

MINING AND INDUSTRIAL SAFETY TECHNOLOGY AND TRAINING INNOVATION (MISTTI) PROJECT

PROJECT NUMBER: H750H9824

TASK: 3.4 MOBILE ROBOT SCOUT

MAPPING AND NAVIGATION SENSORS FOR THE MINE OF THE FUTURE

FOR

WHEELING JESUIT UNIVERSITY
NATIONAL TECHNOLOGY TRANSFER CENTER
NTTC SUBAWARD #0000007131

AWARD DATE: OCTOBER 29, 2010

END DATE: OCTOBER 31, 2011

FINAL REPORT

NOVEMBER 30, 2011

BY

CARNEGIE MELLON UNIVERSITY
5000 FORBES AVE
PITTSBURGH, PA 15213

CMU AWARD #1041656

William Red Whittaker, Principal Investigator

This report was prepared by:

Uland Wong, PhD candidate - CMU
Calibration, image sensor processing, and analysis

Aaron Morris, PhD – AM Holdings
Range sensor processing and analysis

Warren Whittaker, Engineer – CMU
Operations and integration

James Lee – graduate student – CMU
User interface implementation



Mobile Robot Scout

TABLE OF CONTENTS

1. Introduction.....	1
2. Mobile Robot Scout Rescue Tasks	3
Rescue Team Forward Scout	3
Robot Scout Coverage	3
Ad-hoc Communications	3
Robot Scout as Navigator	4
Mapping and Modeling.....	4
Embedded Responder	4
Transport Mule.....	5
3. User Interface.....	5
Image Display Panel	7
Command Tab.....	7
Status Tab.....	8
Notes Tab	8
Implementation	9
4. Mapping Sensor Evaluation.....	10
Introduction.....	10
Sensor Overview	11
Time of Flight LIDAR	12
Triangulation.....	13
Phase Shift LIDAR	14
Focal Plane Array (Flash) LIDAR.....	15
Characterization Methodology.....	15
Laboratory Evaluation and Ideal Target	16
Environmental Evaluation	17

Metrics for Evaluating Models	19
Ideal Metrics	20
Environmental Metrics.....	20
Data Collection	22
Results.....	24
Application to Mine Rescue Scout	27
5. Robot Configuration	28
Physical Configuration.....	29
Sensors	29
3D Modeling	30
Imaging	31
Position and Orientation	33
Environmental Sensors	33
6. Analysis & Demonstration.....	34
Robot-assisted mine rescue.....	34
Robotic inspection	35
Robotic modeling for “accident scene” documentation.....	38
7. Conclusion and Future	41
8. Acknowledgements:.....	42
9. References.....	43
10. Appendix.....	43
Electronic Files	43
Data Files	44

1. Introduction

Carnegie Mellon University in association with AM Holdings undertook 1 year of research in support of the Mining and Industrial Safety Technology and Training Innovation (MISTTI) project for Wheeling Jesuit University. The research focus was on Task 3 Mine of the Future, subtask 3.4 Mobile Robot Scout.

The work established tasks for a Mobile Robot Scout, developed, refined, and operationally tested a mobile robot scout user interface, conducted analysis of sensors for use in underground mapping and navigation, and produced a culminating demonstration showing the following:

- *Robot-assisted mine rescue where a mobile robot forged ahead of a mine rescue team to relay mine conditions.*
- *Robotic inspection where a mobile robot monitored a specific section of mine for abnormal conditions*
- *Robotic modeling where a mobile robot surveyed and documented an “accident scene” and built a mine model for training and virtual reality*

Mobile Robot Scout represents a technology where a robot is equipped to serve as a component for mine accident response and daily monitoring of a working mine. Outfitted with sensing, computing, and communications the Robot Scout is enabled to model and map, sense the environment, locate itself within the mine, navigate the mine, and operate autonomously. The quality of data that is collected results in models that are geometrically accurate, images that are of high resolution, coverage that is complete, and overlays other sensor data such as gas levels and air flow. Processing of the data provides display of the results from simple maps to virtual reality models. A user interface allows a responding rescue team to directly interact to gather information from the Robot Scout and to direct the scout’s mission. The robot can be locally operated by joystick however autonomy allows the operator to command the robot to move forward reverse or to make a turn with a single command and the robot navigates locally around obstacles until the next command. When given a simple plan or previous map the robot can carry out a mission with no human interaction. Integrated into a working mine, the scout will continuously carry out missions to systematically gather information tied with location and time generating data for analysis and identification of changing conditions.

This report is comprised of 5 main sections covering the work from the project:

- Mobile Robot Scout Rescue Tasks - outlines mine rescue tasks using the robot scout technology.
- User Interface – provides the background and the functionality of the developed web based user interface for the robot scout used on this project.
- Mapping Sensor Evaluation – covers the technical approach to characterization of sensors focused on four components: sensor selection, environment selection, data collection, and characterization methodology.
- Robot Configuration – the robot scout was realized by tasking and reconfiguring an existing mine mapping robot used in mines for this project.

- Analysis and Demonstration – presents the data and principals developed and demonstrated for the culminating demonstration for this project.

Following the 5 sections are conclusions, acknowledgements and references. The appendix presents the contents of the electronic files and description of the data files that accompany this report.

Initiatives in parallel support of this project included:

1. Attending the Mining Health and Safety International Symposium April 7-8 at the Charleston WV Civic Center. This symposium was hosted by Wheeling Jesuit University as part of the MISTTI program. Our team showcased the robot, sensing technologies and mobile robot scout scope under development. The symposium provided an opportunity to interact with mining personnel attending the conference. Breakout sessions and presenters provided additional insight into mine rescue with a robot scout.
2. Supporting a group of graduate level students at Carnegie Mellon in gathering background information for the development of a user interface for a mine robot. The Human Computer Interaction Institute course, HCI Methods (05-610 :: A), taught by Dr. Bonnie John and Dr. Matt Kam had a 6 student team that our project team mentored and worked with to gather background information for the mobile robot rescue tasks as well as the user interface. Literature searches and interviews with rescue teams, rescue trainers, government rescuers, and underground robot companies were conducted to gather this information.
3. Submitting, publishing, and presenting a paper on 3D and image base robot sensors to the IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS 2011) based on the sensor work. The conference occurred September 25-30, 2011 in San Francisco, California. An electronic copy of the paper is in the electronic files accompanying this report.

This project has demonstrated many of the technologies needed for a Mobile Robot Scout relative to mine rescue. Adoption of these technologies into mine rescue is foregone.

[The National Academies Press 2007] “Reviews of Research Programs of the National Institute for Occupational Safety and Health” in Chapter 6: EMERGING AREAS IN MINE DISASTER PREVENTION AND CONTROL recommends:

- *“Continuous monitoring of conditions, especially by remote means, will become increasingly vital as the mining environment becomes more complex. The Mining Program should develop the means to continuously monitor data, with the possibility for response via intelligent system analysis, as the level of complexity increases.”*

This Mobile Robot Scout research provides NIOSH verification of progress towards the future of mining in line with the National Academies recommendation.

2. Mobile Robot Scout Rescue Tasks

The mobile robot scout concept is based on robots that have modeling and mapping capabilities that enable navigation and localization within a mine. The robot scout technology is scalable and can be adapted to control many classes of machinery in support of mine rescue. The guiding principles for a robot scout are do no harm, enhance safety, augment human efforts, and document everything. Research for this task was generated from discussions with government and industry experts, robot developers, and our previous research with mining robots: The following are tasks for a Robot Scout:

Rescue Team Forward Scout

Traveling ahead of a rescue team, a robot scout constantly relays conditions ahead providing an advancing team with situational awareness. Robot scouting can range several hundred feet ahead of a rescue team while still maintaining communications underground. When communications are blocked the system is set to gather advance reconnaissance and return to the team according to an incremental advancement plan. Situational awareness provides the rescue team with imagery, gas sensing, temperature and models of what is ahead. The goal of this mode is to advance a rescue team safely and at a faster pace.

Robot Scout Coverage

As a rescue team advances along the most direct route the robot scout provides coverage of parallel entries, crosscuts and dead ends. In this mode the rescue team stays informed of what is happening adjacent to where they are advancing so they do not pass up a miner or advance past a serious hazard such as a fire, gas pocket, or unstable area of roof support. By itself a robot scout can explore an area following algorithms and react to existing conditions to obtain coverage of an area or section. Since all data is recorded mapped and modeled reviews can verify the exact extent of the coverage obtained. Multi robot teams working together plan and replan tasks providing efficient searches and coverage in the event of a discovery or failure.

Ad-hoc Communications

Ad-hoc communications serve as a means to establish a communications network in the mine. Each robot scout robot locally supports high bandwidth wireless 2 way communications. Multiple node radios can be deployed by a single robot scout to establish a chain of communications back to a team, fresh air base, or directly to the command center. Multiple scout robots provide a dynamic Ad-hoc communication network that advances with a rescue team with the capability to analyze and heal communication links. Multiple scout robots can also provide a junction between existing parasitic communication line broken during a mine disaster. Other communication methods such as fiber or very low frequency through the earth radios may also be deployed from the robot scout. When all other communication modes fail the robot scout provides the capability of message passing by traveling back and forth between rescuers and

trapped miners.

Robot Scout as Navigator

As the robot scout drives within the mine it simultaneously builds a model while tracking its location. When given a prior map the scout localizes the mine features and identifies its position within the map. The robot can localize its position very accurately. In smoke filled areas using sensors such as radar the robot scout navigates the obscured mine corridors while maintaining positional accuracy. Having the ability to know where it is means that the robot records positional data for all of the onboard sensors, it can plan paths to navigate from one location to another, and can be used to confirm locations and guide rescuers or trapped miners.

Mapping and Modeling

The robot scout continuously records and builds models as it navigates through the mine. Models from the robot scout can be low resolution showing only path traveled and structure of the mine to very high resolution providing virtual reality of the area modeled.. Miner training research using virtual reality has been conducted with NIOSH for training miners and rescuers in what to do in smoke filled mines [Orr, et al. 2009]. The robot scout technology provides the data sets needed to develop realistic virtual models of mines and accident scenes for training purposes. During rescue review of modeling information transferred can be analyzed by many sets of eyes both in the mine and outside the mine providing confirmation that areas of the mine have been searched. Models document mine conditions existing at a particular time. Comparing multiple models of an area over time provides insight into developing or changing conditions that are indicators of potential hazards. For instance tracking continuing progression in the floor heave, roof slump and rib sloughing is a good indication of potential for collapse. The high resolution models gathered provide invaluable data for documenting, analyzing and recreating an accident scene. With the increase in capabilities for economically gathering 3-D models it is likely that like other industries mining will adopt modeling over mapping as a standard.

Embedded Responder

Robot scouts embedded and working day to day or dedicated units strategically posted within a working mine can be automatically triggered to respond following a mine accident. Triggered by detection of a blast, unsafe gasses, smoke, or by command a Robot scout now starts the responder function and gathers immediate information from within the mine. Once responder mode is establish and sensing through intrinsically safe sensors that conditions are safe to power up, Robot scout starts the responder mission. Before moving the responder will listen for communications from within or outside the mine, turn on a unique visual and audible homing signal, begin mapping and modeling of the environment, and start a rescue mission plan. Linked into the mine network prior to the accident Robot scout should already have knowledge of the last known location of miners, equipment and the current or planned mine operations. Experience and planning will have Robot scout programmed to carry out missions such as go toward the working area, check all rescue shelters, search for miners, search for the accident source, or establish perimeter of safe advance for rescuers. While carrying out rescue missions, Robot scout will continually source the beacon and listen for communications while recording

sensor data, modeling, mapping, and navigating through the mine. When communications are established Robot scout first relays critical information in a short low bandwidth message with location, mission plan, conditions, area traversed, and findings then listens for confirmation and next command. When a miner is found then robot scout communicates; through display and speaker to relay immediate warning of dangers such as bad air. The robot continues with messages of current location and path to shelter or emergency equipment, how many other miners have been found or still have to be located, and a simple options menu for next action. Robot scout now broadcasts a new signature indicating miner found and listens for response from the miner to decide on the next mission.

Transport Mule

During rescue a mobile robot scout serves rescue efforts by transporting materials and equipment alongside a rescue team, to and from a rescue team, or by navigating through hazards blocking rescuers to supply trapped miners. Rescuers are burdened with adverse conditions and the need to carry extra safety equipment, tools, and materials. Having a robot acting as a mule to transport materials lessens this burden. With the autonomous driving capabilities found in the robot scout a robot can travel alongside rescuers with few commands freeing rescuers to focus on the work at hand. The autonomy also enables the robot mule to shuttle equipment and supplies to and from a rescue team. When hazards are too risky for human entry, or response to an area will take rescuers a great deal of time advance then the rescue team can respond by sending the robot scout. Autonomously the robot scout will carry oxygen, medical supplies, food, and rescue equipment to a designated location for trapped miners.

3. User Interface

The robot must relay vital environmental and diagnostic information to the rescue team. To this end, a web-based graphical user interface, shown in Figure 2, was developed. A web-based control system was chosen because it is platform independent and enables not only support for a multitude of devices but also the simultaneous control and observation received from them. The existing interface was also revamped to include graphical elements such as the two large front and back camera views. These gave remote operators a much richer, higher fidelity view of the surrounding environment increasing situational awareness with no cost to human safety. Figure 1 demonstrates Robot scout operation from an unmodified and widely available consumer device.

The user interface display is comprised of two sections. The upper section is the image display panel showing front and rear images. The lower section of the display is managed through the use of a tab bar to save real estate. The current user interface configuration has Command, Status, and Notes tabs however additional tabs can be added to customize and expand the functionality of the user interface in the future.

Note that the robot also has a hand held low level joystick controller that provides direct control of the driving and steering. This simple device allows a rescuer to take direct control of the robot scout for line of sight driving.

