization as part of the SLAM process, both with a 3D and a 2D solution. Branch and Bounce Search (BBS) (Hess et al., 2016) offers faster processing times but is vulnerable to uneven terrain or environments that include elevation and FPFH+RANSAC (Rusu et al., 2009; Fischler and Bolles, 1981) has slower computational times but is more robust when dealing with 3D scenes. More recently, LIO-SAM (Shan et al., 2020) and FAST-LIO (Xu and Zhang, 2021), two popular LiDAR Inertial Odometry modules, got an extension for providing relocalization and global pose estimation capabilities. The former relies on multiple factors, as it utilizes frequent key-frames and a sliding window scan-matching technique to jointly optimize a factor graph. The latter combines low-frequency global localization and high-frequency odometry through a two-stage ICP (Iterative Closest Point) registration and then combines the output with the IMU measurements in the state estimation module. In our framework, we utilize FAST-LIO localization on a known map, where our in-house developed 3DEG Stathoulopoulos et al. (2024a) relocalization framework plays a crucial role by offering the initial transformation estimate between the robot and map frames, which serves as the starting point for odometry initialization. 3DEG leverages learned place recognition descriptors that encompass both place recognition and yaw discrepancy regression capabilities, allowing it to provide an initial estimate for the subsequent ICP registration process. This integration ensures accurate initialization for maintaining precise localization within the map.

2.4 Change Detection

Detecting changes in the environment between different inspection missions is of vital importance since it can be used to update digital twins, monitor the progress of a facility, or even prevent future risks or disasters. In the current literature, the problem of change detection is commonly addressed using multi-temporal images (Goswami et al., 2022; Shuai et al., 2022; Seo et al., 2022) and more specifically using object-based methods or satellite images, since they have shown great performance and robustness. In the scenario of inspecting underground mines, though, the presence of dust particles as well as the lack of illumination can be fatal for such methods to produce a good result. Other methods have turned to occupancy grid maps (Xiao et al., 2013; Gojcic et al., 2021) produced by dense laser scanners, mainly utilized in urban areas. The negative aspect of such methods is the need for dense representations and their computational inefficiency, making them unsuitable for deploying in robotic autonomous missions. Last but not least, as deep learning methods are emerging (Urbach et al., 2020), they still hold their focus on urban areas or isolated objects, leaving a gap in the area of indoor environments or underground mines. To address these challenges, we have developed a framework (Stathoulopoulos et al., 2023b) inspired by machine learning techniques for autonomous robots. Our approach leverages place recognition descriptors, and more specifically, learning based descriptors, to identify changes between two bi-temporal point cloud maps and pinpoint areas of change. By employing a point-to-voxel comparison, we efficiently extract altered objects or obstacles, particularly suited for dynamic scenes and large-scale environments, where we can focus on specific changed regions and avoid costly computations across the entire point cloud maps.

2.5 Contributions

We can thus summarize the contributions stated in the previous sections as:

- The main contribution of this manuscript within field robotics. We demonstrate a direct industry-driven use-case deployed in multiple real and harsh field environments towards increased safety in the mining industry.
- We demonstrate for the first time ever an autonomous aerial gas monitoring mission performed after a real full-scale blast in an operational underground mining area. The outcome of the mission is multimodal sensing data in the form of updated 3D maps, visual imagery, and localized gas level data in the area.
- The manuscripts presents a novel sensing and autonomy stack denoted as the RIA (Routine Inspection Autonomy) framework that enabled fully autonomous inspection missions through risk-aware path planning and global relocalization, deployed from an operator-friendly GUI.
- We show that our learning-based change detection framework is capable of extracting critical information to assist in planning the ore extraction process.

3 Mission specifications

Before delving into the autonomy and hardware stack, we must specify what type of missions it is developed for and in what context. This section will aim to provide the background and specifications that led to the development of the presented autonomy.

3.1 Multi-sensor inspection after blasting

The fundamental aim of the presented mission is to gather the necessary data an operator needs in order to make a decision if a post-blast area is safe or not, and to provide information about the state of the muck pile produced by the blast. To do so, the framework focuses on three types of sensor data: visual camera data to confirm the state of the hanging walls (ceiling) and mining face at the blast site, gas concentration data provided by an onboard gas sensing device that should be synced with onboard localization systems, and finally 3D LiDAR point-cloud reconstruction data that provides the mining operators with an updated 3D point-cloud map view of the area. Especially for the visual and gas data, it is critical that the UAV visits specific viewpoints that an operator knows to be of interest for both gases and potential structural problems. The idea is that the UAV can be located outside of the hazardous area at a point with communication (ex. Wi-Fi). They can then receive inspection waypoints of areas of interest

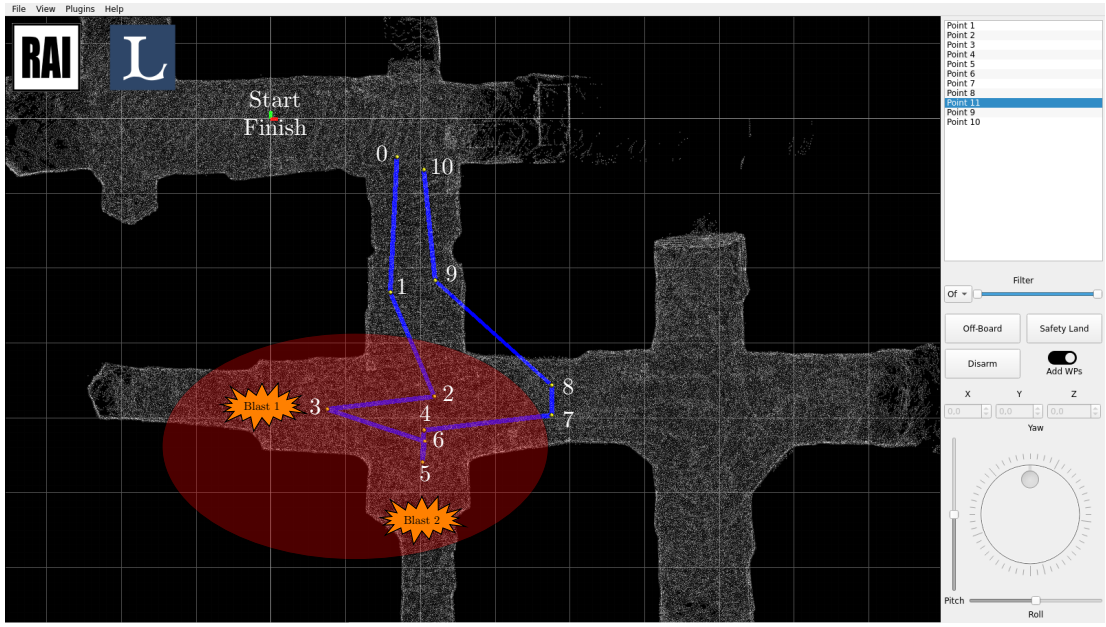

Figure 1: The GUI used by the operator to select a waypoint in the map and create the desired route for the inspection mission. In this figure are also the blast sights marked. The area around the blast sights are dangerous for humans to enter due to the risk of collapse. The areas marked in red were during the experiment no-go zones because of the collapse risk.

regarding the current blasting location from an operator that could be located kilometers away or on the surface. It should then execute the mission without any more communication, while ending the mission back where it started to transfer the data.

3.2 Mission Definition and Operator Interface

Towards these goals, more exploration-oriented navigation stacks such as coverage planning or source localization are in general of less interest, both due to the size of the area and the need to inspect specific locations to ensure safety. We simply want to reliably execute the measurement mission at a series of desired high-level waypoints, or viewpoints, which an operator selects before the mission is initiated. The task then becomes as follows: 1) An operator defines the desired inspection waypoints and sends them to the UAV, 2) relocalize on a global known map of the area, 3) Plan risk-aware and safe paths to the next waypoint in the order they were provided, 4) Execute the path while re-planning at regular intervals, 5) Repeat step 3 when you have arrived at the waypoint pose reference until all areas have been visited, 6) Return to base. As such, the presented RIA stack is developed specifically for this type of mission. The first step in that process is the ability for an operator to select desired waypoints. As part of the RIA stack, we developed an easy-to-use GUI (Graphical User Interface) that enables a non-expert to select the areas of interest to inspect by clicking the desired inspection points in a map. The GUI and the visualized blasting area and related mission which was executed as part of Section 6 is shown in Figure 1, where we have also marked the blasting zones. To select waypoints, denoted as WP, the operator clicks in the GUI that has a 3D point cloud map $\mathbf{M}_{pcd}$ of the area. These are saved and can be sent to the UAV as it is initialized. Once the operator is satisfied with the mission specification, the mission is defined by visiting the list of waypoints $\overline{WP}$ using the D^*_+ path planner. As a note, the blue lines in Figure 1 are not the planned obstacle-free paths, it is simply indicating the order of the clicked waypoints. The result is a generalized and easy-to-deploy framework suitable for industry applications and field conditions for autonomous inspection purposes.

Figure 2: Extreme levels of salt dust during mission execution.

4 Drone and Payload

The UAV platform used for field demonstrations was custom-built at Luleå University of Technology, Robotics & AI Group, and can be seen in Figure 3, with the most important components marked. Most of the autonomy components of RIA rely on 3D LiDAR sensor information as we must navigate in 3D completely dark environments, and for 3D LiDAR based autonomy a critical question in the salt mine environment is that of dust. Throughout the development process we trialed other LiDAR systems but settled on the Ouster OS1 32-beam LiDAR due to the low weight, and importantly that the Ouster could penetrate the salt dust particle clouds kicked up by the propellers enough to enable 3D LiDAR localization and collision avoidance, although far from perfectly. The one thing that the LiDAR system could not manage in the extreme levels of dust was real-time occupancy mapping, as dust particle hits would cause voxels around the UAV to constantly flicker from occupied to free despite tuning and filtering efforts. This was a major restriction and will be discussed more later in section 5.1. A snapshot image from the mission that shows the levels of dust in the post-blast site can be seen in Figure 2.

The UAV also carries a Pixhawk Cube FCU (Flight Control Unit) that handles attitude control and motor mixing, with its internal IMUs (Intertial Measurement Unit) also providing acceleration and roll/pitch data to the Lidar-Intertial localization framework. All computation is handled by an onboard Inter NUC. Not shown in the Figure is a GoPro camera that was added to the front of the UAV to provide visual imagery, selected simply for its internal storage, low weight, and much higher image quality than most small-scale ROS-integrated cameras.

4.1 Gas sensing

Measuring gases requires specific instruments, and in this demonstration we are relying on a Dräger X-am 5600 gas sensor to specifically measure concentrations of CO (Carbon Monoxide), in PPM (Parts Per Million). CO gas is generated by the explosives used for blasting, and is one of the critical gases to monitor, as it is potentially lethal (Ernst and Zibrak, 1998). The Dräger sensor was mounted on the bottom of the platform as shown in Figure 3. Importantly, we had access to the real-time data-stream of gas concentration measurements through a ROS-link with the sensor. This enables the concept of pairing the gas measurements with the real-time onboard localization system such that the measured concentrations can be associated with a position coordinate in the mine.

