Mining Robotics Sensors

Perception Sensors on a Mine Safety Platform

Green JJ[1], Hlophe K[2], Dickens J[3], Teleka R[4], Mathew Price[5]

[1,2,3,4]CMI, CSIR

P.O.Box 91230, Auckland Park, 2006, Johannesburg, South Africa

jgreen@csir.co.za[1], khlophe@csir.co.za[2] ,jdickens@csir.co.za[3] rteleka@csir.co.za[4]

[5]Cogency

4 Swift Street, Salt River, 7925, Cape Town, South Africa

mathew@cogency.co.za

Underground mining robotics has not enjoyed the same technology advances as above ground mining. This paper examines sensing technologies that could enable the development of underground autonomous vehicles. Specifically, we explore a combination of three-dimensional cameras (SR 4000 and XBOX Kinect) and a thermal imaging sensor (FLIR A300) in order to create 3d thermal models of narrow mining stopes. This information can be used in determining the risk of rockfall in an underground mine, which is a major causes of fatalities in underground narrow reef mining. Data are gathered and processed from multiple underground mine sources, and techniques such as surfel modeling and synthetic view generation are explored towards creating visualisations of the data that could be used by miners to monitor areas of risk in the stope. Further work will determine this potential.

Keywords: Underground Mining Robotics, Perception sensors, Sensor Fusion, Infrared Camera,

1. INTRODUCTION

To date, robotics in the mining industry has seen a lot of advancement in automation for above ground applications where positioning can be achieved with GPS, and often enhanced by a combination of DGPS and machine vision techniques.

Underground mining however, has not seen the same advancement due in part to the harsh conditions and inherent difficulties in navigating a robot through a rough, lightless environment. The lack of infrastructure and communication channels across the mine has also hampered the development of autonomous stope systems. Some progress has been made in tramming activities in tunnels with the application of repetitive wall following from load to unload points [1] but no level of autonomy has yet been achieved.

This paper looks at attempts to address the significant challenge areas involved in developing a mine safety robot that will enhance mine safety in the stopes of South African Hard Rock mines [2]. A number of critical developments are required for enabling a robot to sense and navigate in the harsh underground stope environment [3]. We discuss some of the issues and show the progress made to date.

The current objective of the work is to identify the technical risk areas that are barriers to the implementation of underground robots. We focus on a localization solution based on differential time-of-flight (dTOF) beacons, a combined multiple sensor system for

visualization in confined, lightless environments, and thermography for assessing the safety and stability of hanging walls.

Over the last decade approximately 200 miners have lost their lives per year in South Africa in underground mining activities (128 in 2010) [4]- 50% of those deaths occurred in the stopes, and of those, 70% in rock fall related incidents. It is envisaged that a robot could be used to enhance mine safety by gathering data in the stope (e.g. after a blast) and warn miners of potentially unsafe areas, that must be avoided, prior to their re-entry. Furthermore, the collected data can be used to focus making safe activities on unstable hanging walls.

Data will be gathered by the robot using, inter alia, a long-wave infrared (thermal) camera, a short-wave infrared time-of-flight camera, a 3D imaging device, and a sounding device for assessing hanging wall (roof) stability. It will then be combined and processed before being presented to the miner prior to entering the stope. As the eventual deployment of this technology will be on a mobile platform, the goal of realising real-time processing for 3D stope mapping must be kept in mind. While our current implementations process data off-line, the real-time requirement was considered and continues to guide the direction of development.

Section 2 covers the localisation techniques for determining where the data is collected. Sections 3 and 4 then cover the thermal and 3D sensors used before section 5 discusses the combination of data sets and potential visualisation techniques that could make the information about risk available to the miners.

2. LOCALISATION

For localizing the robot underground, which is a major challenge, CSIR are pursuing the implementation of a difference in time-of-flight beacon system [5]. Beacons are surveyed into the underground environment and simultaneously transmit ultrasonic and RF signals. The differential time-of-flight between the ultrasonic signal and the RF signal is used to estimate the distance between the beacon (Figure 1) and the receiver [6].