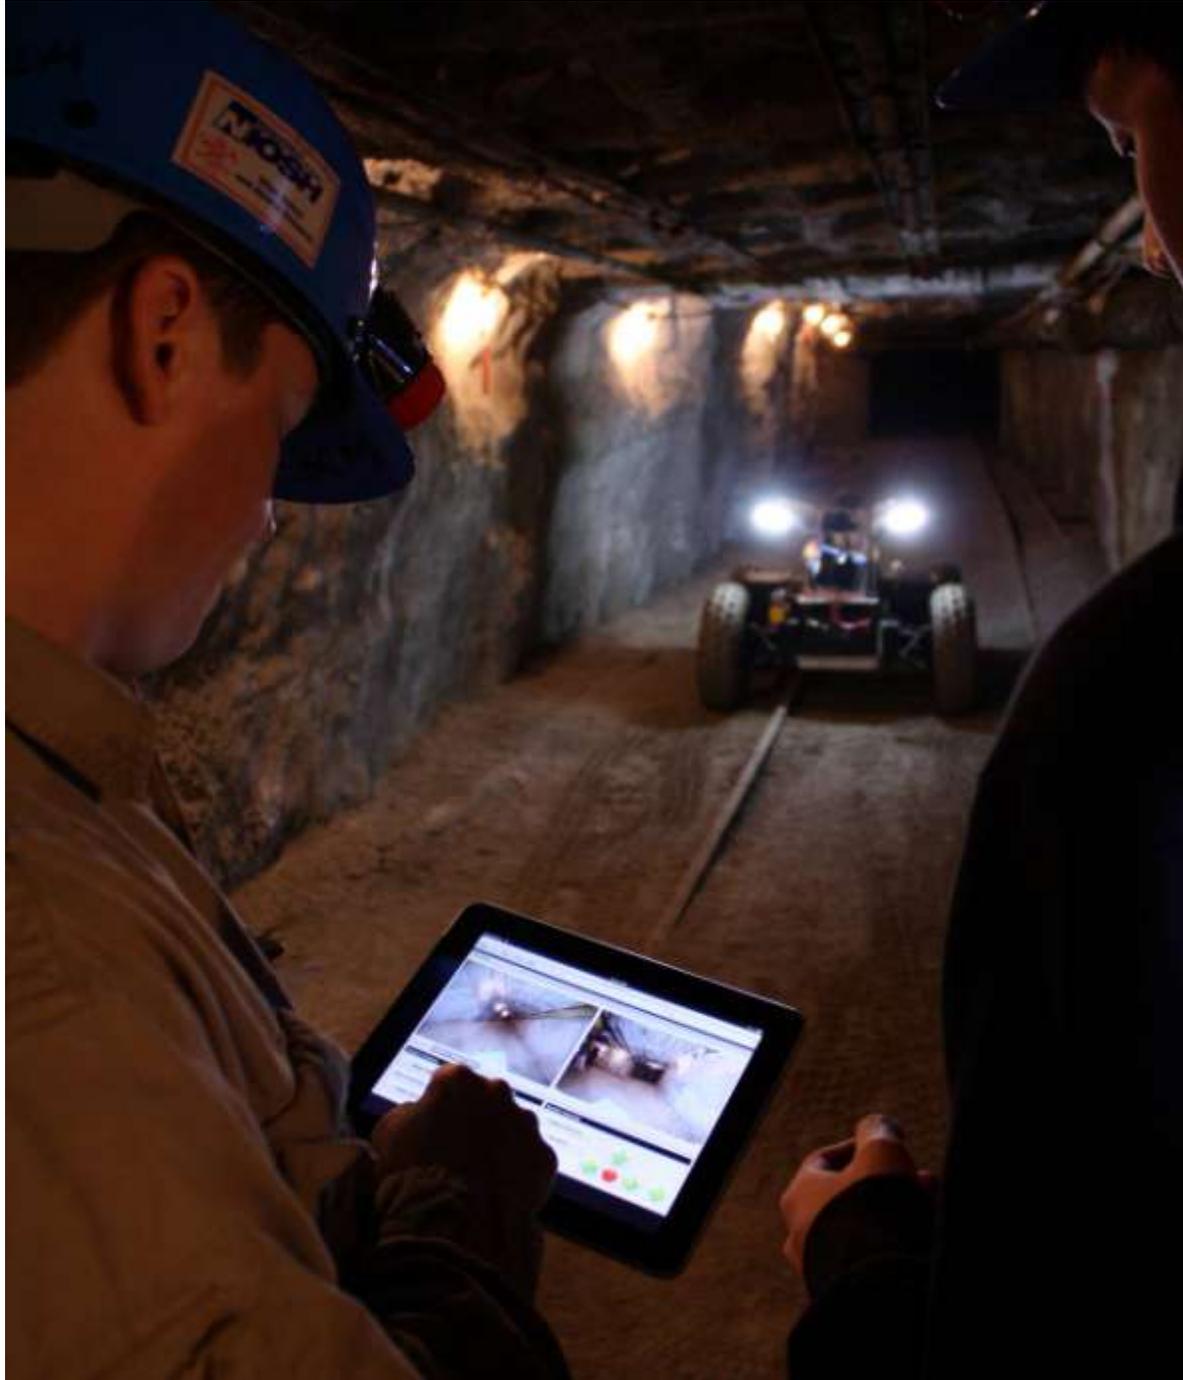


Figure 1. Operation from an Apple iPad

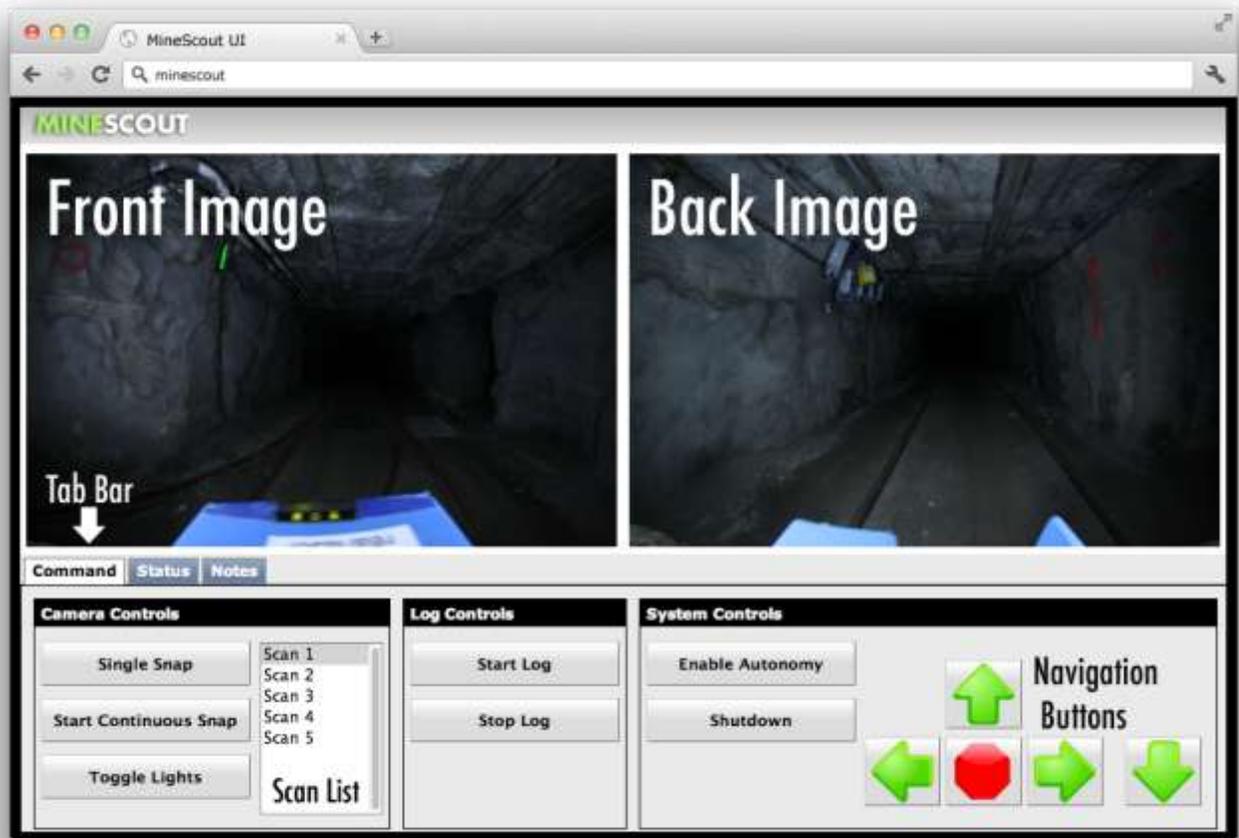


Figure 2. Screenshot of Robot scout web interface.

Image Display Panel

The upper portion of the display is large color Front and Back images providing operators great situation awareness enabling them to inspect surrounding environments and make well-informed decisions. The front and back image displays are updated every few seconds with the most recent photos transferred to the user interface. The section is also used to display and review previously gathered front and back images.

Command Tab

The command tab is the default boot up tab which provides the robot scout operator with the operational controls. This tab has 3 windows for camera control, log control, and system control.

Camera Controls

Single Snap - Snaps a single picture from both cameras.

Start Continuous Snap - A toggle button that enables and disables continuous capture from both of the cameras and enables the lights.

Toggle Lights -A toggle button that enables and disables Robot scout's lights.

Scan List - A scrollable list containing all past images. Clicking on a scan displays the pictures from that scan on the two large camera displays.

Log Controls:

Start Log – Starts the logging of time stamped laser, wheel odometry, inertial status, and commands to a new file.

Stop Log -Stops the logging of data and saves the current file.

System Controls

Enable Autonomy - A toggle button that enables and disables Robot scout's autonomous mode.

Shutdown - Performs a safe shutdown of Robot scout.

Navigational Buttons -Large navigation buttons enable driving control of the robot in autonomous mode. The 5 buttons are turn left, turn right, forward, stop, and reverse. The autonomous mode allows the robot to continue driving intelligently avoiding obstacles from a single command.

Status Tab

The Status Tab displays information about the environment as shown in Figure 3. Date, time, battery voltage, temperature, air velocity and gas readings are displayed and continuously updated. If gas readings, temperature, or battery status reaches predetermined limits then the background color of the display changes and the user screen displays an error message.

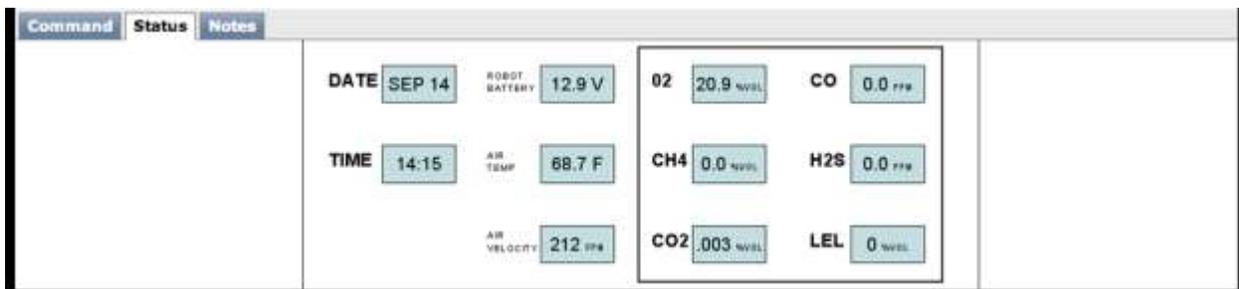


Figure 3. Robot scout status tab

Notes Tab

The Notepad tab, Figure 4, provides an interface for operators to jot notes, document hazards, and share information, messages are logged with time stamps.

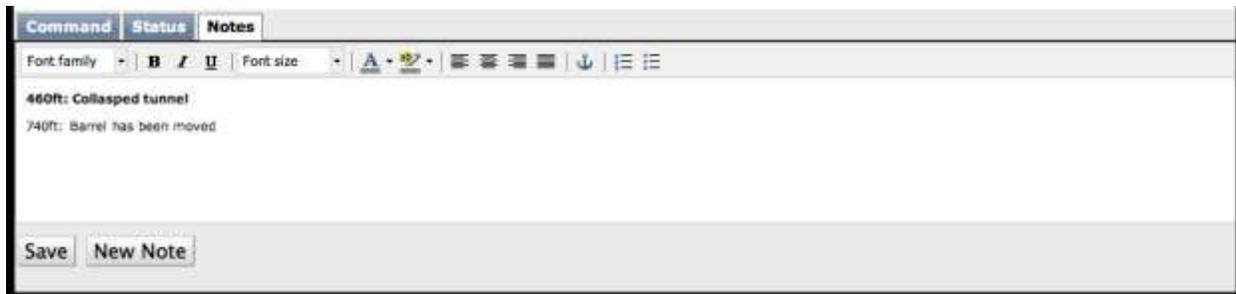


Figure 4. Robot scout notes tab

Implementation

The bulk of the user interface, i.e. control, display, web-server, and robot communication components, were implemented in the C++ coding language. Stylistic elements were designed first in Photoshop and later laid out with CSS and XHTML. The C++ open source libraries OpenCV, Wt, and Boost were also used for some of the image manipulation, web server and process control modules.

4. Mapping Sensor Evaluation

In most general terms, maps are geospatial representations of the information. Robots use maps to navigate, interact and document information about the world. Much like humans who depend on eyesight to get around, all autonomous robots must acquire 3D map information in order to function, even if mapping is not a primary objective. The quality of maps critically affects the dependability and survivability of robots in unknown environments. Sensors which create maps are called *range sensors* and many models, types and manufacturers of range sensors exist. Sensing systems on robots are often designed to balance functionality, robustness and cost by selecting the components most appropriate to objectives; the so-called *sensor selection* problem. However, mine rescue is a largely unstudied application. It is both unclear which sensors are optimal for the required mine rescue tasks and what parameters inform optimality.

As a part of the Robot scout project, a broad evaluation of range sensors was conducted. Specifically, sensors were evaluated on the parameters of mapping capability, the "missing piece" of the sensor selection puzzle that does not yet exist for underground robots. Unlike physical quantities such as mass, power or volume, knowledge of mapping performance in the field cannot be obtained from datasheets or factory calibrations. Only by evaluating and comparing mine models generated from multiple sensors can these quantities be acquired. The following section details the methodology and analysis in characterizing sensors for this project.

This evaluation has advanced the state of the art knowledge in underground sensor selection and the results serve in the design of Robot scout detailed in the rest of this report. It is our hope that future underground robot designs will benefit from the methodology documented herein. Portions of the work funded under this grant were published in the proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011) as "Comparative Evaluation of Range Sensing Technologies for Underground Void Modeling."

Introduction

The technical approach to characterization of sensors has four thrusts: sensor selection, environment selection, data collection, and characterization methodology. Range sensors of interest, best matching rescue application in terms of functionality, were initially down-selected for evaluation to limit scope. Environment selection ensured that the specific locations where evaluation was conducted could be sufficiently generalized to many mine rescue situations. Data collection entails the actual experimentation conducted and data sets logged for evaluation. Lastly, characterization methodology details the process for comparing and evaluating the results from the selected sensors.

Experimentation includes both a laboratory calibration, utilizing ideal targets and a holistic comparison of models generated in representative underground environments. Both these comparisons utilize a common, repeatable methodology across all sensors and metrics of evaluation (sampling accuracy and density) that are inspired by the aforementioned real-world applications. It is noted that only the quality of geometric reconstruction from sensor data is evaluated. Cost, mass, energetics and field robustness are specifically deferred for future work, as addressing all implementation issues would lead to intractability of experimentation in this developing work. Moreover, factory values for these parameters are generally sufficient for

decision making.

The project delivers range sensor data sets from subterranean tunnel environments and analysis of these datasets in accordance to underground navigation practice. Models and imagery are presented in common formats available for import and use in other programs. The baseline models, calibration infrastructure and methodology exist to provide future evaluation of similar range and imaging sensors. A technology roadmap and set of recommendations for underground robotic system integration is provided.

Sensor Overview

Sensors were selected based on prevalence in robotics usage and availability to the authors (refer to Fig. 5 in and Table 1); a total of 10 sensors were evaluated. Experimental configurations (i.e. actuation, physical parameters, and external illumination) were chosen to reflect optimality for void modeling at a critical sensing distance of 2-8 meters as recommended in [Omohundro 2007], a definitive treatise on underground mapping configuration. This study is not intended to be a comprehensive sampling of sensor configuration parameters, but rather a broad sampling of sensor types specific to the application. For example, a baseline of 250mm and infinite focal distances were utilized for stereo vision; a less common configuration found in indoor robotics. There is no claim that results generated herein are strictly valid for any sensors or configurations other than those evaluated.



Figure 5. Illustration of Evaluated Sensors and Configurations – (1) rotating Hokuyo UTM-30LX, (2) rotating SICK LMS111-10100, (3) rotating SICK LMS291-S14, (4) rotating SICK LMS511-10100, (5) rotating SICK LMS200-30106 affixed on a mobile robot, (6) Faro Photon80, (7) IFM O3D 201, (8) custom structured light sensor, (9) custom stereo vision sensor and (10) Microsoft Kinect.

The following section provides a look at the sensors, configuration and data collected in characterization. The underlying technology of each sensor is discussed briefly. Raw data previews presented here are all from the same "unstructured" mine corridor at Bruceton Research

mine for comparison (see environmental evaluation section below). The manufacturer's data sheets for each sensor used are included in the electronic appendix for this report.