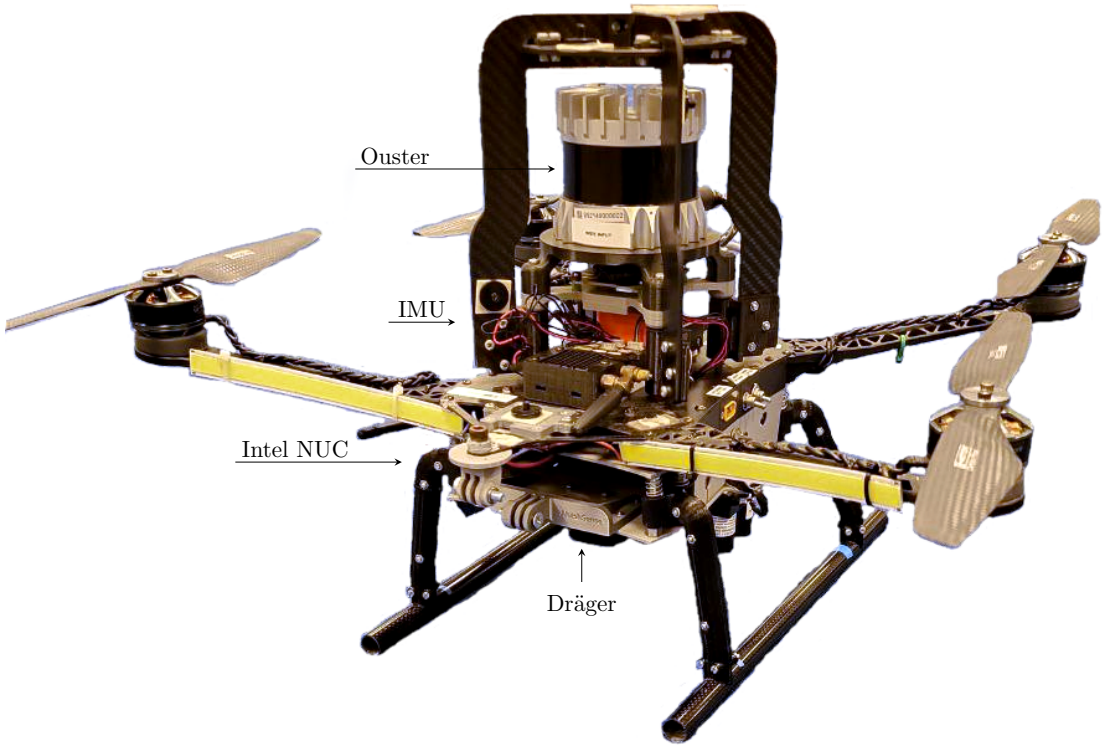

Figure 3: The UAV used in the experiments. In the experiment, a GoPro was mounted in the front of the UAV to capture onboard footage. The Dräger X-am 5600 was mounted below the UAV.

Let us denote the real-time CO measurements as $C_{(t)}$ and together with the estimated world-frame position state of the robot $\hat{p}$ we can generate $C_{(t),\hat{p}}$, or more directly let us denote the time-and-position stamped vector of all such measurements simply as $C_{\hat{p}}$. Those measurements can now be visualized inside the point cloud map $\mathbf{M}_{pcd}$ of the local mining area, and even live-streamed back through ROS in real time to the operator if the wireless communication in the mine allows for it. In the following results visualization in Section 6, $C_{\hat{p}}$ is visualized in a color-coded way so that an operator can see if the area contains any dangerous gases at a glance.

5 Routine Inspection Autonomy (RIA)

Executing a complete mission using only onboard sensing and computation requires a complete autonomy stack. The following subsections will describe the components of the RIA framework in more detail while Figure 4 shows a high level description of how the different parts are connected. In Figure 4, we are assuming that there exists a point cloud map $\mathbf{M}_{pcd}$ and an occupancy map $\mathbf{M}_{bt}$ of the area which in the figure are denoted by their file extensions ".pcd" and ".bt" respectively. To summarize the autonomy described in the following subsections: On the global point cloud map $\mathbf{M}_{pcd}$ we apply the 3DEG (Stathoulopoulos et al., 2024a) relocalization framework using the LiDAR point cloud scans Π_{pcd} to generate an initial guess for the UAV position in the map as x_{init}. The pose x_{init} initializes the FAST-LIO (Xu and Zhang, 2021) LiDAR-Intertial Odometry in the global map frame. The GUI operator sends the inspection waypoint list $\overline{WP}$ to a mission planner, which queries the D^*_+ (Karlsson et al., 2022) path planner to plan a risk-aware path P to the next inspection point WP_i and P is segmented into pose references as $P = [P_1, P_2, \dots P_n]$ (heading and position references). P_i is provided as the attractive force to an Artificial Potential Field (APF) (Lindqvist et al., 2022a), that acts as an inner safety layer and whose repulsive forces are generated from instantaneous raw LiDAR point clouds Π_{pcd}. Finally the APF output is the obstacle-free body-frame state references x_{ref} provided to a reference tracking NMPC (Lindqvist et al., 2020), that in turn generates attitude commands to the onboard FCU.

5.1 The D^*_+ Path Planner

The path planned by D^*_+ is a risk-aware path planner that considers unknown areas and areas close to obstacles, such as walls, as a risk. The risks are considered as a traversal cost c in addition to the standard distance costs/heuristics. D^*_+ is planning the shortest safe path P from the current location $\hat{x}$, according to the sum of traversal costs in the path $\sum \forall c \in P$, to the waypoint WP_i. The gridsearch method is based on the information contained in the occupancy map $\mathbf{M}_{bt}$ provided, in this case, by the OctoMap framework Hornung et al. (2013a). Because $\mathbf{M}_{bt}$ is a voxel map, D^*_+ planning P as the sequence of voxels to visit in order to move from $\hat{x}$ to WP_i in the most optimal (short and safe) manner. The traversal cost c for each voxel is calculated as follows:

$$c_r = \begin{cases} \frac{c_u}{d+1} & \text{if } d < r \\ 0 & \text{else} \end{cases} \tag{1}$$

$$c = \begin{cases} c_r + c_d & \text{if the voxel is free} \\ c_u + c_r + c_d & \text{if the voxel status is unknown} \end{cases} \tag{2}$$

$$\tag{3}$$

where d is the distance from the voxel to the closest occupied voxel and r is the radius in which risk is considered. c_u is the designed cost for unknown voxels, and c_d is the cost for the travel distance. The risk cost is noted as c_r and is a part of the total cost c to traverse through that voxel. In figure 5 the reader can an example of a path planned by D^*_+, where the resulting path P is prioritizing the shortest path in the free area when possible and taking the shortest path through the low-risk green area when there is no free area to pass through. Although the D^*_+ planner is compatible with a continuously updated occupancy map, the extreme levels of salt dust observed in Figure 2 forced us to rely on an offline mapping approach where the occupancy map remained static during runtime. The reason being that the salt dust kicked up by the propellers would at random instants cause voxels to flicker from free to occupied, causing large problems for the path planner. As such an extra reactive safety layer described in section 5.3 was added to ensure robot safety in case of dynamic or changing areas in the map. This restriction points towards further research into real-time applicable and computationally effective dust filters for 3D LiDAR data that would re-enable live occupancy mapping.

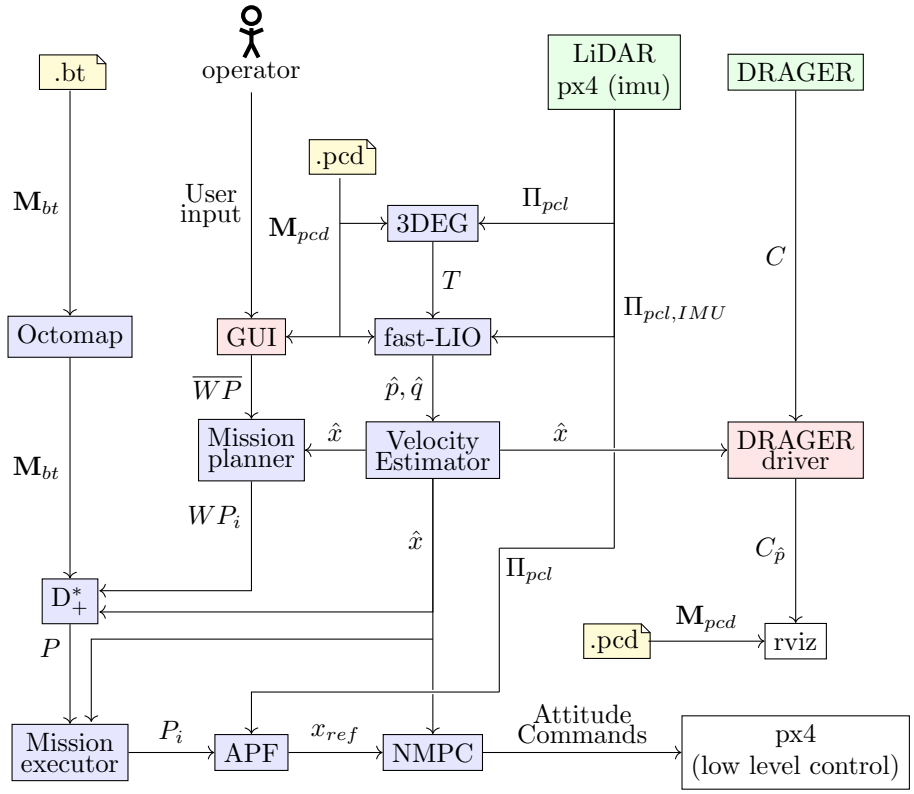

Figure 4: A block diagram of the full system used for the gas inspection mission. The blue boxes are the software that constitutes the autonomy stack. The Red boxes are other software that is not vital for the autonomy but for the mission setup and point of the mission. Marked with green are the sensors that are used and the maps that are loaded from files are marked in yellow. The white components are output related, where rviz is the visualization software and px4 is the onboard low-level controller that is directly controlling the drone.

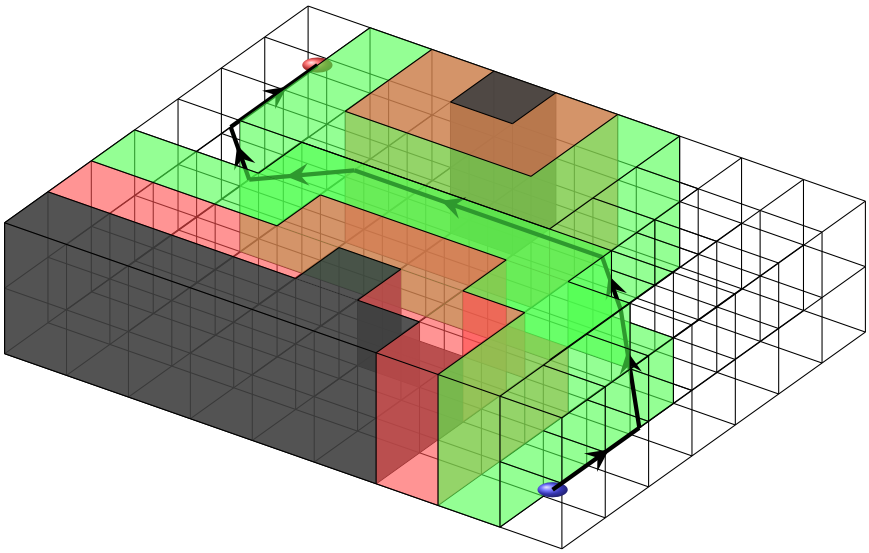

Figure 5: An illustration of how D^*_+ plans a path from the blue ball to the red ball with regards to the risk layer, red voxels are higher risk, green voxels are lower risk, and black voxels are occupied.