Table 1. Evaluated Sensors and Technologies

Sensor Model	Technology	Evaluated Configuration
SICK LMS200-30106	Planar ToF LIDAR	0.5 x 180 degree rotating, 8m mode
SICK LMS291-S14	Planar ToF LIDAR	0.5 x 90 degree rotating, 8m mode
SICK LMS111-10100	Planar ToF LIDAR	0.25 x 270 degree rotating, 20m
SICK LMS511-10100	Planar ToF LIDAR	0.5 x 190 degree rotating, 24m clipped
Hokuyo UTM-30LX	Planar ToF LIDAR	0.25 x 270 degree rotating, 24m clipped
Structured Light*	Triangulation/ Structured Light	PtGrey Scorpion w/ projector (1280x1024), 0.25m baseline
Microsoft Kinect	Triangulation/ Structured Light	Off the shelf configuration, libfreenect, ~5m range
Stereo Vision*	Triangulation/ Stereo Vision	2x Prosilica GC1290 (1290x960), ELAS, 0.25m baseline
IFM O3D 201	Flash LIDAR	Off the shelf, ~8m range
Faro Photon80	Phase-shift LIDAR	Off the shelf, 1/6x, 24m clipped
Velodyne HDL-32E	Actuated 3D ToF LIDAR	Off the shelf, 30m range

* denotes an in-house implementation

Time of Flight LIDAR

Time of Flight LIDAR sends a laser pulse in a narrow beam and measures the time taken for the pulse to return after reflection. Distance is calculated from the delay and knowledge of the constant speed of light. Most ToF lidars used in robotics application today are planar scanners in which a single point laser is rastered horizontally using a spinning right-angle mirror. By rotating or nodding the scanner, a volume of 3D data is obtained.

Multiple planar ToF sensors from SICK AG (LMS111, LMS200, LMS291, LMS511) and Hokuyo Ltd. (UTM-30LX) were characterized in this study to reflect the prevalence of LIDAR scanners in robotics. Three-dimensional data was collected by mounting to a 360 degree rotary actuator. This configuration was described by [Omohundro 2007] as being optimal for void modeling as compared to simpler nodding configurations. The rotary mount was actuated at a slow 1deg/sec to enable selective angular downsampling in post processing. As optical offsets are likely in mechanical actuation, proprietary dewarping software was utilized to estimate offset parameters.

Additionally, one self-actuated 3D ToF sensor (Velodyne HDL-32E) was characterized for this study. This sensor utilizes the same technology as planar ToF LIDAR; however, instead of a single beam, multiple single-beam sensors (32 in the case of the HDL-32E) are utilized in a common housing to produce a vertical spread on the swept axis. This configuration enables sensing a 3D volume with only a single axis actuation (rotational movement of the optical head

for the HDL-32E).

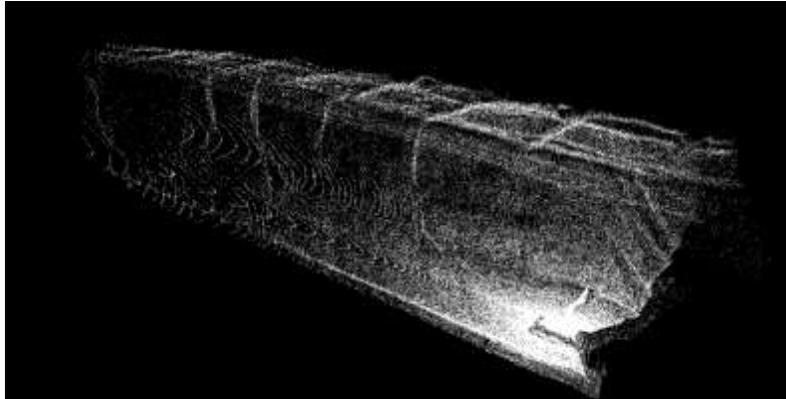


Figure 6. Sick LMS111 point cloud data from the unstructured mine corridor. Time-of-Flight LIDAR data is a balance of density and accuracy, but banding effects are visible from non-uniform mechanical actuation.

Triangulation

Triangulating range sensors use the principle of intersecting geometric rays to measure distance to scene points. These rays must originate from two points with fixed and known translation. Two types of vision-based triangulating sensors were evaluated: stereo vision and structured light.



Figure 7. Stereo Cameras use the parallax between two offset images to recover 3D information. The quality of construction is dependent on the specific lighting, texture and clarity of the images, which changes between views.

Stereo cameras measure the distance to objects by comparing two images of a scene from two cameras and recognizing common features across the images. Triangulation of these “correspondences” with a known baseline produces depth. For testing, two Prosilica GC1290c 1.2MP cameras with 5mm lenses were parallel-mounted with a horizontal baseline of 250mm

and calibrated. Lighting was provided in close proximity to (but behind the field of view of) the camera system during the underground modeling. Feature detection and point cloud generation was accomplished using the LIBELAS software for dense stereo matching



Figure 8. Structured light vertical line projection in unstructured mine corridor.

Structured light sensors solve the correspondence problem by utilizing a light source (the dual of a camera) to unambiguously “paint” the scene with known identifiers. A single, static camera observes the change in brightness of scene points over time, which can then be decoded to generate correspondences between the source and camera. Two structured light systems were evaluated for testing: a custom sensor constructed by integrating a visible-light LCD projector with a Point Grey camera and the infrared Kinect™ sensor from Microsoft. The in-house solution utilizes a 40 degree field of view for both the camera and projector, which are angled 15 degrees inward with a 250mm baseline for optimal coverage and accuracy at two meters from the sensor. The Kinect was used in an off-the-shelf configuration with libfreenect to provide a 1-frame depth estimate.

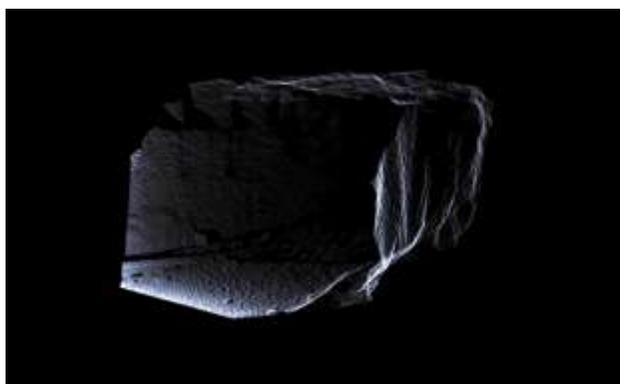


Figure 9. Kinect™ point cloud data was surprisingly dense and accurate, even in the unstructured mine environment, however maximum range is quite short.

Phase Shift LIDAR

Phase shift, also known as FMCW (Frequency Modulated Continuous Wave), sensors

continuously scan the scene with a periodic wave of light, modulating the frequency in a known, repeatable manner. The detected change in phase and frequency of the returned signal give the distance to the scene. This method of measurement produces high sample density and accuracy as there is no need to wait for a return pulse between measurements as in time-of-flight technology. A Faro Photon80 (FMCW LIDAR) was characterized for this study.

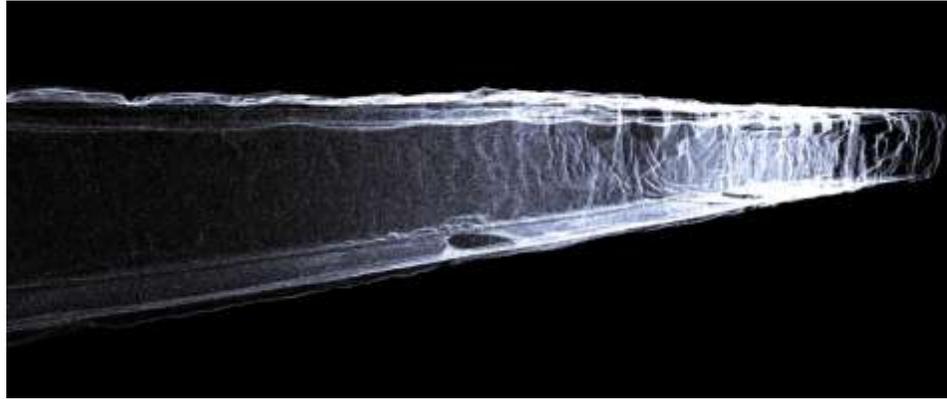


Figure 10. FMCW LIDAR point cloud of unstructured mine corridor. Data from the Photon80 unit is dense, accurate and uniformly distributed.

Focal Plane Array (Flash) LIDAR

Flash LIDAR operates like a digital camera, gathering information in a 2D array of pixels through lensing. Each pixel simultaneously measures a distance and reflective intensity each time a pulsed laser in the unit flash-illuminates the scene. The resulting snapshot provides a 3D model with reflectivity data. The IFM O3D201 sensor is a 50x62 pixel, 10fps flash LIDAR that was used in testing.

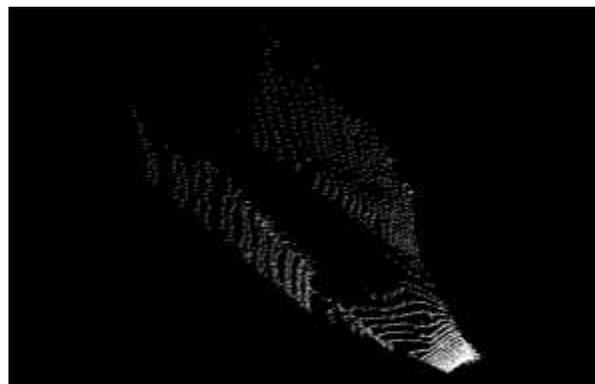


Figure 11 Flash LIDAR point cloud of unstructured mine corridor illustrating characteristically sparse data and short range. Ceiling and floor data is often not captured.

Characterization Methodology

Methodology for comparative analysis of range sensors comprises two parts: (1) characterization

in a lab utilizing an ideal, artificial target and (2) in-situ characterization in representative underground environments. The purpose of the lab calibration is to establish a reliable baseline of comparison and to estimate uncertainty bounds for the sensors. As the true geometry of unstructured void environments cannot be known to arbitrary precision, the other sensors must be compared against the best performing sensor determined in lab trials, with uncertainty values informing a lower bound on error. Environmental characterization effectively stresses sensors with the gamut of materials, features and scenes likely to be found in the underground spaces of interest. Thus, in-situ results enable a holistic view of performance that closely mirrors application intent. Both lab and in-situ characterization are necessary to paint a complete picture of sensor quality.

Laboratory Evaluation and Ideal Target

Laboratory calibration involves scanning a 1.25m x 1.25m, tiled and colored 3D checkerboard from controlled viewpoints. While such “ideal” targets do not exist in field application, their artificial nature enables construction and knowledge of the true geometry to arbitrary tolerance. This information is useful in determining the true error of range sensors, which cannot be surmised in unstructured environments, as well as for testing the rare “edge cases” of sensor error. The checkerboard utilized is constructed to a tolerance of 1mm, beyond the expected accuracy of most contemporary range sensor technologies.

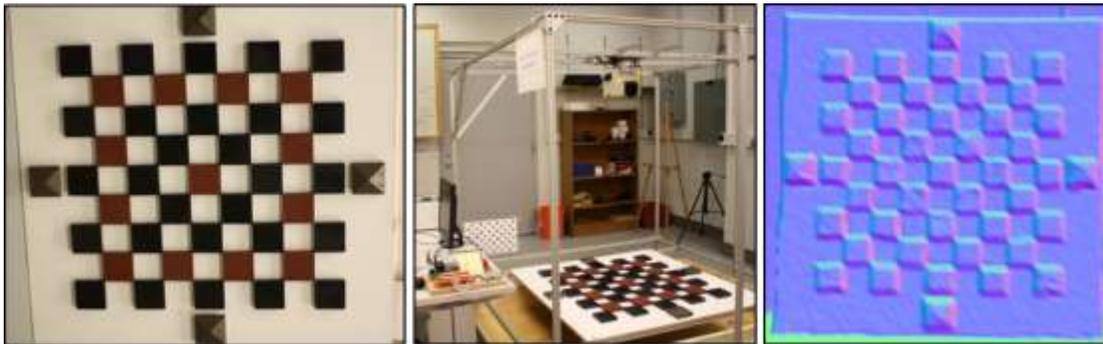


Figure 12. A 3D checkerboard target used for ideal characterization (left), example experimental setup (LMS291 shown) for scanning the checkerboard (middle), and a mesh model of checkerboard generated using range data (right).

Features of the checkerboard are illustrated in Fig. 12. Two colors of semi-gloss dark tiles, raised 1.9cm (0.75”), are mounted on a neutral white diffuse backplane. The colors are chosen to represent a very approximate distribution of underground albedos and reflectances. The varied surface reflectance and color of the tiles enables characterization of sensor error as affected by target material. The pyramidal tiles located on the cardinal points of the target rise 3.8cm (1.5”) from the backplane and are used to automate the process of aligning scans as well as testing pin-point sampling.

The target is centered such that the normal ray of the sensor passes through the middle tile. Scans are taken at a distance of 2.0m from the sensor origin and repeated for primary angles of 90 (normal), 67.5 and 45 degrees. Calibrated mount locations on a support frame provide ground truth for sensor positions from which the checkerboard is scanned (Fig. 12).

Raw output from sensors is first transformed to point clouds with minimal filtering (no-return,

max/min range). Point cloud data is then aligned with the ideal checkerboard model. While approximate sensor and target orientation are known, rotational ambiguities, inaccuracies in mounting and the intrinsic properties of the sensor result in error in raw data. Moreover, while the target may be oriented at a number of angles, the ideal model and error analysis assumes a fronto-parallelism. Utilizing initial estimates of sensor pose, the processing algorithm automatically detects the corner features of the checkerboard and finds a rigid transformation to the known model. A numerical optimization method, iterative closest point (ICP), is then used to fine-tune the alignment in the presence of non-rigid distortions and noise (Fig. 12). Points detected as the raised tiles are colored red, while points detected as part of the back plane are green. Corners of the files are marked with blue '+'s. Several statistical measures of quality are computed.

Checkerboard data was collected for all 10 sensors. This data was evaluated using two metrics inspired by modeling applications, detailed in the metrics section.

Environmental Evaluation

Environmental characterization comprises scanning surveyed underground voids from unknown, yet repeatable views (for example mimicking robot movement). The resulting scans are stitched into a map, if applicable. Three environments were selected for in-situ mapping evaluation (Fig. 13): an unstructured corridor, an unstructured intersection and a closed structured tunnel. These selections were motivated by finding environments whereby the results would be broadly applicable. This was accomplished by generalizing the physical layouts and surface materials found in underground spaces.

Geometric building blocks of most underground voids were first identified. It can be seen that most underground voids are rotationally or translationally symmetric to *corridors* and *intersections*. For example, modern coal mines comprise rectangular grids of corridors and intersections while most caves, lava tubes and tunnels are simply lengthy corridors. Dead ends exist, but these are largely uninteresting as they are simply fronto-parallel walls in the corridor case and allow for minimal robot movement. Thus, environments consisting of these two physical layouts would be sufficient to cover most underground spaces.

Most natural underground spaces consist of rough, diffuse rock which is often covered with a further diffusing dust. This is contrasted with many civic underground spaces such as tunnels or sewers which are smooth in nature and may contain coats of paint. This provides a clear delineation of underground materials: rough and diffuse (henceforth *unstructured*) versus smooth and possibly glossy (*structured*). A combination of these environments may exist in any given space; some mines, for example have structured entryways and unstructured cross cuts.