5.2 Relocalization & 3DEG

For the robot to be able to follow a path to the given waypoints, it needs to be able to localize itself within a given map. As discussed previously in Section 2 the localization aspect is handled by FAST-LIO (Xu and Zhang, 2021), providing accurate odometry at 10Hz. In order for the odometry to be with respect to the given map's coordinate frame, which we will refer to as the world frame $\mathcal{W}$, there is an initialization process that transforms the local robot frame $\mathcal{R}$ to the world coordinate frame $\mathcal{W}$ through a homogeneous rigid transform of the special Euclidean group, $T : \mathcal{R} \rightarrow \mathcal{W}$. The FAST-LIO package requires an initial guess in order to perform the two-step ICP registration and compute the refined transform T before starting the odometry estimation. The initial estimate can be provided manually by the operator in cases where each inspection mission begins from a unique starting point, or it can be set as a fixed starting position. The latter option imposes undesired constraints, such as a predetermined initial location, which diminishes the autonomy of the robotic system. To address this challenge, we employ our previously developed framework, 3DEG Stathoulopoulos et al. (2024a). This framework, when provided with a map, is capable of estimating a transformation matrix $T \in SE(3)$ that aligns the current robot frame, denoted as $\mathcal{R}$, with the global map frame $\mathcal{W}$. Given a point cloud map $M \in \mathbb{R}^3$ and its corresponding trajectory $Tr \in \mathbb{R}^3$, 3DEG can determine the rigid transformation matrix T using just a single LiDAR scan from the current moment. The process begins by partitioning the map and deriving descriptive vectors, denoted as $Q = \{\vec{q}_1, \vec{q}_2, \ldots, \vec{q}_n : \vec{q}_i \in \mathbb{R}^{64}\}$ and $W = \{\vec{w}_1, \vec{w}_2, \ldots, \vec{w}_n : \vec{w}_i \in \mathbb{R}^{64}\}$, where $\vec{q}$ vectors encode location-specific information, while the $\vec{w}$ vectors encode orientation-specific data, enabling the estimation of the yaw difference between two point cloud scans. Subsequently, a sample from the LiDAR is used to extract the current descriptor, $\vec{q}_t$. Immediately following this, $\vec{q}_t$ is used to query the map's vectors, identifying the nearest neighbour through the following process, denoted as:

$$i = \arg, \min_{i \in \mathbb{N}} f(q_i, q_t) \text{ where } q_i \in Q \tag{4}$$

Once the index of the nearest neighbour, denoted as i, is determined, we can employ the vectors $\vec{q}_i$ and $\vec{w}_i$ to construct the initial guess transform, which is subsequently passed on to FAST-LIO. The estimate obtained in this manner is utilized by FAST-LIO to perform the two-step ICP algorithm, which refines the alignment between the robot frame and the current map, ensuring precise localization.

5.3 Control and Safety

The local navigation stack of RIA is based on the framework presented in Lindqvist et al. (2022a), where an Artificial Potential Field (APF) that operates directly on raw 3D LiDAR point clouds is proposed and combined with a Nonlinear Model Predictive Controller Lindqvist et al. (2020) (NMPC) to perform position set-point tracking as is required when following inspection trajectories. The idea is that while the D^*_+ trajectories are risk-aware, the artificial potential field can provide a backup safety layer by generating repulsive forces from raw LiDAR scans, and as such it is resistant to poor occupancy mapping and can save the platform in instances of momentarily incorrect path planning or state-estimation jumps when near a wall or obstacle. For completeness, the following section will shortly describe the APF used in the mission demonstration. The force field formulation follows closely the foundational works on APFs (Park et al., 2001) with the difference of not considering discrete obstacle surfaces and instead applying weak repulsive forces for each LiDAR hit inside the radius of influence of the repulsion r_F and summing them to get the total. Let us denote the 3D point cloud input to the algorithm as $\{\Pi_{pcl}\}$ consisting of points relative to the LiDAR frame as $\rho = [\rho_x, \rho_y, \rho_z]^T$. The subset of such points inside r_F is $\boldsymbol{\rho}_F \in \{\Pi_{pcl}\}$ where $\| \rho_F^i \| \leq r_F$ and $i = 0, 1, \ldots, N_{\rho_F}$ are the points that we want to use to generate a repulsive force. Let us then directly also consider an inner safety-critical radius r_c, with related points $\boldsymbol{\rho}_c \in \{\Pi_{pcl}\}$ that are N_{ρ_c} in number. To repulse the UAV away from any nearby LiDAR hits, the repulsive force field function becomes

$$F^r = \sum_{i=1}^{N_{\rho_F}} L^r (1 - \frac{\| \rho_F^i \|}{r_F})^2 \frac{-\rho_F^i}{\| \rho_F^i \|} + \sum_{i=1}^{N_{\rho_c}} L^c \frac{-\rho_c^i}{\| \rho_c^i \|}. \tag{5}$$

Here L^r and L^c describe repulsive gains. We can see that the "outer" first term has a radius-dependant repulsive force, where the force equation becomes zero for $\| \rho_F^i \| = r_F$ increasing up to L^r at $\| \rho_F^i \| = 0$. This generates a smooth repulsion as an obstacle enters the radius of influence as to not generate too aggressive maneuvering, while the inner safety radius imposes a static force-per-point L^r as to directly repulse the UAV away from walls or obstacles entering the safety-critical radius. Due to the huge amounts of dust we restricted $\boldsymbol{\rho}_c$ by also enforcing that

$N_{\rho_c} \geq N_{c,min}$ or the safety-critical force is set to zero, where $N_{c,min}$ is a cut-off number to filter out dust hits inside the critical safety radius to differentiate between dust particles and obstacles. Assuming the goal is to reach the next point $P_i = [wp_x, wp_y, wp_z]^T$ in P generated by D_+^* we define the attractive force $F^a = [F_x^a, F_y^a, F_z^a]^T$ such that $F^a = P_i^{\mathcal{B}} - \hat{p}^{\mathcal{B}}$, with $P_i^{\mathcal{B}}$ and $\hat{p}^{\mathcal{B}}$ denoting a yaw-compensated coordinate frame measurement of the robot position $\hat{p}$ and the next goal point P_i. From an intuitive point of view this can be seen as the attractive force F^a being the vector from the current position to the next given way-point with a unitary gain, while the repulsive force F^r is the shift in the next way-point required to avoid obstacles. The preliminary work Lindqvist et al. (2022a) also describes a force normalization and saturation, as well as an adaptive weights scheme, which is also applied in this work. Summing the resulting repulsion and attraction, we form the state reference x_{ref} sent to the set-point reference tracking controller, which in the RIA framework is a high-performance NMPC Lindqvist et al. (2020) implemented in the Optimization Engine Sopasakis et al. (2020).

5.4 Change detection

An important aspect of the autonomous inspection routine is the ability to fully automate the detection of potential over time changes in the environment. Assuming regular inspection missions in the same area, at regular intervals, the reconstructed data can be utilized to detect and track changes over time. As an example, it could be broken or misplaced equipment, rock falls, deformation or drifting of the tunnels, and so on. In order to avoid the limitations mentioned in Section 2, we utilize our own algorithm (Stathoulopoulos et al., 2023b) that works directly on the point cloud data and provides a fast and scalable solution. In a 3D space $\mathbb{R}^3$, a robot generates a point cloud map M_t in iteration t with respect to its local coordinate frame $\mathcal{W}^t$, along with the robot's trajectory defined as Tr_t, containing position points relative to $\mathcal{W}^t$. The problem can be formulated as, finding the pair of regions in M_t and M_{t+1} with the most change, defined as:

$$\max_{(i,j)\,\in\,\mathbb{N}} \bar{d}(S_{k_i}, S_{k_j}), \tag{6}$$

Here, S is a spherical subset of a map M, and $\bar{d}(X, Y)$ measures the difference between regions. The goal is to extract objects of interest O_n that either weren't present in the first map or have been displaced. The developed framework first aligns the maps using a homogeneous rigid transformation of the special Euclidean group $T : \mathbb{R}^3 \to \mathbb{R}^3$, derived from our previous work on point cloud map merging (Stathoulopoulos et al., 2023a). The next step is to identify the changed regions, via the 3DEG (Stathoulopoulos et al., 2024a) place recognition descriptors. The descriptors denoted as $\vec{q} \in \mathbb{R}^{64}$ are multidimensional vectors that encode orientation-invariant and place-dependent information and can be used to query similar places. In this scenario we have reverted the problem of place recognition and instead of querying descriptors to find the pair with the minimum distance, the regions with the most change are determined by finding the pair of vectors $\vec{q}_{t,i}$ and $\vec{q}_{t+1,j}$ with maximum distance in the descriptors' space, a process denoted as:

$$(k_i, k_j) = \underset{(i,j)\,\in\,\mathbb{N}}{\arg\max}\, f(Q_{t,i}, Q_{t+1,j}) \tag{7}$$

Since the robot r performs routine inspection missions, the trajectories Tr_t and Tr_{t+1} are similar and therefore the corresponding descriptor pairs are similar, except the ones that contain severe changes. Then, in order to extract the objects of interest $O_n = S_{k_j} \backslash S_{k_i} = \{x \in S_{k_j} : x \notin S_{k_i}\}$, we seek for the points in S_{k_j} that are not in S_{k_i}. To optimize this process, the reference spherical regions S_{k_i} are voxelized in an occupancy map V_{k_i}, and then a point-to-voxel comparison follows to extract the objects of interest. This process is described as:

$$O_n = \{x \in S_{kj} \mid \forall\ v \in V_{k_i},\ x \notin v\}, \tag{8}$$

Finally, a statistical outlier removal filter is applied to O_n to eliminate noise or misclassified points, where outliers are determined based on a threshold λ from the mean μ and standard deviation σ.