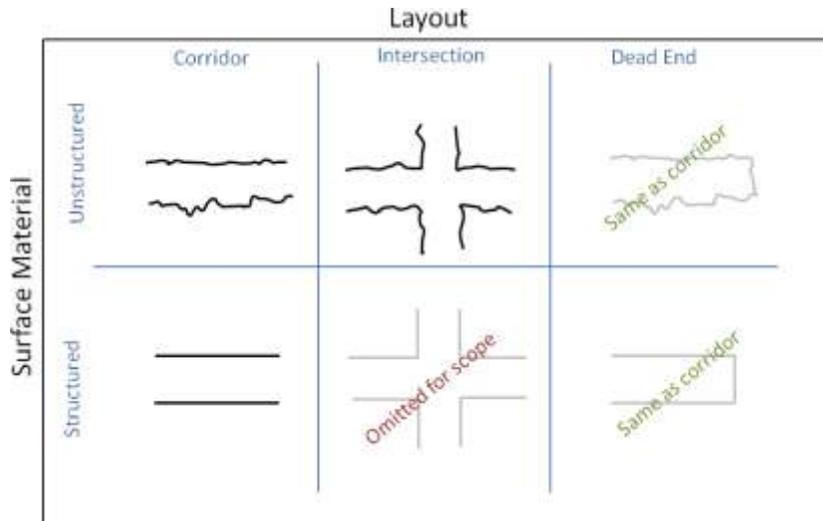


Figure 13. Generalization of Underground Environments. Several of these environments are duplicates from a sensor perspective and one uninteresting environment was omitted for scope.

By plotting the discriminative parameters of material and geometry, four possible generic environments are identified (Figure 13). Of the four, one of these (structured intersection) is rare and largely uninteresting to this project, and thus was omitted for reasons of scope. Experimental locations for the rest were determined by accessibility and generality.

The unstructured environments mapped are located at the Bruceton Research Coal Mine in Pittsburgh, Pa. The corridor is a 60m tunnel with approximately 1.75 x 2.0m cross section, of which 20m was mapped with sensors. The intersection scene is a complex cross intersection of a wide corridor (3m) and a narrow corridor (1m). The primary construction in the mine environments is coal-dust covered gunnite with a polyester mine curtain in the corridor and a wooden cribbing roof support in the intersection.

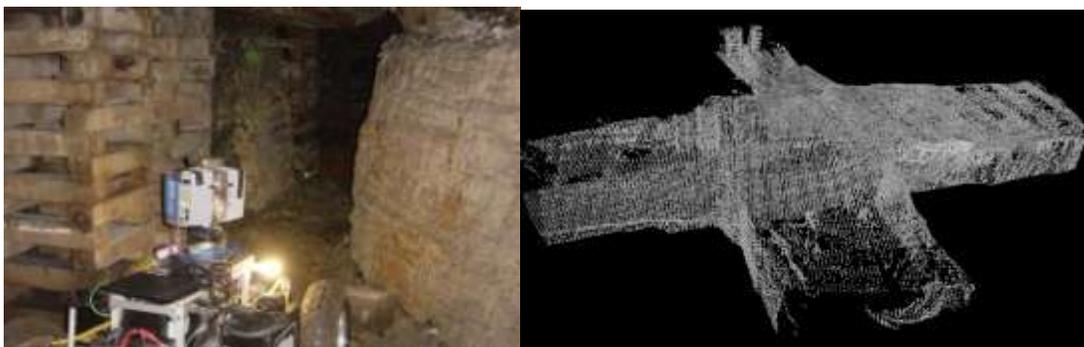


Figure 14. Unstructured Corridor photograph (left) and point cloud (right). The Photon80 sensor and mobile robot are visible in the image.



Figure 15. Unstructured Corridor photograph (left) and point cloud (right). Blue calibration cubes are visible in the image.

The unstructured corridor was mapped by incrementing sensors longitudinally along a length of 20m. The number of scans (i.e. incremental distance) was calculated from the range and modality class of each sensor and ranged from 3, 7 or 20 scans. Blue fiducial cubes were placed in surveyed locations to enable fast stitching of incremental models, though the final alignments were tuned using iterative closest point (ICP). The intersection was mapped from a single static location, but sensors with narrow field of views were rotated to cover the scene. The structured corridor was mapped from a single location, aimed at the dead end. Dimensional and usage restrictions on some of the sensors prevented data collection from every sensor in all environments (ex. the LMS200 was affixed to a large mobile robot).



Figure 16. Structured Corridor photograph (left) and point cloud (right).

The structured corridor is an indoor underground tunnel of 2.0 x 2.0m cross section at Carnegie Mellon University, on the site of the former Bureau of Mines buildings. It is made of a smooth concrete, painted white, and approximately 20m in length with a dead end (ending at an elevator). A series of heat and water pipes line one of the walls. The blue fiducial cubes were required for stitching multiple scans as the entire tunnel is largely translationally symmetric.

Metrics for Evaluating Models

Metrics are transformations of high-dimensional, complex, multi-modal data to a common low-dimensional framework (which may consist of several scalar values) where they can be compared. The most important differences in the data must be amplified and summarized for

clarity. Properly designed metrics are perhaps the most important element in evaluating maps from different sensors. Scores should reflect the overall desirability of a specific map, factoring differences in map size, distribution and erroneous information. Yet, a map's score must not be biased with regards to how the data is generated or form factor of the sensor. This information is incorporated later in the sensor selection process, thus map-making metrics must be wholly independent from other selection parameters.

Metrics selected for sensor evaluations were inspired by common use of underground mine maps in inspection and production estimation. Two questions arise frequently: how accurate are the maps and how much of the area is mapped? The intuitive metrics of *coverage* and *accuracy* are attempts to answer these questions.

Similar but separate metrics are utilized in laboratory and in-situation environmental analysis. The ideal metrics establish an approximate baseline ordering of performance, such that the same sensors can be evaluated on environmental data without ground truth. The following section delves into the specific details of how these metrics are calculated and how estimation differs in the ideal and the environmental case.

Ideal Metrics

Ideal evaluation of accuracy and coverage comprises comparing scans of the checkerboard target to its ideal fabrication model. The full target was imaged within the field of view of every sensor, while background and other areas were not utilized in calculating performance.

Accuracy (Range Error). The range error is the error between an observed data point and its known true location for a single measurement. The mean of the error distribution is a common measurement of *accuracy*. The range error used in this paper is calculated by aligning sensor data of the target to the ideal model using ICP and then ray tracing the datapoints from the sensor origin. The L_2 -norm of the difference is the reported value. A large range error indicates an inaccurate or poorly calibrated sensor. The standard deviation of the range error is also an important measurement known as *precision*.

Coverage (Inter-vertex Distance). A frequent objective of 3D scanning is to create a mesh model or to infer surface geometry for object recognition. Both these applications require dense and regularly distributed surface samples. Inter-vertex statistics are generated by performing a 2D Delaunay triangulation (*delaunaytri* function in matlab) on the surface points and measuring the distribution of resulting triangle side lengths. Large inter-vertex distances are indicative of “holes” in the model while a large variance in the metric is indicative of badly shaped triangles. This statistic reflects the *coverage* of measurements on the target, which is an amalgam of angular density, sample rate, and field of view. Many actuated sensors generate gratuitous readings but lack angular resolution in one or more axes, and thus exhibit inferior performance in resolving objects as compared to lower-throughput, fixed-resolution sensors.

Environmental Metrics

Metrics employed to assess data collected from the sample environments are extensions of those used in ideal analysis with a few differences. These differences arise primarily from the fact that

“ground truth” geometry of the environment is unknown and unobtainable. Chiefly, calculations are patch-based (quantized) to enable simple analysis on stitched models of multiple viewpoints without requiring 3D Delaunay triangulation (a time consuming and error-prone process) and to establish approximation of true values from densely distributed local data. Where data density was insufficient for patch based analysis, less-accurate aligned data from another sensor was used for comparison. Secondly, as no sensor captured the entirety of the surface area in any environment, it is necessary to include the notion of total coverage as well as density in determination of the coverage metric. Analysis must consider not only how far apart measurements are spaced, but also what areas have been left unmapped due to short range or small field of view.

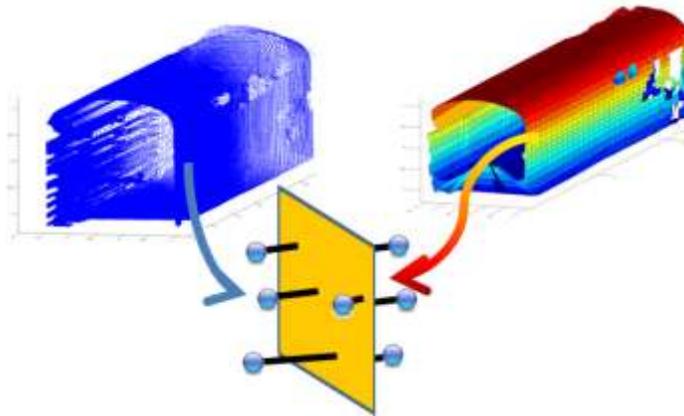


Figure 17. Range Accuracy.

Accuracy over surface patches (VPSP). Though dense, mean-centered samples may correctly approximate volumetric properties of the environment, high variance (excessive noise) results in “blurry” models, making minute features difficult to distinguish. Range variance is computed by applying a squared error statistic to the 10cm x 10cm surface patches. The error used is the normal distance from each measurement to the fitted plane for each patch. Results summarized in Table 4 describe this metric in terms of Variance-Per-Surface Patch (VPSP).

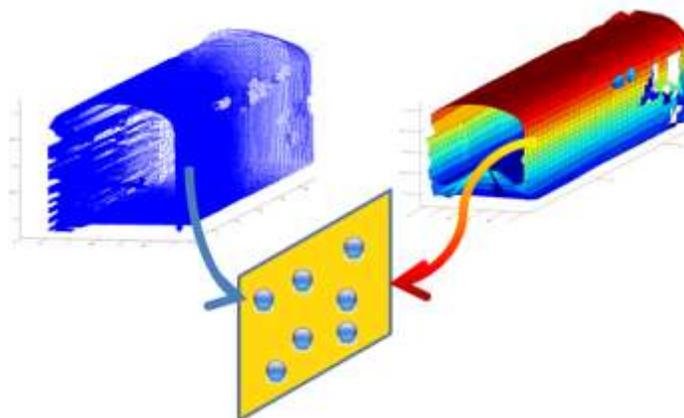


Figure 18. Surface Density.

Density on surface patches (PPSP). Quality feature extraction from range data requires sufficient sample density. Sparse range data presents limited geospatial content and difficult determination of point associations. Moreover, robotic navigation processes such as obstacle detection, trajectory planning and other perception approaches rely on sufficient surface coverage to work properly. Surface coverage density is computed by counting the number of points that fall within two dimensional surface patches. The patches are approximately 100cm^2 or $10\text{cm} \times 10\text{cm}$ in size (see Fig. 5). An upper threshold of 20 points (designated as “sufficient” coverage) was set, with the remaining cells containing 20 or less points. Thus, a curve of density vs. distance is generated. Range sensors typically display concentrated density in the near field, which tapers off with increasing distance. Results summarized in Table 4 describe this metric in terms of Points-Per-Surface Patch (PPSP).

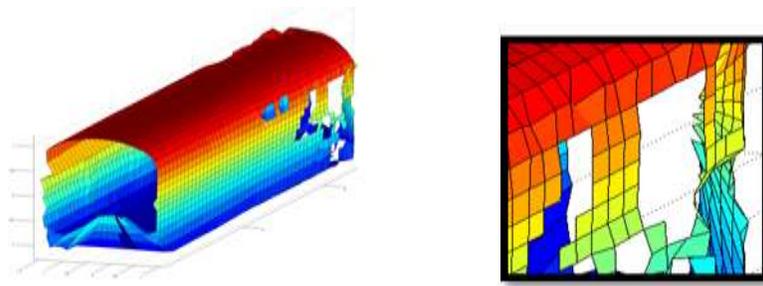


Figure 19. Surface Coverage.

Total surface coverage (PSC). Long range feature identification can increase the accuracy and reliability of position estimation and fusion of multiple models. Short range or small field of view not only results in fewer distinguishing features for extraction, these features also have a smaller baseline, resulting in greater uncertainty. Moreover, short sensing horizons indirectly affect quality in automated model building by inducing erratic trajectory behavior. Lastly, complete surface coverage is required for accurate volumetric analysis of void models. Surface coverage is calculated by tallying patches that have been sampled. Patches associated with one or more points is labeled as “occupied.” Patches with no points are labeled “unoccupied.” Surface coverage is therefore a ratio of occupied patches to a total patch count (determined manually from the size of the environment). This metric can be seen as a scalar summary of the density curve. Results summarized in Table 4 describe this metric in terms of Percent Surface Covered (PSC).

Data Collection

Sensor data collection occurred on 10 dates over a period of 8 months. Setup and primary experimentation occurred mostly during the first two months of the project, with further field experimentation as new sensors were acquired. Table 2 and Table 3 below document collection experiments and specific breakdowns of data by environment respectively. This program resulted in 48GB of registered 3D range and image data and almost 60m combined of corridors mapped with 10 sensor types. These represent the densest, co-registered, multi-modal underground maps ever generated to date.

Table 2. Summary of Data Collection Activities

Date	Location	Activity
09/01/2010	Bruceton Mine	O3D201 data collection in unstructured corridor and intersection
09/10/2010	Bruceton Mine	Photon80 LIDAR data collection in unstructured corridor and intersection for baseline comparative model
10/19/2010	CMU Tunnel	ToF LIDAR Collection in structured corridor (LMS111, LMS200, LMS291)
10/20/2010	CMU Lab	LMS200 and LMS291 data collection on ideal checkerboard target
10/22/2010	Bruceton Mine	ToF LIDAR, O3D201 and Structured Light data collection in unstructured corridor and intersection
10/26/2010	CMU Lab	LMS111, structured light, stereo, O3D201, Photon80 data collection on ideal checkerboard target
02/22/2011	CMU Tunnel	UTM-30LX, Kinect, LMS511 data collection in structured corridor
02/25/2011	Bruceton Mine	UTM-30LX, Kinect, LMS511 data collection in unstructured corridor and intersection
02/28/2011	CMU Lab	UTM-30LX, Kinect, LMS511 data collection on ideal checkerboard target
05/26/2011	Bruceton Mine	HDL-32E data collection in unstructured corridor and intersection

Table 3. Environment and Data Information

Sensor	Structured Corridor (20m)	Unstructured Intersection	Unstructured Corridor (20m)
LMS200	n/a	1 scan	1 scan
LMS291	3m increment	1 scan	1 scan
LMS111	3m increment	1 scan	1 scan
LMS511	3m increment	1 scan	1 scan
UTM-30LX	n/a	n/a	1 scan
Structured Light	1m increment	3 scans	1 scan

Kinect	1m increment	3 scans	1 scan
Stereo Vision	1m increment	1 scan	n/a
O3D 201	1m increment	3 scans	n/a
Photon80	7m increment	1 scan	1 scan
HDL-32E*	7m increment	1 scan	n/a

*HDL-32E data was collected, but not evaluated due to time and data usage constraints from the manufacturer

Results

The following section describes the results of analyzing map data utilizing the described metrics. A scatter plot that summarizes relative sensor performance on the ideal target using the metrics described in the previous section is shown in Figure 20. The x-axis (range error) is the empirical value of the accuracy and the y-axis (tri-neighbor inter-vertex distance) represents density. Sensors closer to the origin (zero) have better performance. The colored ellipses represent the uncertainty in the estimation of this value and are scaled by a factor of 2 for clarity. Experimental error, such as deviations in mounting and data capture, as well as noise generated in the physical sensing process contribute to greater uncertainty.

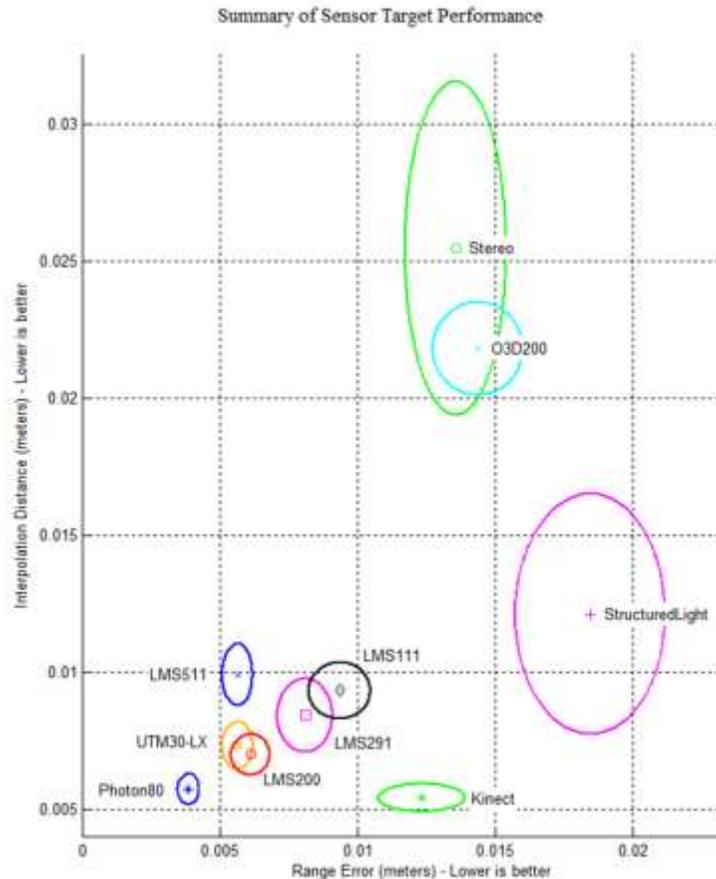


Figure 20. Summary of Ideal Target Performance

The results show a natural grouping of the sensors into three performance classes. The Faro Photon80 was in a class of its own in regards to both metrics: a conclusion consistent with its pricepoint. As-built and survey LIDARs such as the Photon80 are designed to trade portability for maximal modeling performance.

All five planar time-of-flight sensors exhibited similar performance in a class below the Photon80, which is consistent with manufacturer specification and intended application. The LMS200, which has been a staple on underground modeling robots due to its lack of built-in filter, ties the LMS511 in accuracy and nominally wins out over the others. It should be noted that software issues prevented the LMS511 from operating at the highest angular resolution, though accuracy was unaffected. Had the sensor been capable of the factory maximum 0.125deg resolution, it likely would have been the best performing ToF LIDAR.

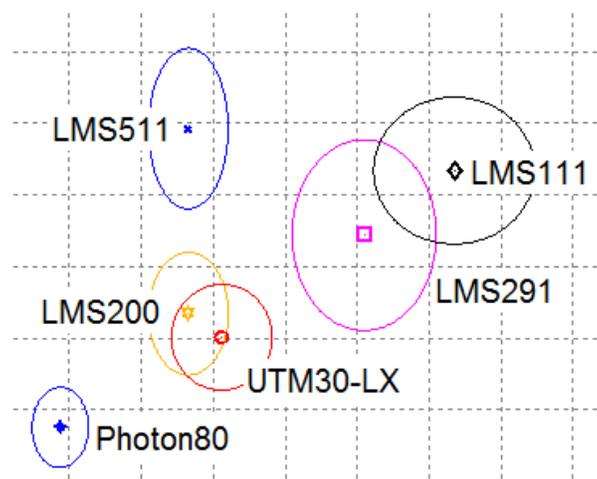


Figure 21. Closeup of planar Time-of-Flight LIDAR performance.

Inconsistent performers comprise the last class of range sensors. These sensors feature notable shortcomings in one or both of the metrics. Flash LIDAR is a nascent technology for outdoor sensing; the IFM O3D has comparable performance to the in-house designed stereo and structured light sensors, though it exhibits marginally better balanced performance and higher robustness. The structured light sensor has high range error arising from poor reflectivity that affects localization of light stripes at the highest scale, but the pattern and consistency of identified points is uniform and dense resulting in better target coverage. Stereo vision is strongly affected by the lack of texture and the repetitive tiling on the checkerboard. Depth estimation is generally accurate near the edges and corners of tiles and poor in the middle. The ELAS algorithm automatically rejects these ambiguous areas, leaving accurate points, but with large holes in between. Results from the Kinect sensor are intriguing. While the density score is skewed due to fortuitous combination of narrow field of view and high density of the CCD, the pixel samples are not truly independent due to interpolation. However, the Kinect functions admirably as a low-cost volumetric mapper in the ideal case, greatly outperforming its pricepoint.

Analysis was also conducted in field environments to assess the impact of natural and artificial surfaces on each sensor. Each environment exhibits varying environmental materials, levels of

occlusion and surface conditions that collectively affect performance of the laser scanners. Examined here (as described in Section 4.2) are the Points-Per-Surface Patch (PPSP), Variance-Per-Surface Patch (VPSP) and Percent Surface Covered (PSC). PPSP looks at the average number of points residing in a surface patch over all surface patches considered (spanning approximately 20m from the sensor). VPSP is the averaged variance of a surface patch over all surface patches considered. Lastly, PSC is the average Percent Surface Coverage, a ratio of the number of patches with data over the total number of patches that are averaged across all scans in a particular environment (see Table 4).

Table 4. Environmental Surface Comparison

	Unstructured Corridor			Unstructured Intersection			Structured Corridor		
	PPSP*	VPSP	PSC	PPSP	VPSP	PSC	PPSP	VPSP	PSC
Photon80	9.97	9.0e-4	52.08	8.34	0.09	30.05	9.16	3.1e-3	39.39
LMS111	18.58	2.0e-3	29.16	18.98	1.5e-2	23.71	18.49	1.0e-3	36.35
LMS200	-	-	-	16.52	0.12	24.75	14.18	9.0e-3	36.44
LMS291	16.98	5.2e-3	24.52	16.7	3.0e-3	18.83	17.27	4.9e-3	22.85
LMS511	19.10	6.2e-2	21.99	-	-	-	18.35	5.7e-2	42.30
UTM30	-	-	-	-	-	-	18.89	9.4e-3	33.96
O3D201	1.87	1.0e-3	3.14	2.04	3e-3	3.09	-	-	-
Stereo	19.58	7.31	0.45	18.00	1.7e-2	6.27	-	-	-
S. Light	-	-	-	-	-	-	12.84	4.9e-3	1.65
Kinect	13.01	8.7e-3	15.10	19.64	0.04	5.57	16.64	9.0e-3	11.44

*PPSP in units of points, VPSP in units of meters², PSC in units of percent.

The LMS111 came closest to the baseline sensor (Photon80) across all metrics. In fact, it exhibits a marginally lower measurement variance (VPSP) score than the Photon80 in several cases. Possible explanations for this anomaly include the presence of moving average filtering internal to the LMS111 (independent of the single-shot mode utilized) resulting in over-smoothed surfaces or a non-uniform angular bias of samples in each patch due to rotational actuation. The Photon VPSP should be taken as the best estimator to the true surface roughness value. The conclusion reached here is also contradictory to the ideal analysis, which nominally placed the LMS511 as the best performing time-of-flight sensor in accuracy. Generally speaking, the specific ordering of these sensors is mostly within the intra-class variance of ToF technology seen in this study. However, as density and accuracy estimation are not entirely independent, this effect could stem from lower numerical stability when fitting patch planes to the less angularly dense LMS511 data. The other time-of-flight sensors follow the LMS111 closely, with the particular LMS291-S14 unit generating anomalously high PPSP score due to a maximum range

threshold of 8 meters. This threshold introduces a bias in regards to PPSP, which is a distribution-averaged value.

The PSC results arguably paint a more complete picture of density than PPSP. In this metric, the 8 meter range was a limitation for the LMS 291, while there was ample surface coverage by the LMS111 until approximately 20 meters. The Photon80 consistently captured the majority of the surfaces considered (with the maximum occurring in the unstructured corridor with approximately 52% of the available surface). Though the Photon80 is the baseline sensor, coverage values are not 100% due to holes and occlusions in the convex hull of measurement. In terms of relative surface coverage, the time-of-flight LIDARs achieved between 40-60% of the Photon80's coverage in the unstructured corridor, 60-75% coverage in the unstructured intersection and 75-100% in the structured corridor. These results correlate with the maximum possible sensing horizons (in decreasing order), where the photon performs particularly well in the lengthy unstructured corridor.

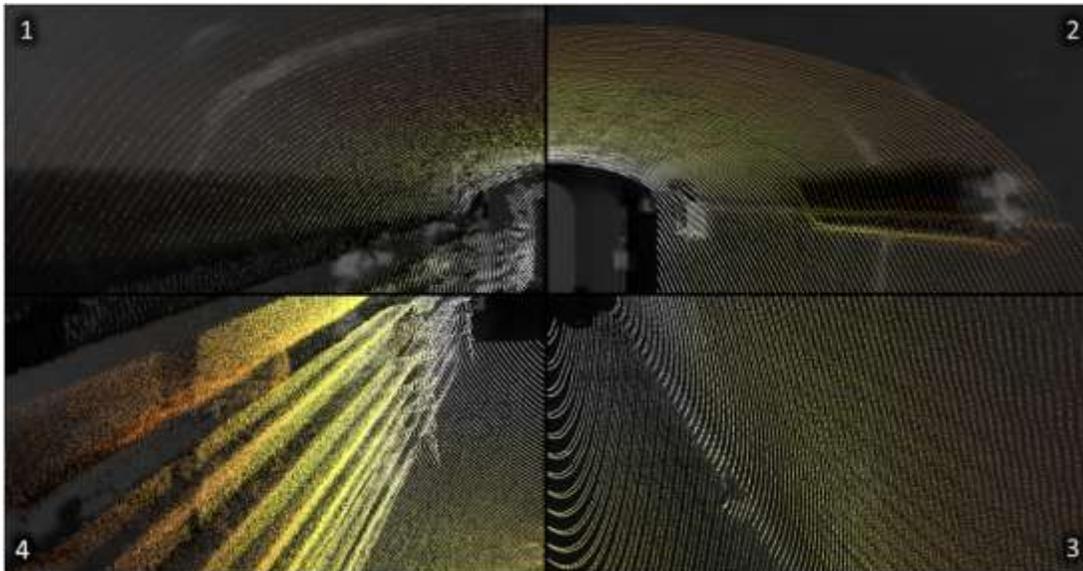


Figure 22. Rendered Faro Photon80 mesh models from the structured (top) and unstructured corridor environments (bottom) are shown in the left image. A montage of point clouds (right) illustrates differences between the (1) Hokuyo UTM-30LX (2) Sick LMS291 (3) Sick LMS200 and (4) Sick LMS111 sensors. The points are mapped onto a grayscale albedo map of the structured corridor generated using the Photon80.

As a group, the inconsistent performers measured particularly poorly in density and coverage metrics in absolute terms, but did not underperform significantly in the accuracy (VPSP) metric. Once again the Kinect outperformed expectation with an accuracy value of about 50% less than the ToF LIDARs and coverage between 33-50%. It should be noted that while LIDAR sensors were thoroughly represented in this study, results for triangulation sensors could be considered data deficient. A multitude of possible configurations exist in camera and lens selection and matching algorithms. Shrewd enhancements to the evaluated setups could shift the results of triangulation sensors closer to LIDAR or decrease the uncertainty ellipses.

Application to Mine Rescue Scout

We do not advocate a particular "best" sensor, as selection must consider the necessities of each

individual application. As detailed in the following Robot Configuration chapter, our Robot scout experimental platform was built on a repurposed mobile robot with existing sensors: a pair of rotating Sick LMS200s. Sensor characterization work showed this to be a solidly performing configuration in terms of mapping capability. Therefore, in the interest of cost and development time savings in this early proof-of-concept work, these were retained for the project. Applying a full selection matrix utilizing mass, power and ruggedness parameters along with mapping capability determined in this work, we believe that the Hokuyo UTM30-LX will be a particularly interesting candidate for future Robot scout development. It exhibits similar mapping performance to the LMS200, while being 1/10th the mass and power. While the UTM30-LX has a plastic housing (meaning it is less rugged than other sensors reviewed), no current commercial sensors are mine certified. Its small form factor makes it advantageous for conversion, such as sealing in a rugged explosion-proof box.

5. Robot Configuration



Figure 23. The Robot scout platform is configurable with a spectrum of sensors suited to each task: 1) radar configuration for smoke-filled response, 2) mounted with survey scanner for inspection modeling and 3) reconnaissance scout.

This project resulted in the development of the Robot scout platform as a proof of concept robotic rescue aid. It is capable of mounting a multitude of sensing configurations tailored to applications ranging from rescue to inspection (Figure 23).

Robot scout was realized by upgrading and repurposing an older mine mapping robot (Cavecrawler) developed at CMU with rescue capability. Repurposing CaveCrawler was determined as a low-risk, cost effective alternative to ground-up design in physical demonstration of project objective. The original chassis and much of the electronics were preserved, while computing, sensing and some of the external hard points for mounting were upgraded.

The Robot scout platform is not mine certified and lacks longevity for prime-time operations, it is intended to demonstrate the efficacy of robotic technologies and the rescue scout idea. We believe future ground-up development of a robotic chassis with the principles detailed in this study will result in a much more robust system. The following section details the features of the Robot scout robot.

Physical Configuration

Robot scout is an all wheel drive, body averaging robot capable of driving at a walking pace with the sensing and computing abilities to enable autonomous operations. This all electric system is powered by two 12 volt 100 amp hour batteries allowing driving operations up to ½ mile in mild terrain. The body averaging configuration allows Robot scout to take on aggressive terrain and obstacles however high centering can be a problem with ground clearance only 7". An on board inverter provides 120vac power to the 2 drive motors, 2 steering motors, and convenience outlets. Additionally the bus 12vdc and regulated 24vdc are available voltages to power sensors. Configured as Robot scout the robot is 42" wide x 62" long x 30" high weighing 450lbs. Reconfiguration for this project moved the primary LIDAR sensors up to accommodate mounting batteries horizontally. A structural framework and deck were added to the chassis to provide mounting locations for the sensing study. The primary PC-104 computer was supplemented with a second Atom computer to handle the user interface development. The existing 802.11 radio was upgraded with a new wireless access point offering 4-10 times more bandwidth. This node is capable of the 802.11s Wireless Mesh standard which is useful for future development of post accident establishment of communications.

Sensors

Robot scout features a baseline sensor suite of range, imaging and position sensors. This is the minimum set of sensors necessary for spatial mapping and for human awareness. This setup enables data from other sensors - such as gas, temperature and acoustic - to be easily geo-registered in the robot's map.