6 Results

This section will detail the subsequent field demonstration evaluations of the RIA framework in the context of the considered use-case of safety inspection and gas monitoring. The experimental evaluation starts with initial smaller scale missions, still in relevant mining environments, to test out the aerial gas monitoring system and the developed autonomy stack. Next, the RIA stack was evaluated in a large-scale mission scenario for visual inspection with the goal of testing the autonomy and sensing stack on a large scale and in a very harsh environment. Finally, we

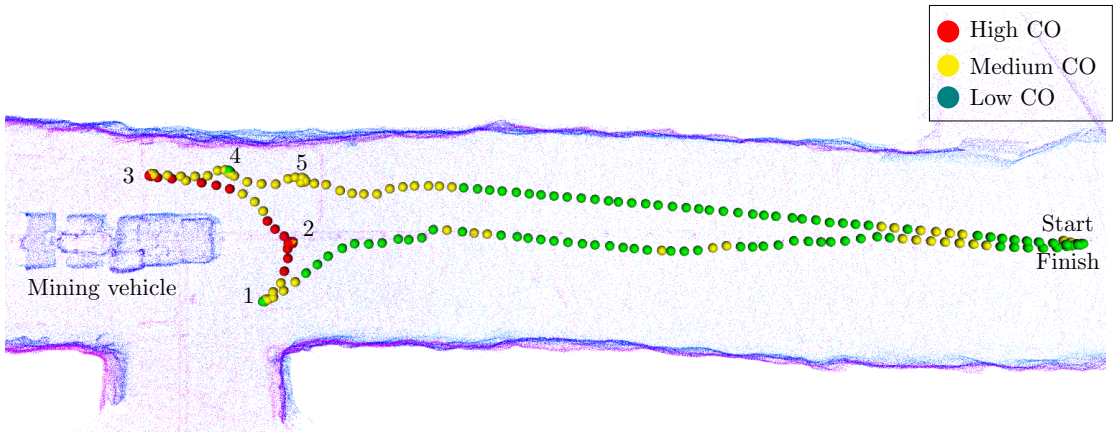

Figure 6: The detected gas levels plotted in the map, with low levels as green dots, medium yellow, and high as red. In this test where the low limit 1.3 and the high 2.5.

demonstrate the gas monitoring ability in the direct use-case scenario of a large-scale mission executed shortly after a real blasting operation to inspect for lingering gases, changes in the area caused by the blast (muck pile and rockfalls), and visual safety inspection for structural integrity. The best way to view the final demonstration result is through the experiment video, provided at `https://www.youtube.com/watch?v=jzg-KsWKPp8&t=3s`, showing the real-time behavior of the robot performing fully autonomous inspection in a potash/salt mine environment. For clarity, we should note that the initial maps used to set up the mission $\mathbf{M}_{pcd}$ and $\mathbf{M}_{bt}$ were built relying on the SLAM (Simultaneous Localizaion and Mapping) framework LIO-SAM (Shan et al., 2020) to build a dense point cloud map which can later be used for relocalization, and using OctoMap Hornung et al. (2013a) to build a 3D voxel occupancy map of the area used for path planning. These maps can then be updated from sensor data collected with each consecutive routine inspection mission in the area.

6.1 Initial gas monitoring evaluation

The first evaluation was performed in order to evaluate the aerial gas sensing capability in a small-scale mission. This mission was performed at the Epiroc Test Mine near Örebro, Sweden. A large mining vehicle was left running for a few minutes to produce CO gas, and then the RIA-enabled UAV executed a waypoint inspection mission for several waypoints around the truck, which can be seen in Figure 6. In this Figure, the gas values are color coded with arbitrary cut-off values just to visualize where higher concentrations were detected. As anticipated, elevated gas levels were correctly observed in close proximity to the machinery's exhaust. Additionally, heightened concentrations were detected near a ventilation exhaust along the path which can be seen in the Figure. This small-scale test demonstrated the ability to accurately measure the position of high gas concentrations with an aerial platform through the developed autonomy and sensing stack, and that gas pocket concentrations are detectable despite the air disturbance from the propellers and located at reasonable locations (e.g. near the vehicle exhaust).

6.2 Gas monitoring after test blast

Next, we evaluate the RIA stack for gas monitoring in a smaller mission but generating CO gas from an explosives source, mimicking a real post-blast inspection scenario. This test was enabled by mining company LKAB in northern Sweden, in the largest underground iron mine in the world. A modest charge (for mining standards) of $8 - 10\,\text{kg}$ of explosives was strategically placed in a mining drift to generate detectable gases. Here, the mission set-up was simple: navigate approximately 50 meters into a mining drift where the blast happened and return. Starting the mission, the first step is the relocalization process, which can be seen in Figure 7, as an example of the 3DEG relocalization process, aligning the initial localization frame with the pre-built map frame. Although the starting position is nearby the position of the UAV when it built, the alignment is crucial for proper localization within the map. The whole process of the relocalization takes under 0.5 seconds. The mission outcome can be seen in Figure 10, where we can see that the UAV detected a gas pocket inside the drift where the blast happened. The additional waypoint shown in Figure 10 was added to also test the risk-aware navigation capabilities - which can be observed by the generated path

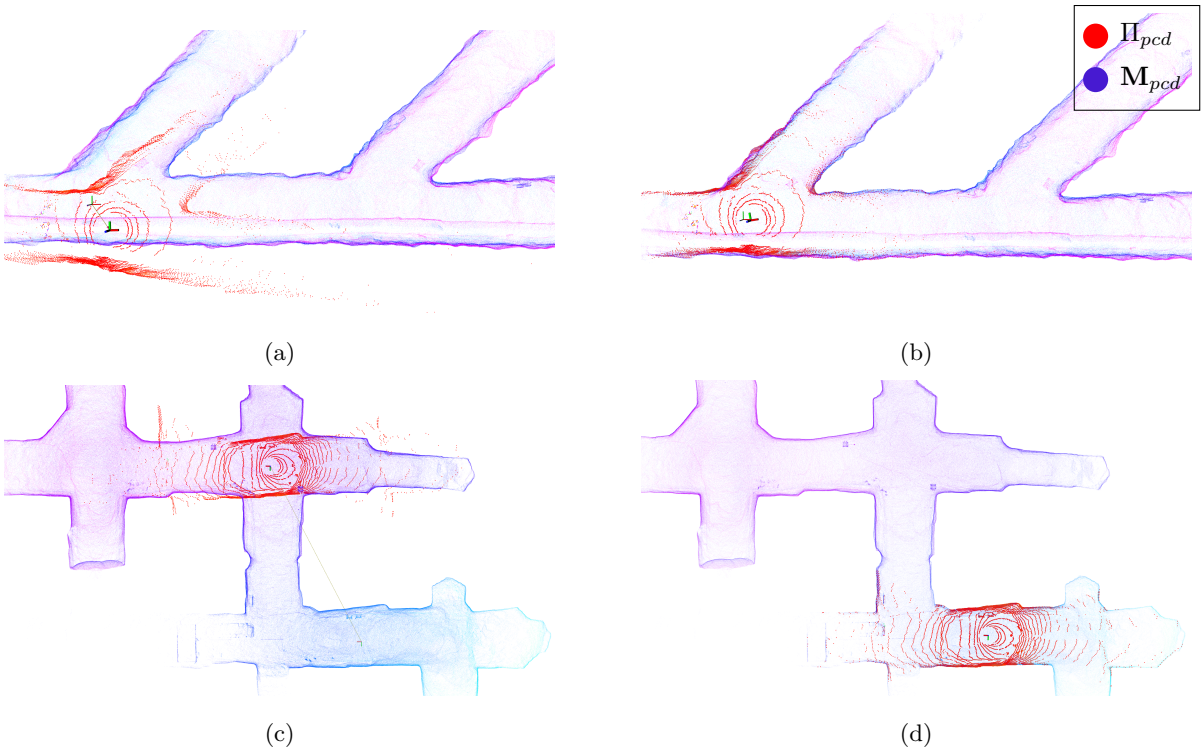

Figure 7: In (a) and (c) are the raw scans shown on the map before the relocalization took place. And in (b) and (d) are the raw scans shown properly aligned to the map after the relocalization. (a) and (b) are from the gas test in LKAB while (c) and (d) are from the full scale gas mission.

not cutting corners and keeping a safe distance to the tunnel walls at all times. A snapshot of mission execution can be seen in Figure 9, showing the UAV also operating in the lingering smoke left from the blast and in the complete darkness of the mining environment. The figure shows the UAV entering the mining drift where the test blast took place approximately 40 minutes after the blast, but while there was still a presence of smoke and detectable elevated levels of CO gas as shown in Figure 8.

6.3 Large-scale autonomy test of RIA

To assess the navigation capabilities of RIA, a comprehensive large-scale navigation test was conducted at a K+S Group potash mine in Germany. The test involved strategically placing 9 WP's, as illustrated in figure 11, throughout the mine's environment with the underlying mission purpose of visual inspection of the mining drifts. The test was performed at the same site as the main demonstrator for gas inspections described in the following section 6.4 to mimic the environment conditions. The objective was to extend the drone's coverage, challenging the planning capabilities of D^*_+ by incorporating waypoints that were deliberately positioned out of the line of sight and inside the mining drifts. The Figure 11 show safe and efficient navigation in a large-scale mining area with a mission length of around 280 m, a significant mission length as compared to previous works deployed in the field, while dealing with significant levels of salt dust.

During the test, the "mission executor" set up to maintain a specified heading while hovering at the waypoints for a duration of 10 s. This intentional pause allowed for a visual inspection of the drift. The resulting images from this test can be observed in figure 12, showcasing dark but useful images facilitated by the drone's stable hovering, showing the visual safety inspection of the ceiling (hanging wall) and tunnel walls. Operators at the site denoted the imagery as useful for determining the safety of the drift.

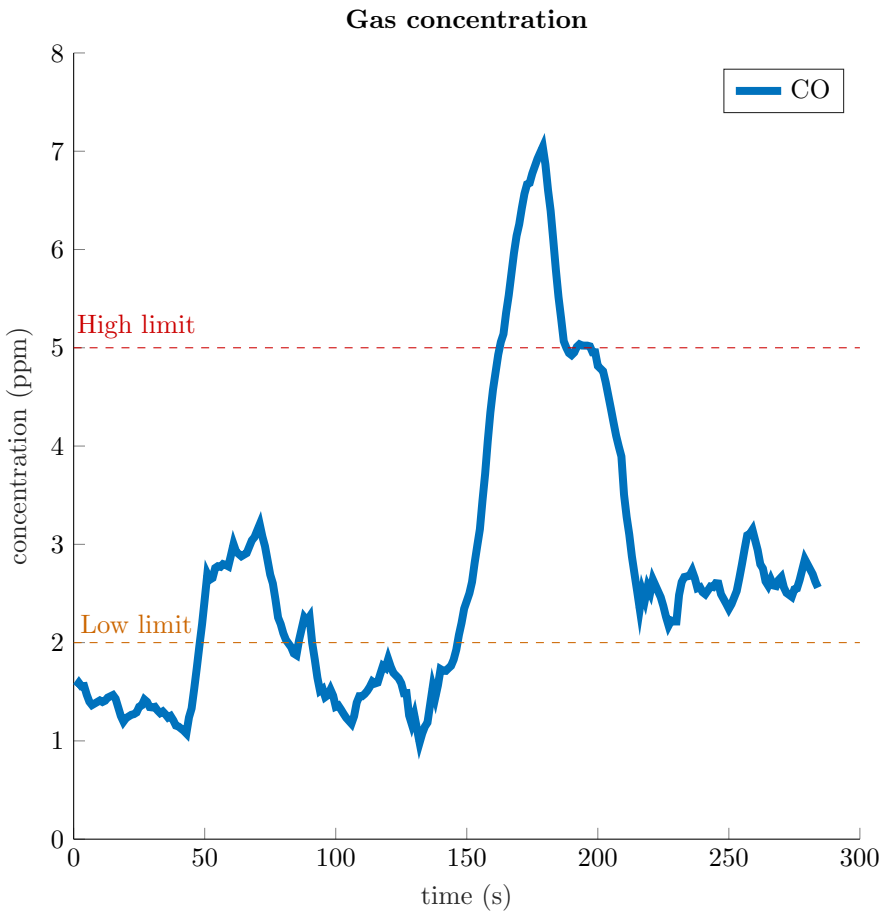

Figure 8: Gas concentration (CO) in parts per million (ppm) measurements over time during the test mission in LKAB.

Figure 9: Snap shot from the gas monitoring after test blast mission. Even after 40 min the smoke levels are high.

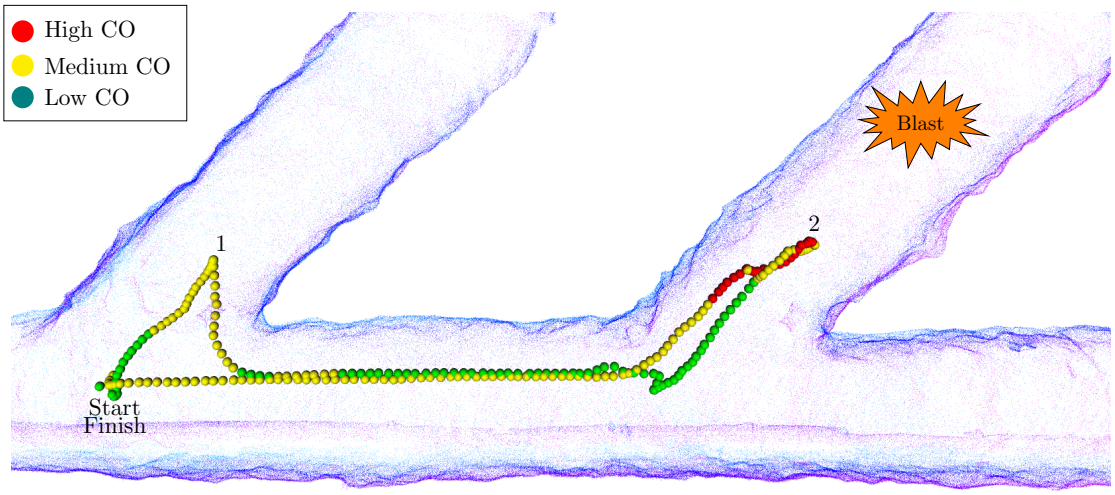

Figure 10: The detected gas levels plotted in the map, with low levels as green dots, medium yellow, and high as red. In this test the low limit 2.0 and the high 5.0 ppm.

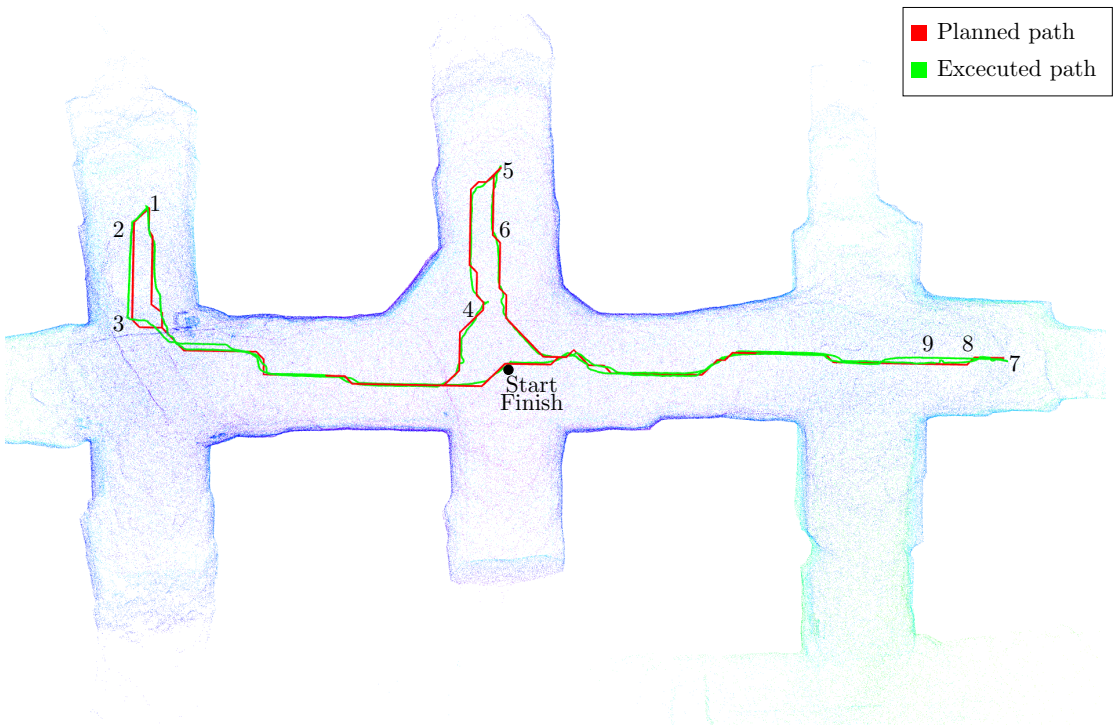

Figure 11: Test mission of the RIA framework. 9 WP (10 with the return WP)designed to test D_+^* and provide visual inspection imagery for a operator. The green path is the one the drone actually took while the red is the planned path.

6.4 Field Demonstration after real large-scale blasting operation

6.4.1 Mission Set-up

As previously mentioned, the demonstration took place at a K+S Group-operated potash mine in Germany, at a depth of 700 m. Two drifts, as indicated in Figure 1, were loaded with explosives as per standard blasting procedures. The crew was stationed in a safe zone some kilometers away during the blast, and arrived at the site together with the mine workers who usually perform manual safety inspections around 40 min after the blast. The UAV mission was deployed from a zone declared to be safe, to enter into the area with blasted mining drifts.

The inspection waypoints, shown in Figure 1 (and later with the inspection route in Figure 13) were selected by the mine operators to provide them with interesting and useful data points regarding ventilation, gas pockets, and visual images at specific locations. E.g. the mission was deployed as if it was performing a real routine inspection mission after blasting, not just as an experiment to demonstrate our autonomy. The exact same mission was also performed prior to the blast to collect data used for the change detection framework.

6.4.2 Mission overview

The best way for the reader to see the field demonstration experiment is through the following video link: `https://www.youtube.com/watch?v=jzg-KsWKPp8&`. In the video, the reader can experience the harsh mine environment and see the UAV during operations, while we show the real-time data from path planning, localization, scanning, and video from the onboard camera.

The inspection route calculated by D_+^* can be seen in Figure 13. Due to the large number of inspection points, the risk-aware navigation does not necessarily shine in this mission, as most viewpoints are in the line of sight of each other. The only risk-aware aspect is around viewpoints 1, 9, and 10 where the UAV was commanded to inspect a

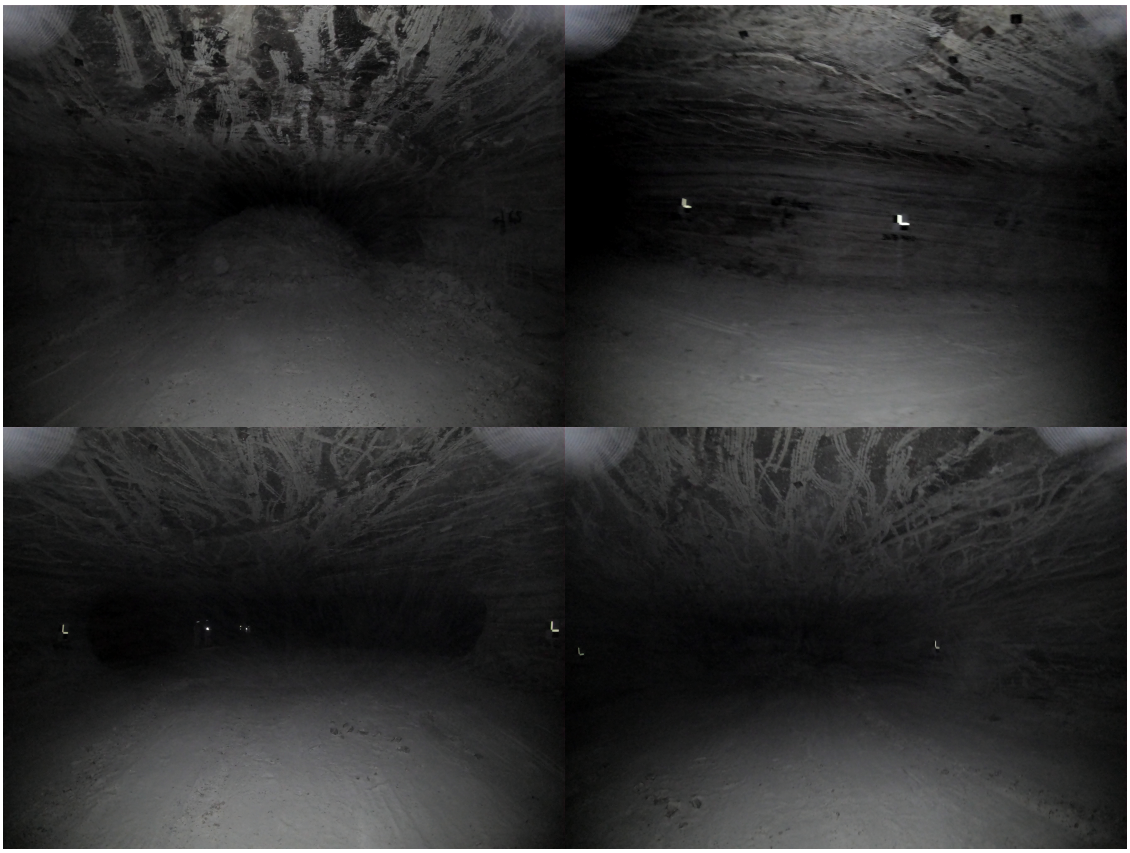

Figure 12: Images for the RIA module test showing results from the visual inspection. The UAV is inspecting the structural integrity of the ceiling and tunnel walls for any cracks produced by a blast, and the top-left image shows the muck pile indicating a successful blast.