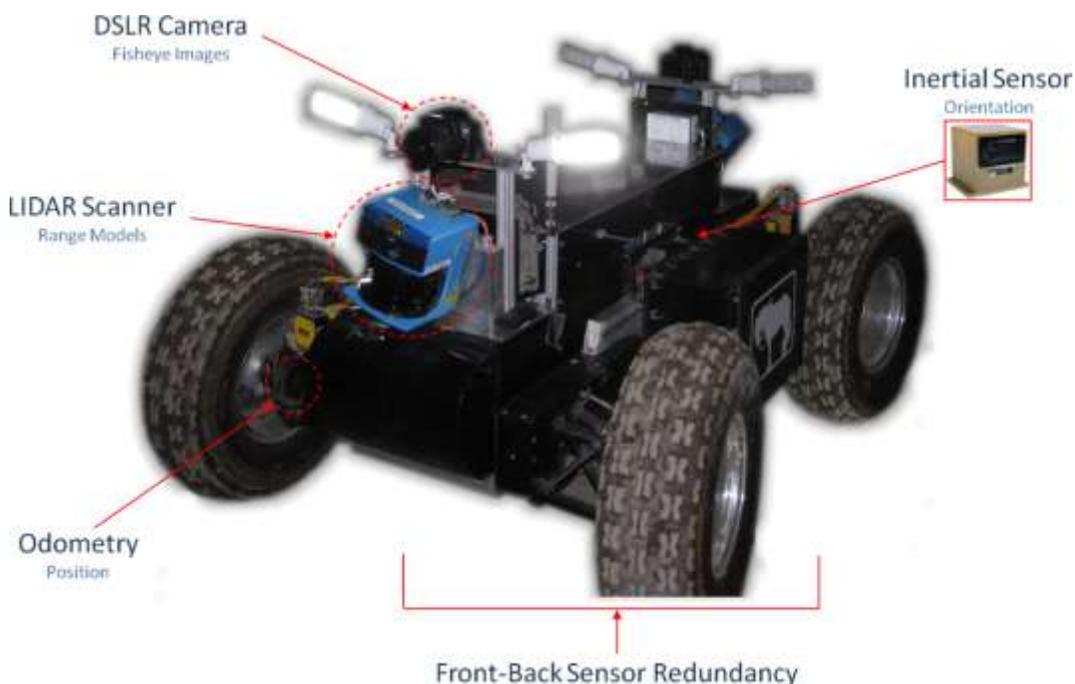


Figure 24. Sensors on Robot scout configured for rescue reconnaissance operations.

Robot scout features a front-back symmetric, redundant sensing configuration (see figure 24), which enables the robot to operate in narrow areas without turning and to reverse with the same fidelity as forging ahead. The redundancy also enables continuation of the mission, should any one sensor fail. Many of the sensors were preserved from the Cavecrawler platform during the upgrade to Robot scout, but with remounting to more advantageous view points.

3D Modeling

Robot scout's primary mapping sensors are dual spinning LMS200 ToF LIDAR scanners (*Spinner*) from Sick Ag. The design of these sensors is detailed in the prior work of our research group [Omohundro 2007], and the particular modules used are carry-overs from the Cavecrawler robot. During the project, the Spinners on Cavecrawler were moved up higher from their original very low viewpoints. It was determined that the low, glancing perspective provided short sensing horizon and poor data of the ground plane. Care was taken to preserve the total exterior dimensions and low-profile of the robot by swapping volume with other electronics.

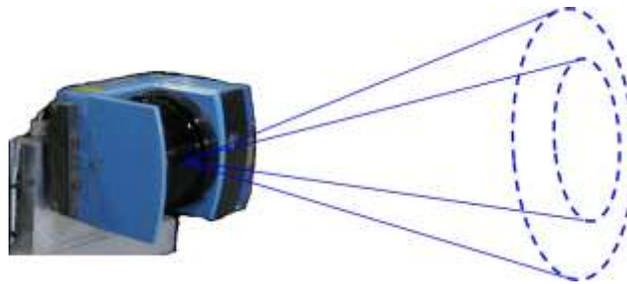


Figure 25. Configuration of the rotating Sick LMS200 mapping sensors

Much of the justification for utilizing a continuous rotary action for tunnel mapping (as opposed to a tilting action popular in surface robotics) is an advantageous measurement density distribution. Specifically, with rotary action, an undesirable density singularity occurs only down the infinite length of the tunnel (where there is little useful information). Whereas with a nodding configuration, two singularities occur on each side wall, and the walls and ceilings feature uneven distribution of data. Moreover, with the addition of slip rings, continuous rotation reduces power draw, motor fatigue, and results in fewer actuation anomalies in map data.



Figure 26. Mesh Models generated from Robot scout

Each of the Spinners provides a hemisphere of map data (180 x 360 degrees), and both of them put together provide spherical coverage, with a small gap resulting from the length of their displacement. Under normal operations (medium quality, medium speed map), the rate of rotation is approximately 30 RPM, resulting in 800,000 point measurements per minute.

Imaging

Most often, humans make decisions based on photographic information, not 3D point clouds, which are essential to the robot and useful to humans in a non-urgent manner. This meant rethinking the way rescuers would use information from the robot. The Robot scout imaging system was designed from the ground up to provide the best quality imagery and greatest situational awareness for rescuers. Consequently, a major objective for the design of the imaging system was to provide the clearest imagery for rescuers in the most adverse of conditions like dust, smoke and total darkness. Critically, utilizing the technology of today, this meant moving from 30 frame-per-second video capability to single-shot images at a lower rate. However, what is gained is a guarantee of no "bad images," images that would never be blurry, or too dark, too noisy, low resolution, or out of focus.



Figure 27. Configuration of Imaging Sensors on Robot scout

The imaging system on Robot scout is shown in Figure 27. It consists of a digital single-lens reflex (DSLR) camera with fisheye optics and programmable LED lights. The DSLR camera is controlled by the robot over USB using the Canon SDK software, which can remotely adjust image parameters (i.e. shutter speed) to suit the scene, and capture and download images. While the camera is capable of recording 24FPS HD video to onboard memory, it is not capable of streaming this to the robot (and thus the human user), due to the limited bandwidth of the USB connection.

The lens is a full-frame 180 degree fisheye, meaning that the viewable area takes up the whole image (some fisheyes, for example circular ones, have wasted areas where there is no image information) and that the diagonal field of view is 180 degrees (the horizontal FOV is about 160 degrees). Fisheyes are capable of imaging such wide angles because they distort (squeeze) areas

of the image to accommodate peripheral data. Different fisheye optics result in different areas of the image being distorted. The particular Rokinon 8mm lens used on Robot scout uses a stereographic projection, which maximizes peripheral resolution at the cost of central resolution (where it is often pitch black in tunnels).

Illumination is the third and oft-ignored component of an imaging system. The Robot scout system features two rectangular LED arrays on each side of the camera, separated by 8 inches. LED sources are highly efficient, meaning they consume minimal power and generate minimal heat for their brightness. These particular arrays have the equivalent output of a 40W incandescent bulb while consuming 7W of power. Unlike florescent lights -which are also highly efficient - LEDs do not flicker at 60Hz (which creates imaging problems) and ramp quickly to full brightness such that they can be utilized as flashes. The separation of the lights to the side of the camera provides for uniform lighting (no dark self-shadowed areas) and reduces blinding in dusty environments (see below).

There are many reasons for using a custom DSLR, lens and illumination combination instead of an integrated commercial "low-light" CCTV camera (popular on many underground robots). CCTV cameras often feature plastic "wide angle" lenses (usually 90-120 degree), they do not compare to DSLR fisheyes in angular coverage or optical performance. Figure 28 below illustrates the benefit of fisheye optics for rescuers who need to quickly assess situation and look for important information like other humans and obstacles. In addition to enhanced field of view, fisheyes have large depths of field (pictures are always in focus) and large apertures (collect more light) than rectilinear lenses.

CCTV cameras, even those marked as low light capable, often result in underexposed images with large amounts of noise. This is due to insufficient CCD sensor area, which itself is a result of small packaging and low cost. Modern DSLR cameras have over 10 times the sensor area as CCTV cameras, translating to much better image quality, resolution, and most importantly, low light performance. The cameras utilized on Robot scout shoot 10 megapixel images and are capable of 12 bits of dynamic range over three color channels. Most low light CCTV cameras are monochromatic, feature 640x480 equivalent resolution, and are only capable of discriminating less than 8 bits of brightness (due to analog conversion). Famously, borehole images shot during the Crandall Canyon disaster using CCTV cameras resulted in images that were either almost totally dark or blown out due to inadequate dynamic range!



Figure 28. Images from a standard 90 degree "wide angle" lens (left) versus a 180 degree fish eye lens (right) illustrate the enhanced situational awareness in fish eye optics.

Lastly, integrated LED rings on off-the-shelf cameras are notorious for creating terrible images. The close proximity of the lights to the lens creates two problems. Firstly, the camera is easily blinded by any dust or particulates in the air by backscattering (see also Mie scattering). This is the reason why fog lights on vehicles are mounted low and away from the driver's field of view. Secondly, stray light at the fringes of lenses create undesirable phenomenon such as lens flares and ghosting (multiple reflections of the scene). While high-quality lenses use lens hoods designed to mitigate glancing rays, CCTV lenses are often simply embedded in flat plates.

The caveat to using commercial photography grade equipment is that it is usually not packaged for robust field operation. We recognize the need for a roadmap to operations from the proof of concept system proposed. Thus, the use of industrial-grade machine vision (MV) cameras is recommended for future designs (these were not selected in this program due to budget constraints). MV cameras feature the imaging quality and lens selection of SLRs, but are packaged minimally for industrial use. Since they communicate over gigabit Ethernet, it is also possible to stream full-motion video in real time while retaining high quality.

Position and Orientation

Robot scout has encoders on the drive motors that provide feedback on the wheel rotations and the distance traveled. Integrated in the system is a Crossbow Inertial Measurement Unit (IMU) model 400CC, a high performance solid-state six degree-of-freedom (6DOF) Inertial Package intended for OEM navigation and control applications. This high reliability strap-down inertial system provides accurate measurement of angular rate and linear acceleration. With the encoder and IMU Robot scout can provide an estimate of position heading and roll, pitch, and yaw orientation data which is logged continuously.

Environmental Sensors

Robot scout comes equipped with a digital voltage and ambient temperature readout. Serial USB and RS-232 interfaces exist to integrate other sensing. CaveCrawler has previously supported monitoring readings from commercial gas units. During the course of this project we did not integrate a gas sensor, anemometer or other environmental sensors. Multi gas monitors from Industrial Scientific and Mine Safety Appliances were investigated and would easily interface with our system. NIOSH has recent research published using solid state anemometers in the

mine. Anemometer data from Robot scout would provide insight into air flow during rescue response. The sensitivity of these solid state units can identify very low velocities of air flow and flow direction. Additionally RFID sensors and microphones were identified sensors that could supplement the system. Mines that use RFID systems to tag and locate equipment and personnel would be read and identified by the Robot scout then relayed with time and position stamping to rescue personnel. Microphones deployed from the Robot scout could listen for sounds which would can be recorded and made available for playback. Rescuer's in close proximity could use the microphone to log verbal descriptions at a location, algorithms could be developed to tag and record unnatural sounds with position and time stamp for rescuer review.

6. Analysis & Demonstration

The intent of this program was to demonstrate the capabilities of a mobile robot scout to perform tasks specific to mine rescue, mine safety, and supply data for next-generation training tools. A final demonstration marked the culmination of program activities, which occurred at the Bruceton Research Coal Mine on September 15, 2011. During this demonstration a mobile robot research platform (called Cave Crawler) autonomously traversed, mapped and modeled sections of a mine live in front of project sponsors and mine officials. Specifically, this demonstration showcased a control interface, sensing, autonomous robot navigation, data & information, and communications. Three scenarios captured key program objectives, which included:

- *Robot-assisted mine rescue where a mobile robot will forge ahead of a mine rescue team to relay mine conditions.*
- *Robotic inspection where a mobile robot monitored a specific section of mine for abnormal conditions*
- *Robotic modeling where a mobile robot will survey and document an "accident scene" and build a mine model for training and virtual reality*

This section details each of the scenarios and describes the relevance to mine operations.

Robot-assisted mine rescue

Robot-assisted mine rescue allows mine rescue teams to more effectively gather information about mine conditions without being in harm's way. This has the advantage of allowing rescue teams to make better decisions at a faster pace and thereby increase the likelihood of a desirable outcome.

The concept is to augment teams of mine rescuers with mobile robot scouts that can forge ahead of a rescue team and relay mine conditions. The robotic scout operates autonomously with high-level commands from a rescue coordinator. In this manner, the rescue team can operate unencumbered from remote controls that are tedious to operate in rescue situations.

For the demonstration, the robotic scout, Cave Crawler, was tasked to autonomously explore a section of mine, relay conditions and summarize findings with in an interactive map. Starting just outside of the mine entrance, Cave Crawler voyaged up the main entry, turning right at the first cross cut and made its way to a far man door that prohibited further exploration. While traversing

this path, it relayed snapshots of mine conditions to an interface over wireless communication. Results from this exploration are captured in Figure 29.



Figure 29. Map and Photographs of Mine Exploration

The map shows the route Cave Crawler took to reach the man door. The map is a 2.5D presentation of the mine; meaning, 3D elements of the mine (e.g. barrels, columns, tracks, etc.) are visible. Photographs are linked to map positions to provide a way to “click” on the map and see data relayed to the interface during exploration. In this particular figure, barrels and the man door are shown, which represent unique features located at the beginning and end of the robot’s mission, respectively.

Robotic inspection

Robotic inspection has applications related to mine operation as a means to prevent mine accidents. The advantage of robotic inspection is in its ability to provide consistent, repetitive, and thorough data collection in specific areas of a mine. In addition, the robot’s localization and mapping system provides the means to geo-spatially tag data to very accurate locations in a mine.

The concept presented in this demonstration is to place a robot into a mine and have it repeatedly map and take environmental measurements. These measurements can then be placed into a database and queried for reports or organized into visualization layers on top of a mine map.

For the demonstration, the robotic scout, Cave Crawler, was tasked to autonomously traverse a section of mine and document conditions with findings summarized in an interactive map. Starting at the man door, Cave Crawler traversed along the corridor to a specified destination and

returned. During its return, participants at the demonstration were asked to become part of the mine environment. As such, a group consisting of nine individuals dispersed into the robot's inspection route and became part of the scene. The resulting data sets are show in Figure 31, and Figure .

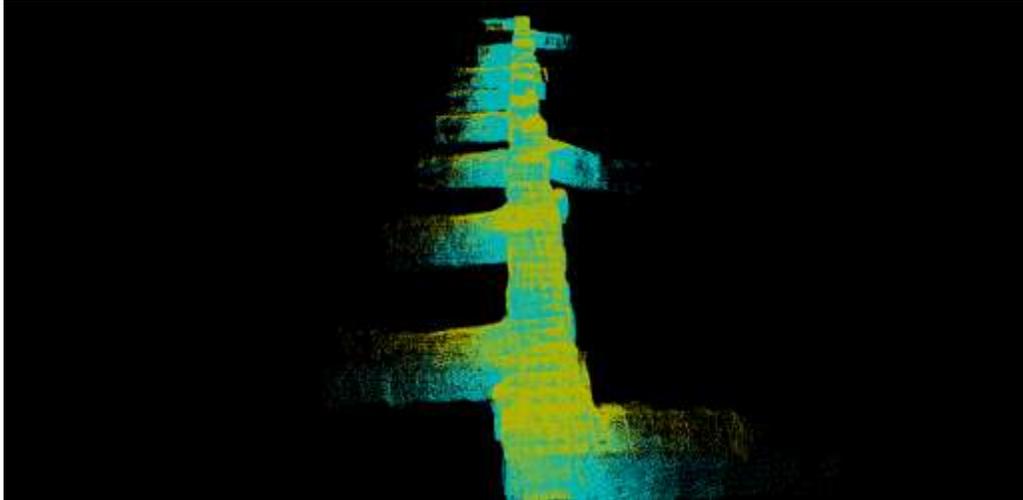


Figure 30. 3D Mine Model. The different colors represent the source laser that produced the model data.

During each traverse, 3D laser data was fused into a 3D model of the mine surface. The first traverse provided the reference frame, or “base map” that would be compared to data collected during the second pass. Differences between the two maps where represented as red highlights as show in Figure 32.

In addition to the data collected at the demonstration, a number of data sets were collected over the course of the project. Within this collection are data sets with closed loops, which create additional challenges for modeling algorithms, and very dense traverses that produced high-resolutions maps. Examples of these data sets are show in Figure 31.