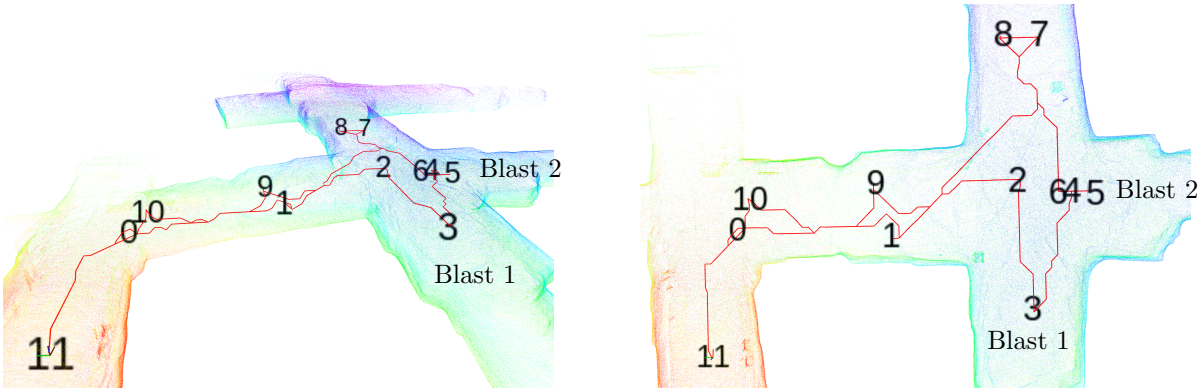

(a) Isometric view

(b) Top view.

Figure 13: The waypoints and planned route for the gas inspection missions.

point close to the walls. Here we can see the risk costs in action, as the UAV spends minimal time close to the walls, and only enters the unsafe area to reach the desired inspection point. In the mission the UAV flew approximately 205 m to visit the 11 operator-provided inspection waypoints. The 3DEG relocalization process is also visualized in Figure 7 for this mission, aligning the localization with the mission start position in the map.

6.4.3 Gas Monitoring

During the mission, the CO gas concentration level was plotted in the point cloud map of the area, shown in the resulting Figure 14, while a plot of raw gas concentrations can be found in Figure 15. In Figure 14, the gas levels are color-coded as green, yellow, and red for low, medium, and high values respectively. It should be noted that the color coding thresholds were set to visualize the detected peak in gas concentrations (which can be seen in Figure 15), and the real dangerous levels are around 50 ppm, and as such the area was free of hazardous levels of gas at the time of the mission. (note: measuring gases right after the blast is pointless as there will always be too high concentrations for workers to enter, as such the mission was performed at the regular time for manual inspection when gases have usually dissipated to the level where workers can enter, but an inspection must first approve the area). The important and critical data that is visualized in Figure 14 is that there were pockets of higher gas concentrations lingering around the second blast zone, which was also the drift without a ventilator. The ability to measure at many locations and associate the gas measurement with a position coordinate in the mine enables a more thorough inspection of lingering gases, and our demonstration shows that such gas pockets can be detected using an autonomous UAV.

6.4.4 Change detection

The change detection process can be considered crucial after the blasting operation, since the safety inspection crew can quickly access the status of the expanded drifts and identify any potential hazards caused by rockfalls. The change detection process begins by aligning the two-point cloud maps from before and after the blasting, as depicted in Figure 16. In more detail, Figure 16a presents the map before the blast, in yellow points, Figure 16b presents the map after the blast in blue points and finally, Figure 16c presents both maps merged into one, resulting in a green color due to the overlapping of the yellow and blue points. The next steps for the change detection algorithm are to identify the changed regions through the place recognition descriptors and then extract the changed points by the point-to-voxel comparison. The results of this process are depicted in Figure 17. In Figure 17a, the reader can see a top view of both blasted drifts, where the walls have been expanded by approximately 10 meters on each side. The walls are highlighted in red color by the change detection algorithm along with the muck pile as well as some of the humans that were observing the experiment. The reader can further see a more detailed view of the expanded drifts in Figures 17b and 17c that correspond to the waypoints WP1, WP3, and WP2, WP5 respectively. Overall, the algorithm needs 0.301 seconds to align the two point cloud maps, then 0.068 seconds to detect the changed areas and finally, 0.213 seconds to extract the objects. As a result, within a second, we are able to automatically identify the expansion of the walls as well as detect the addition of approximately 120 m^3 of new volume from the blast. This information can then be used to plan the extraction process of the ore (or in this case salt) instantly after the

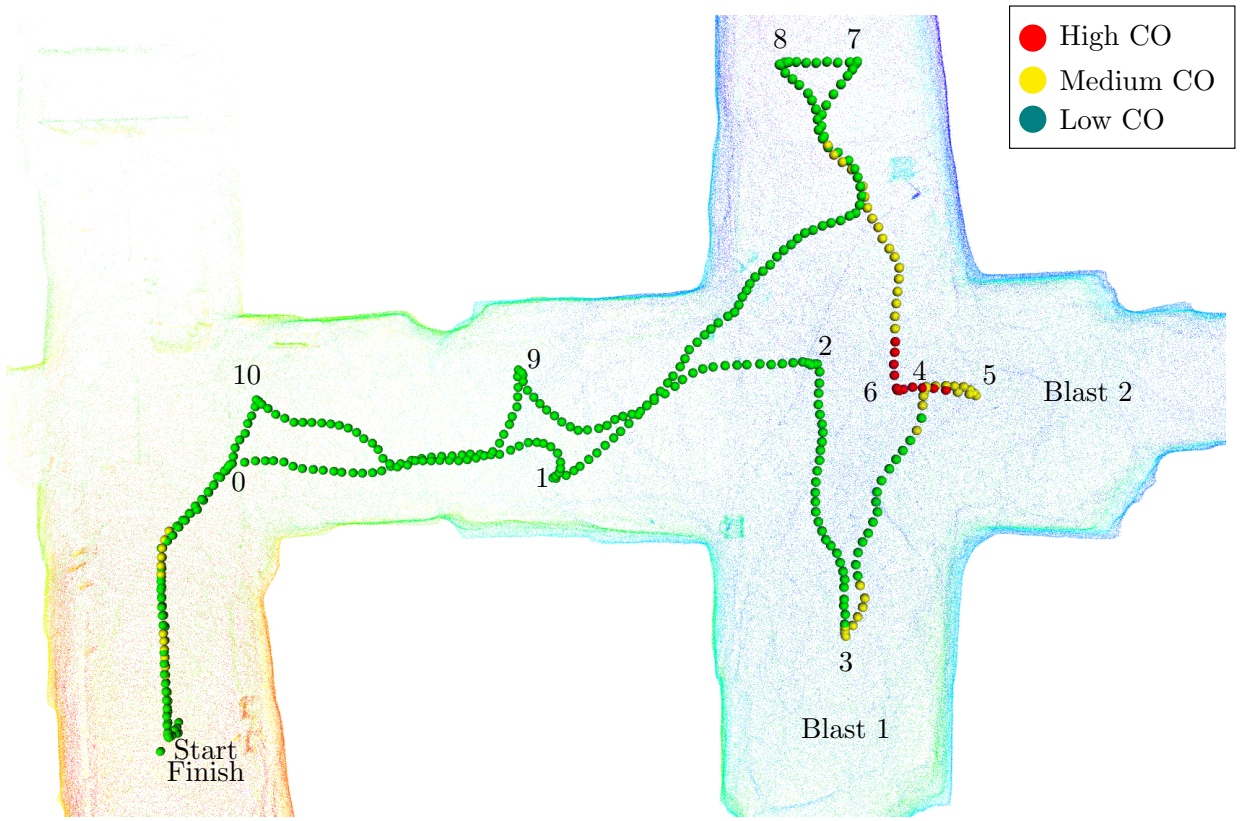

Figure 14: The detected gas levels plotted in the map, with low levels as green dots, medium yellow, and high as red.

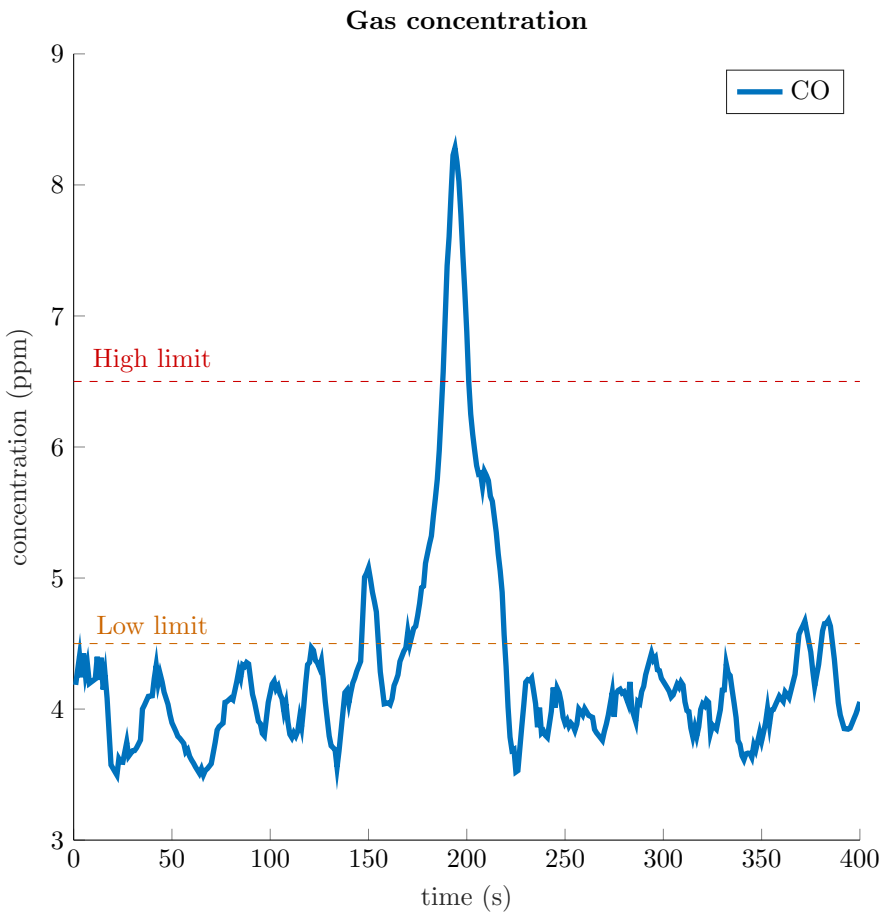

Figure 15: Gas concentration (CO) in parts per million (ppm) measurements over time during the mission.

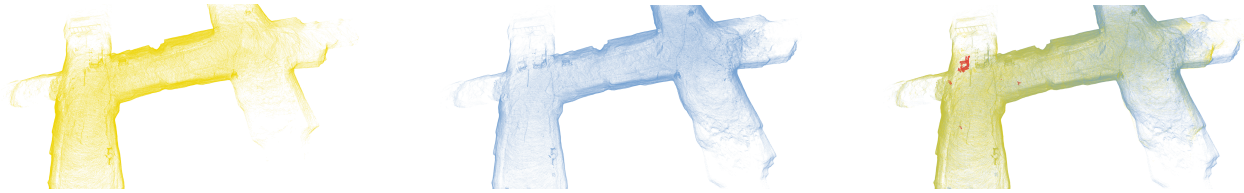

(a) The point cloud registered in the mission before the blast.

(b) The point cloud registered from the mission after the blast.