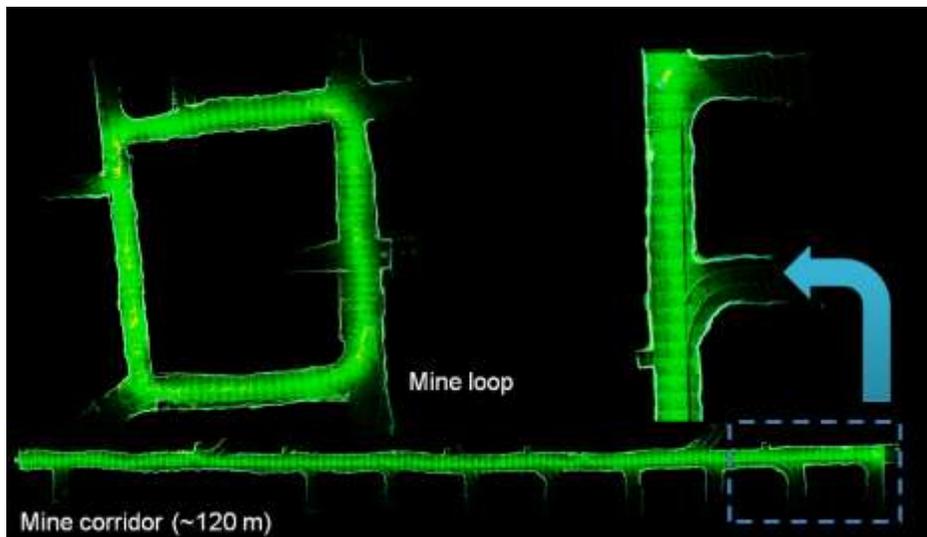


Figure 31. Other mine maps and models collected during project.

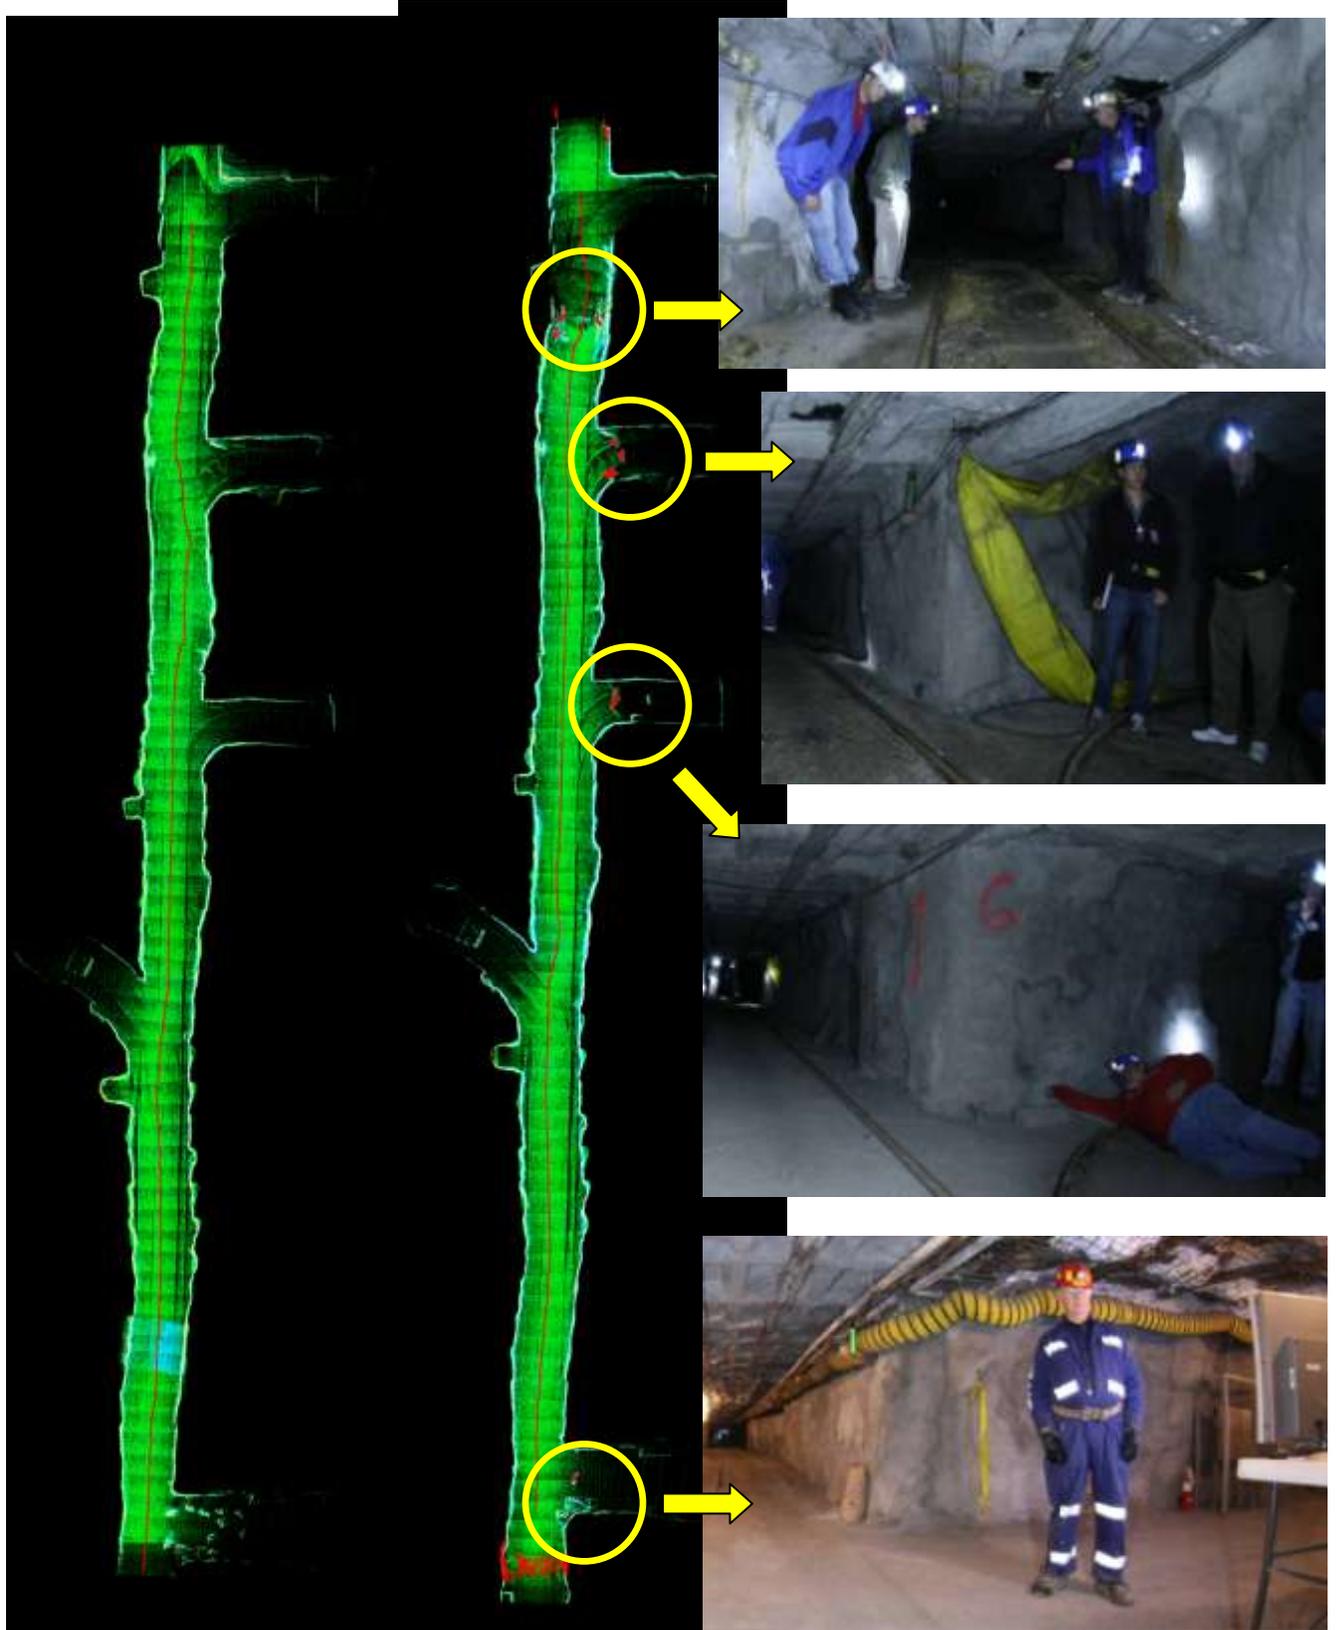


Figure 32. Model comparison.

The images shown in Figure 33 capture the final demonstration result for robotic inspection. The image on the far left is the base map generated as Cave Crawler drove the corridor from the bottom of the image to the top (path is indicated by the red line). The middle image shows changes detected from the first pass to the second in the form of red markers. The images on the far right are photographs recorded by Cave Crawler, which correspond with the detected “changes.”

Robotic modeling for “accident scene” documentation

Robotic modeling also has applications related to documentation, specifically following accidents. The sensors employed by robots are meticulous in their ability to capture and archive even the minutest details of a scene. Combined with the capabilities such as difference tracking, the robotic modeling provides an unrivaled means to document an environment.

The concept presented in this project was to show how quickly and effectively a robot could capture details of a scene. Going at the walking pace of a person, a mobile robot can reconstruct the entire 3D geometry of a mine, capture video and photographs, and be equipped to support any additional sensor of interest.

For the demonstration, Cave Crawler created 3D models in near real-time. The models shown in Figure 31 and Figure 32 were all captured, processed and displayed during the demonstration, which spanned approximately 1.5 hours. Furthermore, results shown in Figure 35 and Figure 36, illustrate the detail that can be capture from a continuously moving robotic scout.

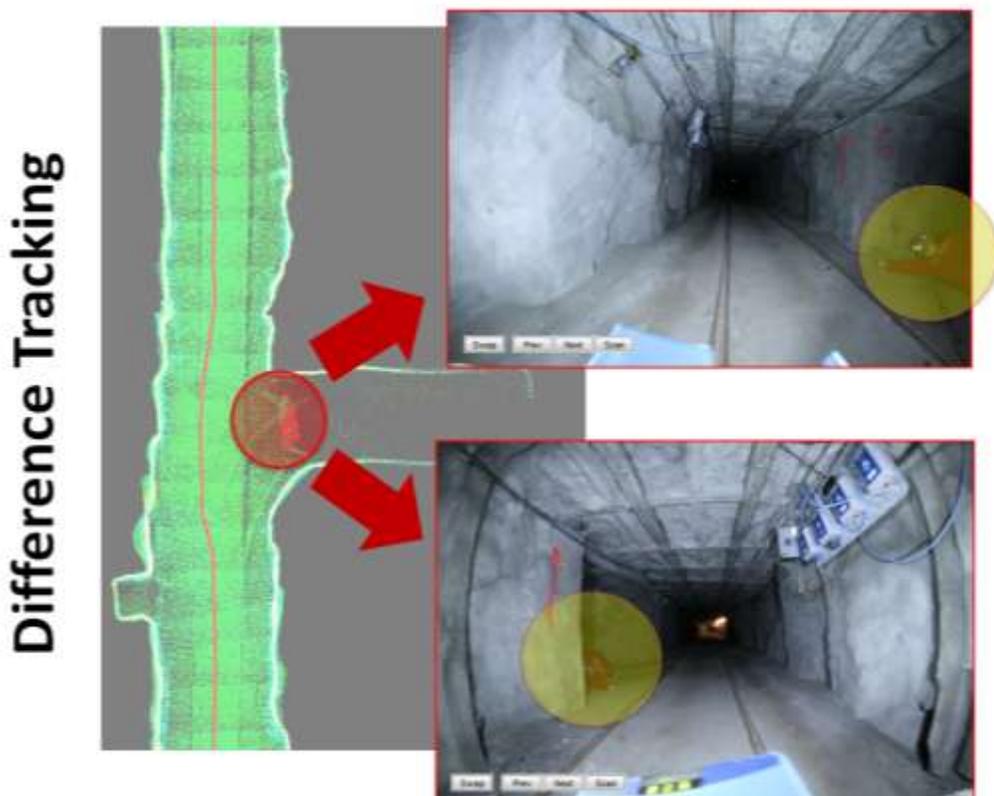


Figure 33. Example difference tracking for accident scene documentation.

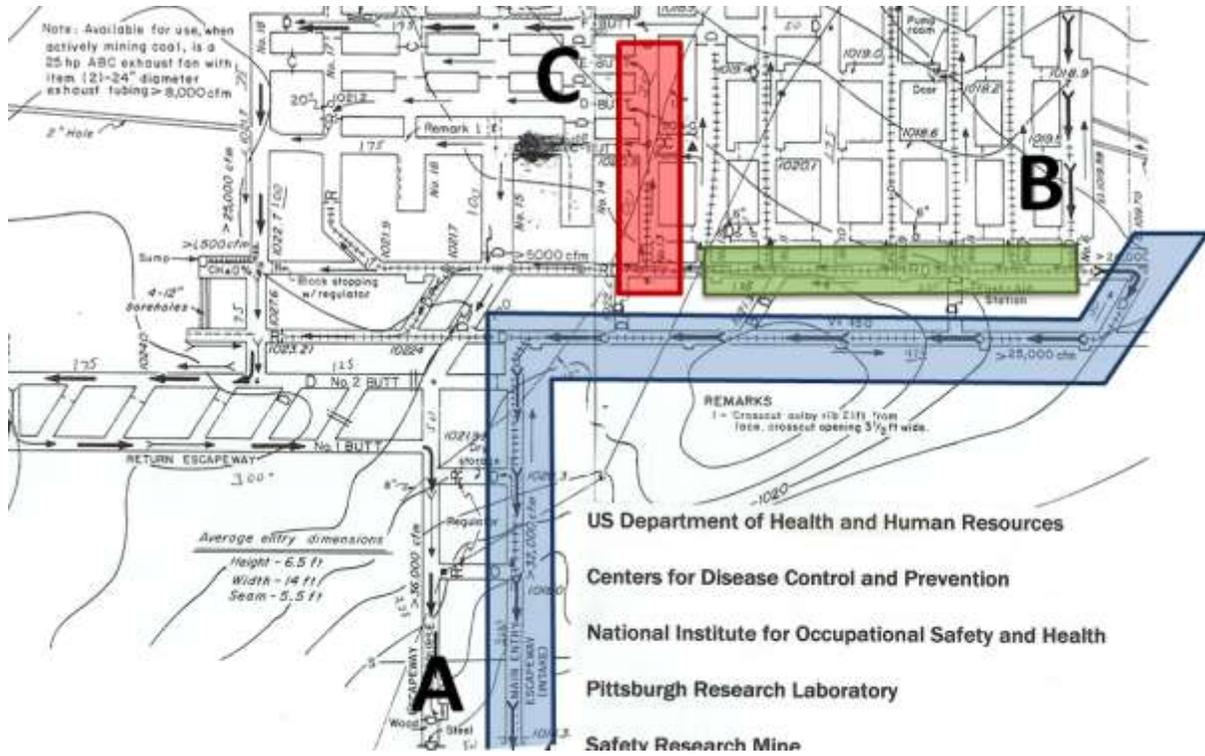
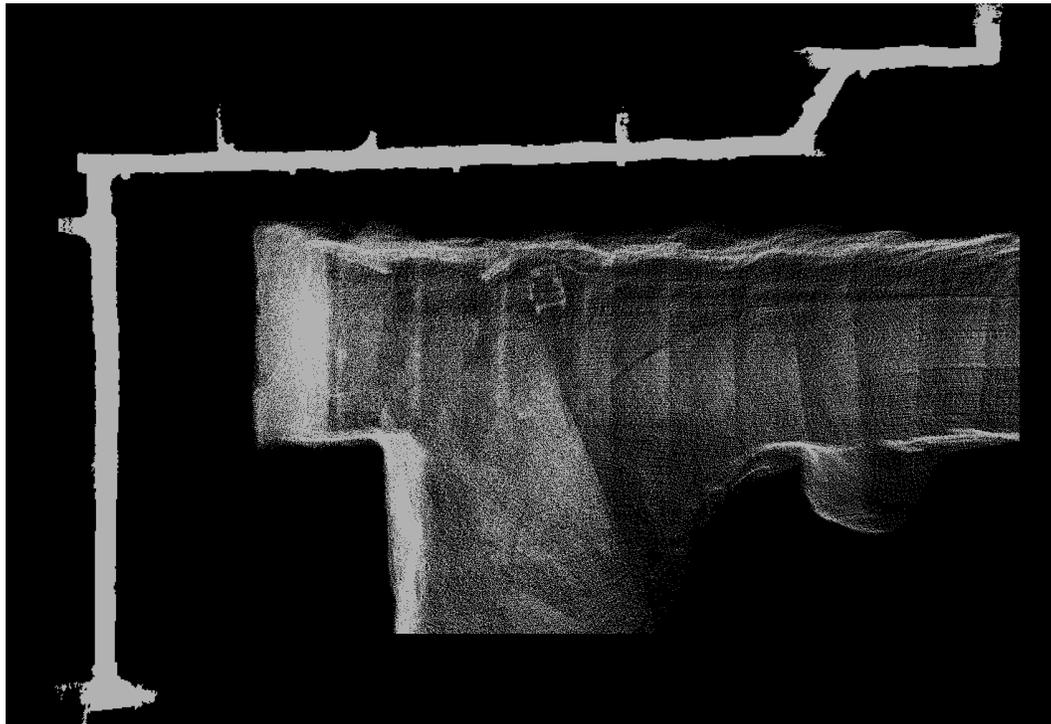


Figure 34. Corridor segments mapped and modeled during project.