(c) The two-point clouds merged into one.

Figure 16: The two-point clouds recorded during the missions were visualized as well as merged.

inspection mission is complete.

6.4.5 Visual Inspection

The final inspection aspect comes from onboard visual camera data from the blasting sites. As is clear from the experiment video, the huge amounts of dust combined with the darkness of the mine makes the image stream less usable despite onboard illumination provided by the UAV. But, as was also indicated by mine operators, some snapshots of clear images did provide useful data in regards to if the blast happened correctly and to the state of the muckpile. Figure 18 shows two snapshot images from the mission highlighting the dust levels in the blasted drift but also an instance of a clear image that includes the pile generated by the blasting operation. In general, raw standard RGB image data easily becomes corrupted in these types of environments. This points towards the use of other visual sensing or lighting technologies that could reduce the dust reflectivity and create a clearer image, as dust filtering approaches are more difficult to apply to visual data as opposed to the LiDAR pointclouds.

7 Conclusions

This article has examined the use of autonomous UAVs to perform inspection and gas monitoring tasks in underground mines. To enable the fully autonomous mission the manuscript has described the developed RIA stack, as well as developed data processing algorithms. Through our successful field demonstrations in multiple mines we have shown the potential for autonomous aerial gas monitoring using UAVs. The main experiment showcasing an inspection mission after a real blasting procedure serves as a technology demonstrator of the capabilities of field robots and autonomous UAVs for deployment into harsh environment for increased safety. The development process and field trials also identified directions for future research mainly in terms of dealing with high amounts of dust and smoke. Real-time point cloud filtering of dust particles, as well as sensors like radars that can penetrate the smoke and dust can have a massive impact on improving the navigation capabilities in these types of scenarios - and in our case the straightforward extension of re-enabling the use of real-time occupancy mapping which was a major restriction. In terms of the RIA stack, we aim to extend the framework into the deployment of multiple robots through task assignment Dahlquist et al. (2023), as well as developing methods based on semantic scene recognition Baheti et al. (2020) that would allow the system to on its own identify inspection points of interest after an operator provides the general area to be inspected. Despite these limitations, the developed autonomy and sensing stack combined with the high TRL (around TRL-7) demonstrator, shows that these systems are ready to be integrated into the mining sector.

References

Agha, A., Otsu, K., Morrell, B., Fan, D. D., Thakker, R., Santamaria-Navarro, A., Kim, S.-K., Bouman, A., Lei, X., Edlund, J. A., Ginting, M. F., Ebadi, K., Anderson, M. O., Pailevanian, T., Terry, E., Wolf, M. T., Tagliabue, A., Vaquero, T. S., Palieri, M., Tepsuporn, S., Chang, Y., Kalantari, A., Chavez, F., Lopez, B. T., 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., Lee, C.-A., Gao, J. L., Merewether, G. B., Maldonado-Contreras, J., Salhotra, G., Silva, M. S. D., Ramtoula, B., Fakoorian, S. A., Hatteland, A., Kim, T., Bartlett, T., Stephens, A., Kim, L., Bergh, C. F., Heiden, E., Lew, T., Cauligi, A., Heywood, T., Kramer, A.,

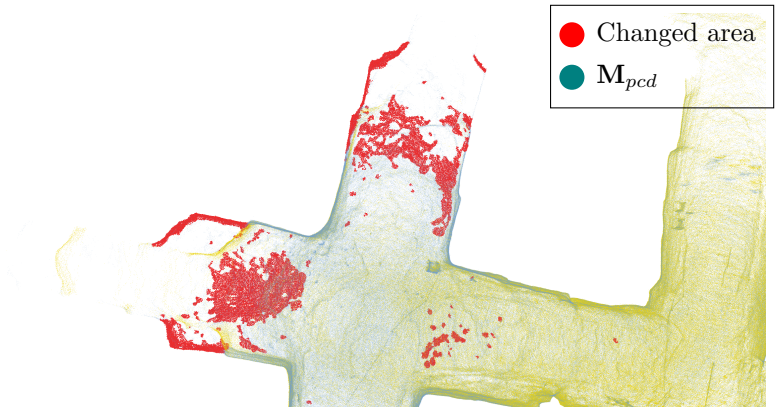

(a) A top view of the two drift where explosions took place, with the changes marked with red dots. In the tunnel to the right, a few dots are visible from the humans who observed the experiment.

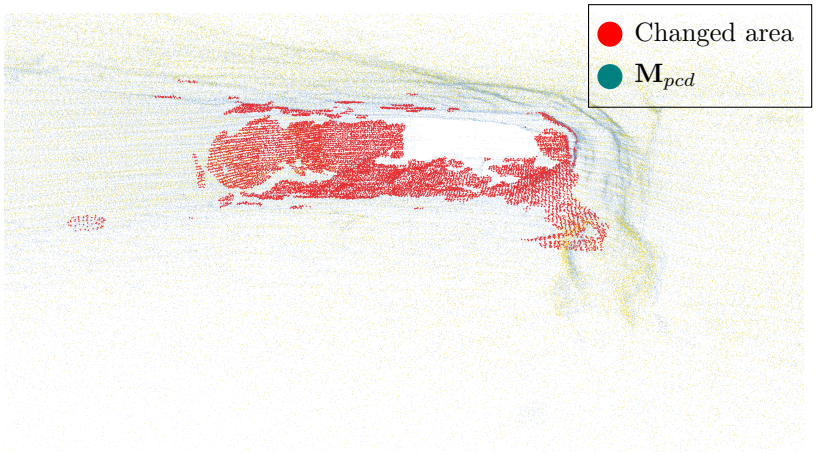

(b) An inside view of the drift with blast 1 and waypoint 3.

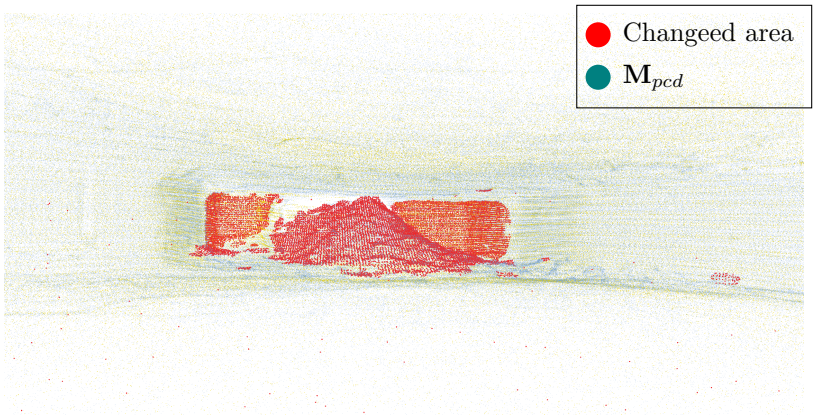

(c) An inside view of the drift with blast 2 and waypoint 5.

Figure 17: A visualization of the differences in registered point clouds between the before and after the blast shown in red.

(a) On instance of imagery from a GoPro mounted on the drone during the experiment where there was no visibility at all

(b) On instance of imagery from a GoPro mounted on the drone during the experiment where the dust clears enough to get some information out of the image.

Figure 18: Snapshots for the drone's point of view during the experiment, showing a good and a bad instance of dust levels.

Leopold, H., Melikyan, H., yong Choi, H., Daftry, S., Toupet, O., Wee, I., Thakur, A., Feras, M., Beltrame, G. A., Nikolakopoulos, G., Shim, D.-Y., Carlone, L., and Burdick, J. W. (2022). NeBula: TEAM CoSTAR's Robotic Autonomy Solution that Won Phase II of DARPA Subterranean Challenge. *Field Robotics*, 2:1432–1506.

Arain, M. A., Hernandez Bennetts, V., Schaffernicht, E., and Lilienthal, A. J. (2021). Sniffing out fugitive methane emissions: autonomous remote gas inspection with a mobile robot. *The International Journal of Robotics Research*, 40(4-5):782–814.

Atapour-Abarghouei, A. and Breckon, T. P. (2018). A comparative review of plausible hole filling strategies in the context of scene depth image completion. *Computers & Graphics*, 72:39–58.

Baheti, B., Innani, S., Gajre, S., and Talbar, S. (2020). Semantic scene segmentation in unstructured environment with modified deeplabv3+. *Pattern Recognition Letters*, 138:223–229.

Bamford, T., Esmaeili, K., and Schoellig, A. P. (2017). Aerial rock fragmentation analysis in low-light condition using uav technology. *arXiv preprint arXiv:1708.06343*.

Basha, P. C. and Patel, M. V. K. (2022). Applications of Robotics in Mining Industry. *Thematics Journal of Applied Sciences*, 11(6):1–6.

Bhatt, A., Maity, A., and Das, K. (2018). Path planning for uav using sniff-dog-algorithm. In *2018 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 692–701. IEEE.

Dahlquist, N., Lindqvist, B., Saradagi, A., and Nikolakopoulos, G. (2023). Reactive multi-agent coordination using auction-based task allocation and behavior trees. *arXiv preprint arXiv:2304.01976*.

Dang, T., Tranzatto, M., Khattak, S., Mascarich, F., Alexis, K., and Hutter, M. (2020). Graph-based subterranean exploration path planning using aerial and legged robots. *Journal of Field Robotics*, 37(8):1363–1388. Wiley Online Library.

DARPA (2020). DARPA Subterranean (SubT) challenge. Accessed: February 2021.

Ernst, A. and Zibrak, J. D. (1998). Carbon monoxide poisoning. *New England journal of medicine*, 339(22):1603–1608.

Fischler, M. A. and Bolles, R. C. (1981). Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Commun. ACM*, 24(6):381–395.

Forooshani, A. E., Bashir, S., Michelson, D. G., and Noghanian, S. (2013). A survey of wireless communications and propagation modeling in underground mines. *IEEE Communications surveys & tutorials*, 15(4):1524–1545.

Freire, G. and Cota, R. (2017). Capture of images in inaccessible areas in an underground mine using an unmanned aerial vehicle. In *UMT 2017: Proceedings of the First International Conference on Underground Mining Technology*. Australian Centre for Geomechanics.

Gojcic, Z., Schmid, L., and Wieser, A. (2021). Dense 3D displacement vector fields for point cloud-based landslide monitoring. *Landslides*, 18(12):3821–3832.

Gomathi, V., Sowmeya, S., Avudaiammal, P., and Ganeshia, R. (2015). Design of an adaptive coal mine rescue robot using wireless sensor networks. *International Journal of Computer Applications*, 975:8887.

Goswami, A., Sharma, D., Mathuku, H., Gangadharan, S. M. P., Yadav, C. S., Sahu, S. K., Pradhan, M. K., Singh, J., and Imran, H. (2022). Change Detection in Remote Sensing Image Data Comparing Algebraic and Machine Learning Methods. *Electronics*, 11(3).

Harabor, D. and Grastien, A. (2014). Improving jump point search. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 24, pages 128–135.

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.

Horizon, N. S. (2020). Next Generation Carbon Neutral Pilots for Smart Intelligent Mining Systems.