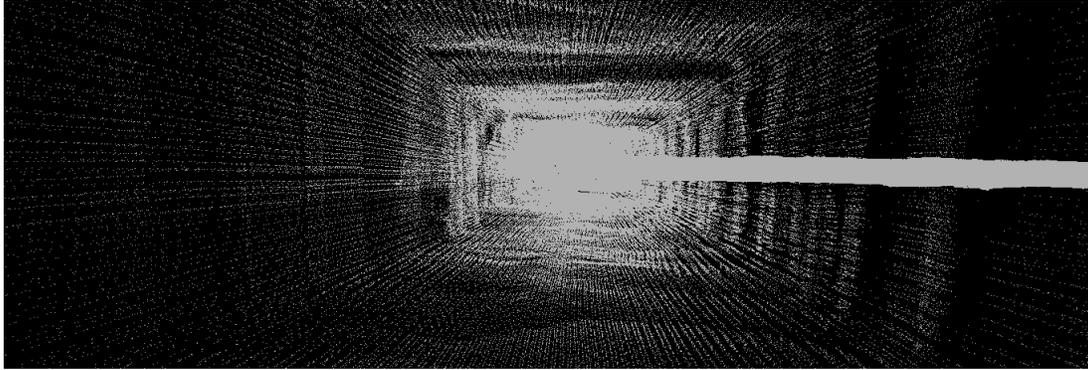


Figure 35. Model of Traverse in Segment A (Mine Entry)

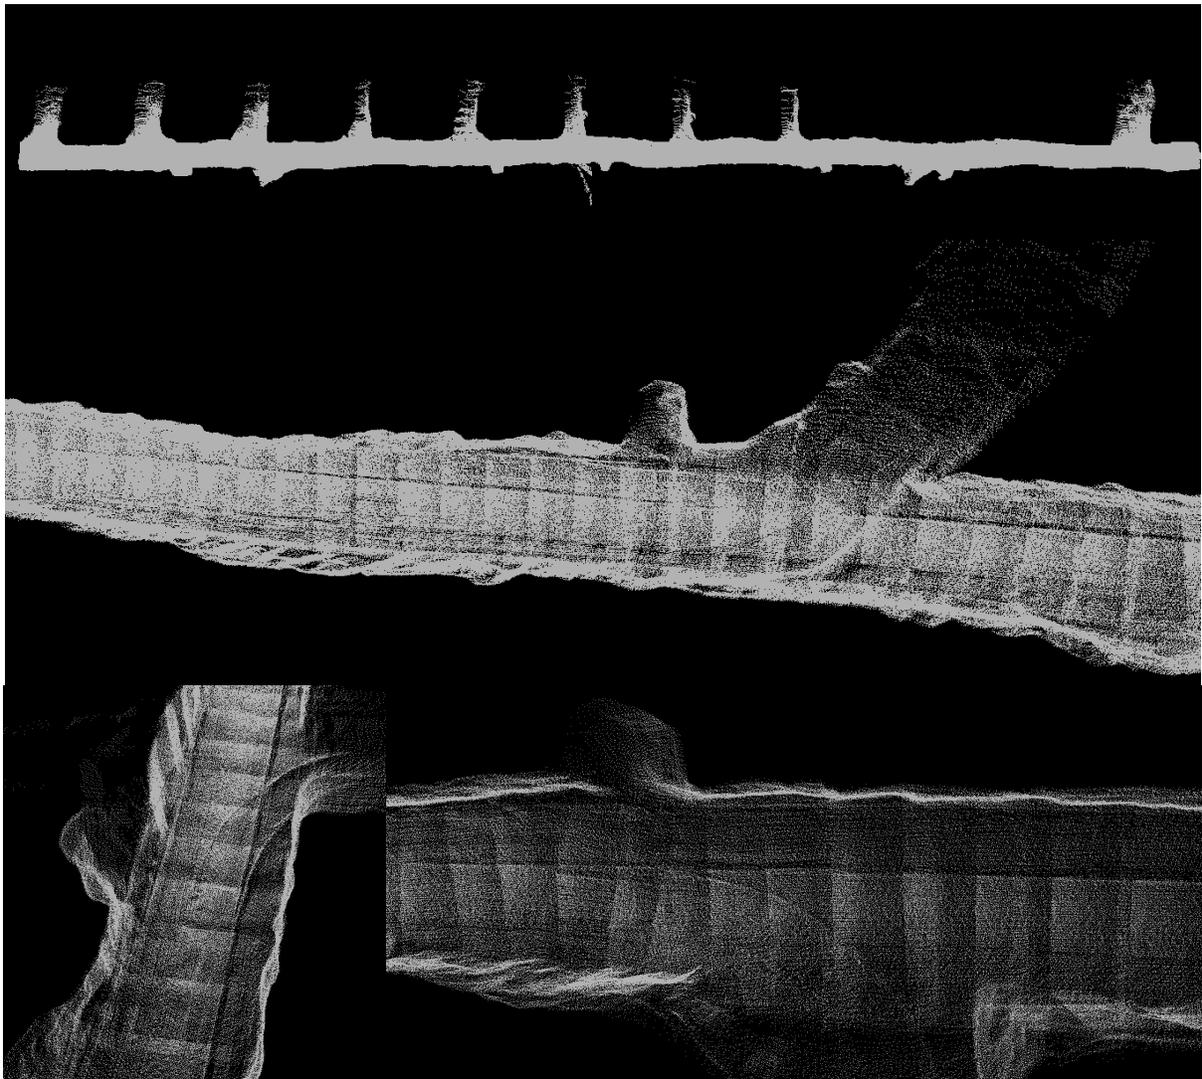


Figure 36. Segment B Model

7. Conclusion and Future

The research work for this project has produced the following results:

- Defined rescue tasks for a robot scout.
- Implemented a user interface to control the robot scout,
- Refined methodologies for evaluating modeling sensors for mines.
- Evaluated several classes of sensors.
- Integrated the robot scout system onto a CMU robot.
- Using the robot and user interface demonstrated in full the objectives for the robot stated at the onset of the project:
 - *Robot-assisted mine rescue where a mobile robot forged ahead of a mine rescue team to relay mine conditions.*
 - *Robotic inspection where a mobile robot monitored a specific section of mine for abnormal conditions*
 - *Robotic modeling where a mobile robot surveyed and documented an “accident scene” and built a mine model for training and virtual reality*

The Mobile Robot Scout future

The recommended next step required in order to advance the system is to set up a long term (3 month or more) program where the Robot scout technology routinely maps a designated mine. This will provide testing and assessment for a pre-production system. This will also provide the opportunity for increased mine industry awareness and visibility of the technology. Mining personnel associated with rescue need to have a chance to gain exposure to the system in this setting and provide feedback on the implementation.

Miner training simulator developers need realistic models gathered from mines to make simulators closer to virtual reality. The modeling data from the robot scout contains the information to produce the virtual reality model. The goal is to establish a methodology to quickly update training simulators with fresh models from the field. Ultimately the goal is to train miners in a virtual mine environment that represents the mine where they will be working.

Mine rescuers need to have a second generation Robot Scout to integrate into mine rescue training. There needs to be a program that rigorously tests the Robot Scout functionality in performing the mine rescue tasks outlined in this paper.

Mine models produced over time from the same mine need to be statistically analyzed to determine if subtle changes within the mine are indicators of developing hazards. During this phase both government and industry need to look at the Robot Scout technology and identify ways that the system could be used for inspection, operations, as we move toward the automated mines of the future.

8. Acknowledgements:

Carnegie Mellon University and our project team wish to thank Wheeling Jesuit University for this opportunity to support the MISTTI project. We especially thank Davitt McAteer and his project team for their support throughout the project.

We also wish to thank:

- National Institute of Occupational Safety and Health, NIOSH, for granting Carnegie Mellon access to the Pittsburgh Lab, Research Mine at Bruceton, PA. and to the Mine Foreman Paul Stefko, and miners Joe Sabo, and Jack Teatino for supporting our field work and demonstration.
- Mine Safety and Health Administration, MSHA, personnel for their input and in particular the contributions of John Urosek and his MSHA robotic team, plus Virgil Brown with members of his mine response team
- Danny Spratt, Mine Rescue Coordinator, Office of Miners' Health, Safety, and Training State of West Virginia for his insight into mine rescue teams and operations
- Mine Technology and Training Center, Prosperity, PA for the opportunity to observe and gather information on mine rescue during the fall competition.
- RedZone Robotics, Inc. and their field team for information and observation of field operations with their robots for the sewer industry.
- Mine Rescue User Interface students Tim Andrianoff, Nisha Kurani, Andrew Lee, Ward Penney, and Justine Yang from the Carnegie Mellon University, Human Computer Interaction Institute course, HCI Methods (05-610 :: A), taught by Dr. Bonnie John and Dr. Matt Kam for their contributions.
- CMU project staff including Abhinav Valada



Figure 37 Image taken by Robot Mine Scout during the final demonstration at the NIOSH research mine in Bruceton, PA on Sept. 15, 2011

9. References

- [1] **[Omohundro 2007]** Z. Omohundro. Robot Configuration for Subterranean Modeling. PhD Dissertation. Robotics Institute, Carnegie Mellon University, 2007.
- [2] **[Morris 2007]** A. Morris. Robotic Introspection for Exploration and Mapping of Subterranean Environments. PhD Dissertation. Robotics Institute, Carnegie Mellon University, 2007.
- [3] **[Wong, et al. 2011]** U. Wong, A. Morris, C. Lea, J. Lee, C. Whittaker, B. Garney, W. Whittaker. Comparative Evaluation of Range Sensing Technologies for Underground Void Modeling. Proc. International Conference on Intelligent Robotics and Systems, 2011.
- [4] **[Orr, et al. 2009]** T.J. Orr, L.G. Mallet, K.A. Margolis] Enhanced Fire Escape Training for Mine Workers Using Virtual Reality Simulation NIOSHTIC-2 No. 20036395, 2009.
- [5] **[The National Academies Press 2007]** Mining Safety and Health Research at NIOSH: Reviews of Research Programs of the National Institute for Occupational Safety and Health” Committee to Review the NIOSH Mining Safety and Health Research Program, Committee on Earth Resources, National Research Council, ISBN 9780309103428, 2007.

10. Appendix

Electronic Files

This report is accompanied with electronic files from the project which includes:

Final Report

- This final report

IROS Paper

- Intelligent Robots and Systems conference, IROS, paper

Final Demonstration

- Movies and images
- Power point presentation
- Task 3.4 demonstration handout

HCI User Interface Class

- Power point presentation
- Web research materials

Sensor Characterization

- Sensor Specifications from Manufacturers
- Calibration data
- Laboratory Data

- Environmental Data
- Integrated Experiments

Data Files

Movie: Making 2-D map. Shows the construction process behind the making of a 2.5-D Map, which includes

1. Construction of a 2-D map
2. Optimization of the robots path
3. The “filling” in of 3-D data
4. 2-1/2 D map (height differential top removed)

Movie: Map Interaction. Shows concept of an interactive map with

- Photo display upper right
- The 2.5D map on top left 2
- A navigation window on the left

Data: The root folder has a series of data collection deployments defined as <date>_<location>, such as 2011.05.26_Bruceton. These document the time and place for each data set accrued during this project.

Inside each deployment folder, s a series of processed deployments are labeled as

`<process date>-ccproj[0|1]_<num>`

where

- “process date” is the day the deployment was processed
- [0|1] is 0 if the reference laser is the front or 1 if the reference laser is the rear
- num is the deployment identifier

Each processed deployment has the raw log data and an "Output" directory. The "Output" directory contains the processed results that include:

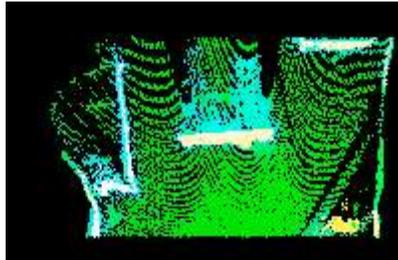
- map (folder) that contains a 2.5D map of the deployment
- report (folder) that contains a plan view (mmap.png) with distance, the plan view (planView.pts) in <x,y> (meters), and SUMMARY.txt that states the time and distance traveled
- vrmf (folder) that contain a 3D point model. NOTE: not all deployments (especially the large ones) have these models

The raw data is kept for repository completeness and, for example, will look like

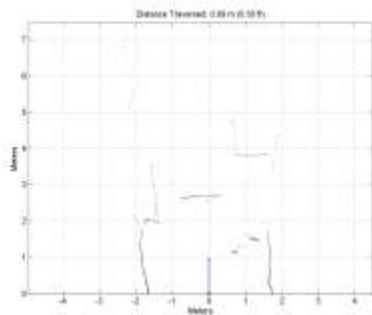
- 2011.09.15-ccproj0004 – Main project directory
 - async-0004.log – robot internal communication log
 - cc-command-0004.cmd – log of commands issued to robot

- cc-frontlaser-0004.sls – raw data for front laser
- cc-inertial-0004.log – raw data for inertial sensor
- cc-joystick-0004.log – raw data for joystick
- cc-odometry-0004.odo – raw data for odometry
- cc-pc-0004.log – internal process log
- cc-rearlaser-0004.sls – raw data log for rear laser
- cc-status-0004.log – status log for robot
- Output - output directory for results

- Map – contains the 2.5 map such as



- Report – contains results such as floor plan map



- Vrmf – contains the 3D point cloud data in vrmf format (not for every project). This data can be viewed and manipulated by software such as Meshlab:
<http://meshlab.sourceforge.net/>