Hornung, A., Wurm, K. M., Bennewitz, M., Stachniss, C., and Burgard, W. (2013a). OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous robots*, 34:189–206.

Hornung, A., Wurm, K. M., Bennewitz, M., Stachniss, C., and Burgard, W. (2013b). OctoMap: An Efficient Probabilistic 3D Mapping Framework Based on Octrees. *Autonomous Robots*. Software available at `https://octomap.github.io`.

Jiang, A., Zou, L., Wei, Y., Shen, G., and Jiang, Y. (2022). Complete Coverage Path Planning for Industrial Mobile Robots. In *Proceedings of the 2022 4th International Conference on Robotics, Intelligent Control and Artificial Intelligence*, pages 42–46.

Kanellakis, C., Mansouri, S. S., Georgoulas, G., and Nikolakopoulos, G. (2018). Towards autonomous surveying of underground mine using mavs. In *International Conference on Robotics in Alpe-Adria Danube Region*, pages 173–180. Springer.

Karlsson, S., Koval, A., Kanellakis, C., and Nikolakopoulos, G. (2022). D+*: A risk aware platform agnostic heterogeneous path planner. *Expert systems with applications*, page 119408.

Kim, H. and Choi, Y. (2023). Development of Autonomous Driving Patrol Robot for Improving Underground Mine Safety. *Applied Sciences*, 13(6):3717.

Koenig, S. and Likhachev, M. (2002). D* lite. *Aaai/iaai*, 15.

Koide, K., Miura, J., and Menegatti, E. (2019). A portable three-dimensional lidar-based system for long-term and wide-area people behavior measurement. *International Journal of Advanced Robotic Systems*, 16(2):1729881419841532.

Koval, A., Kanellakis, C., and Nikolakopoulos, G. (2022). Evaluation of Lidar-based 3D SLAM algorithms in SubT environment. *IFAC-PapersOnLine*, 55(38):126–131.

LaValle, S. M., Kuffner, J. J., Donald, B., et al. (2001). Rapidly-exploring random trees: Progress and prospects. *Algorithmic and computational robotics: new directions*, 5:293–308.

Li, F., Zlatanova, S., Koopman, M., Bai, X., and Diakité, A. (2018). Universal path planning for an indoor drone. *Automation in Construction*, 95:275–283.

Li, Y., Li, M., Zhu, H., Hu, E., Tang, C., Li, P., and You, S. (2020). Development and applications of rescue robots for explosion accidents in coal mines. *Journal of Field Robotics*, 37(3):466–489.

Lindqvist, B., Haluska, J., Kanellakis, C., and Nikolakopoulos, G. (2022a). An adaptive 3D artificial potential field for fail-safe UAV navigation. In *2022 30th Mediterranean Conference on Control and Automation (MED)*, pages 362–367. IEEE.

Lindqvist, B., Kanellakis, C., Mansouri, S. S., Agha-mohammadi, A.-a., and Nikolakopoulos, G. (2022b). Compra: A compact reactive autonomy framework for subterranean mav based search-and-rescue operations. *Journal of Intelligent & Robotic Systems*, 105(3):49.

Lindqvist, B., Karlsson, S., Koval, A., Tevetzidis, I., Haluška, J., Kanellakis, C., Agha-mohammadi, A.-a., and Nikolakopoulos, G. (2022c). Multimodality robotic systems: Integrated combined legged-aerial mobility for subterranean search-and-rescue. *Robotics and Autonomous Systems*, 154:104134.

Lindqvist, B., Mansouri, S. S., Agha-mohammadi, A.-a., and Nikolakopoulos, G. (2020). Nonlinear MPC for collision avoidance and control of UAVs with dynamic obstacles. *IEEE robotics and automation letters*, 5(4):6001–6008.

Papachristos, C., Khattak, S., Mascarich, F., and Alexis, K. (2019). Autonomous navigation and mapping in underground mines using aerial robots. In *2019 IEEE Aerospace Conference*, pages 1–8. IEEE.

Park, M. G., Jeon, J. H., and Lee, M. C. (2001). Obstacle avoidance for mobile robots using artificial potential field approach with simulated annealing. In *ISIE 2001. 2001 IEEE International Symposium on Industrial Electronics Proceedings (Cat. No. 01TH8570)*, volume 3, pages 1530–1535. IEEE.

Protopapadakis, E., Stentoumis, C., Doulamis, N., Doulamis, A., Loupos, K., Makantasis, K., Kopsiaftis, G., and Amditis, A. (2016). Autonomous robotic inspection in tunnels. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 3:167–174.

Ranjan, A., Sahu, H., and Sahu, H. (2014). Communications challenges in underground mines. *Communications*, 5(2):23–29.

Reddy, A. H., Kalyan, B., and Murthy, C. S. (2015). Mine rescue robot system–a review. *Procedia Earth and Planetary Science*, 11:457–462.

Rossi, M. and Brunelli, D. (2015). Autonomous gas detection and mapping with unmanned aerial vehicles. *IEEE Transactions on Instrumentation and Measurement*, 65(4):765–775.

Rouček, T., Pecka, M., Čížek, P., Petříček, T., Bayer, J., Šalanskỳ, V., Heřt, D., Petrlík, M., Báča, T., Spurnỳ, V., et al. (2020). Darpa subterranean challenge: Multi-robotic exploration of underground environments. In *Modelling and Simulation for Autonomous Systems: 6th International Conference, MESAS 2019, Palermo, Italy, October 29–31, 2019, Revised Selected Papers 6*, pages 274–290. Springer.

Rusu, R. B., Blodow, N., and Beetz, M. (2009). Fast Point Feature Histograms (FPFH) for 3D registration. In *2009 IEEE International Conference on Robotics and Automation*, pages 3212–3217.

Seo, J., Park, W., and Kim, T. (2022). Feature-Based Approach to Change Detection of Small Objects from High-Resolution Satellite Images. *Remote Sensing*, 14(3).

Shahmoradi, J., Talebi, E., Roghanchi, P., and Hassanalian, M. (2020). A comprehensive review of applications of drone technology in the mining industry. *Drones*, 4(3):34.

Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., and Daniela, R. (2020). LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5135–5142. IEEE.

Shuai, W., Jiang, F., Zheng, H., and Li, J. (2022). MSGATN: A Superpixel-Based Multi-Scale Siamese Graph Attention Network for Change Detection in Remote Sensing Images. *Applied Sciences*, 12(10).

Sopasakis, P., Fresk, E., and Patrinos, P. (2020). OpEn: Code generation for embedded nonconvex optimization. *IFAC-PapersOnLine*, 53(2):6548–6554.

Stathoulopoulos, N., Koval, A., Agha-mohammadi, A.-a., and Nikolakopoulos, G. (2023a). FRAME: Fast and Robust Autonomous 3D Point Cloud Map-Merging for Egocentric Multi-Robot Exploration. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3483–3489.

Stathoulopoulos, N., Koval, A., and Nikolakopoulos, G. (2023b). Irregular Change Detection in Sparse Bi-Temporal Point Clouds Using Learned Place Recognition Descriptors and Point-to-Voxel Comparison. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8077–8083.

Stathoulopoulos, N., Koval, A., and Nikolakopoulos, G. (2024a). 3DEG: Data-Driven Descriptor Extraction for Global re-localization in subterranean environments. *Expert Systems with Applications*, 237:121508.

Stathoulopoulos, N., Koval, A., and Nikolakopoulos, G. (2024b). A Comparative Field Study of Global Pose Estimation Algorithms in Subterranean Environments. *International Journal of Control, Automation and Systems*, 22(2):XXX.

Tordesillas, J. and How, J. P. (2021). FASTER: Fast and Safe Trajectory Planner for Navigation in Unknown Environments. *IEEE Transactions on Robotics*, 38(2):922–938.

Tordesillas, J., Lopez, B. T., and How, J. P. (2019). FASTER: Fast and Safe Trajectory Planner for Flights in Unknown Environments. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1934–1940. IEEE.

Tranzatto, M., Mascarich, F., Bernreiter, L., Godinho, C., Camurri, M., Khattak, S., Dang, T., Reijgwart, V., Loeje, J., Wisth, D., et al. (2022). Cerberus: Autonomous legged and aerial robotic exploration in the tunnel and urban circuits of the darpa subterranean challenge. *arXiv preprint arXiv:2201.07067*.

Turner, R., Bhagwat, N., Galayda, L., Knoll, C., Russell, E., and MacLaughlin, M. (2018). Geotechnical Characterization of Underground Mine Excavations from UAV-Captured Photogrammetric & Thermal Imagery. In *52nd US Rock Mechanics/Geomechanics Symposium*, pages ARMA–2018. OnePetro.

Urbach, D., Ben-Shabat, Y., and Lindenbaum, M. (2020). DPDist: Comparing Point Clouds Using Deep Point Cloud Distance. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12356 LNCS:545–560.

Wang, L., Ye, H., Wang, Q., Gao, Y., Xu, C., and Gao, F. (2021). Learning-based 3D Occupancy Prediction for Autonomous Navigation in Occluded Environments.

Wang, W., Dong, W., Su, Y., Wu, D., and Du, Z. (2014). Development of search-and-rescue robots for underground coal mine applications. *Journal of Field Robotics*, 31(3):386–407.

Warren, C. W. (1993). Fast path planning using modified a* method. In *[1993] Proceedings IEEE International Conference on Robotics and Automation*, pages 662–667. IEEE.

Xiao, W., Vallet, B., and Paparoditis, N. (2013). Change Detection in 3D Point Clouds Acquired by a Mobile Mapping System. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, II-5/W2:331–336.

Xiong, J., Gao, H., Wang, M., Li, H., and Lin, W. (2021). Occupancy Map Guided Fast Video-based Dynamic Point Cloud Coding. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1.

Xu, W. and Zhang, F. (2021). FAST-LIO: A Fast, Robust LiDAR-inertial Odometry Package by Tightly-Coupled Iterated Kalman Filter, github: `https://github.com/HViktorTsoi/FAST_LIO_LOCALIZATION`.

Yan F., Liu YS., . X. J. (2013). Path planning in complex 3d environments using a probabilistic roadmap method. *International Journal of Automation and Computing*, 10:525–533.

Zeng, J. (2022). The development trend of robots. In *2021 International Conference on Social Development and Media Communication (SDMC 2021)*, pages 1261–1265. Atlantis Press.

Zhang, N., Zhang, M., and Low, K. H. (2021). 3D path planning and real-time collision resolution of multirotor drone operations in complex urban low-altitude airspace. *Transportation Research Part C: Emerging Technologies*, 129:103123.

Zhao, J., Gao, J., Zhao, F., and Liu, Y. (2017). A search-and-rescue robot system for remotely sensing the underground coal mine environment. *Sensors*, 17(10):2426.

Zimroz, R., Hutter, M., Mistry, M., Stefaniak, P., Walas, K., and Wodecki, J. (2019). Why should inspection robots be used in deep underground mines? In *Proceedings of the 27th International Symposium on Mine Planning and Equipment Selection-MPES 2018*, pages 497–507. Springer.