Figure 1: Multi directional dTOF Beacon.

Resulting positions will be further refined with on-board inertial sensors using an unscented Kalman Filter (UKF). This will allow the positional uncertainty of the robot to be bounded and will improve registration of sensor data for real-time navigation and mapping with large data sets and multiple sensors.

3. THERMAL IMAGING

Electronic perception underground is hampered by adverse conditions, especially in the deep-level hard rock mines of South Africa where humidity reaches saturation, temperatures are in excess of 40°C, and the threat of explosion requires intrinsically safe construction of equipment. There is also no ambient light available, so any illumination must be carried on-board. To complicate matters, the vast distances underground require tetherless operation and all power must be supplied by battery. Finally, the presence of dust hampers the use of optical techniques, and the abundant use of water to combat the dust creates extreme conditions that are not conducive to the operation of most machinery.

Infrared sensing poses a potential solution to some of these problems. Infrared light can penetrate dust very well due to its long wavelength. If the wavelength of an electromagnetic wave is larger than the diameter of an obscuring particle then the wave will tend to pass through it instead of being reflected or scattered. The long wavelength (7 – 14 µm) of thermal infrared allows it to penetrate airborne dust having smaller particle sizes.[7]. All objects radiate infrared at a wavelength dependant on their temperature, As the rock is already heated, the thermal emissions can be used to view the surroundings which negates the need for a standalone illumination system. Past Research [8] and [9-11] have explored the use of thermal imaging cameras in mining environments for the analysis of hanging wall stability.

This is made possible by the use of ventilation air which is a vital component in most active South African hard rock mines. Cooled ventilation air is used to keep the stope environment within acceptable limits as the rock can reach temperatures in excess of 65°C. This cools the rock surface, creating a thermal gradient through the host rock. If a crack exists, the heat flow is interrupted and the surface rock will cool preferentially. Analysis of the thermal gradient can therefore be used to identify potentially loose rocks that pose a danger to mine workers. The same technology is applicable for mapping stope areas in order to create a virtual mine that could be used by a small machine to make operational decisions. Three experimental data sets of increasing size and complexity are discussed below.

3.1. Initial data set.

A Wuhan Guide MobIR M2 hand-held thermal imaging camera was used for initial work. The field of view was limited for the application of inspection of a hanging wall in a 1m stope. Since the resolution of the camera is low (120x120 pixels) the resulting images do not provide sufficient coverage for robust detection of thermal features (associated with loose rock). Therefore, multiple overlapping images of an area of interest were collected and an automatic image stitching method was used. Figure 2 shows the first data set collected at Driefontein mine at a depth of approximately 3500m below surface.

Stitching is achieved by extracting invariant features [12] for each image and finding a series of 2D homographies (a 3x3 mapping between 3D planes) for pairs of images. Even though the camera is translated during capture, a 2D homography was found to be sufficient for stitching due to the approximately planar nature of the rock surface.

Unlike visible-light images in indoor scenes, invariant features are very sparse in the thermal images. This motivated us to use a simplified camera model based on only 4

parameters: 2D translation, scale, and rotation about the optical axis. The approximation worked well, because the images were captured by moving the camera in a parallel manner over the surface.
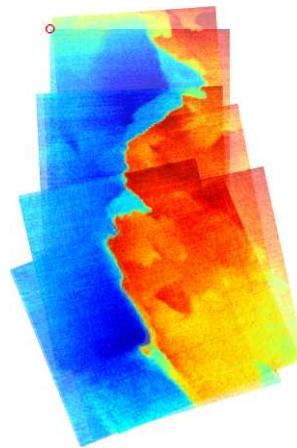
*Figure 2. Stitched thermal images from first dataset (Driefontein Mine).*

The experiment clearly showed how the ability to stitch thermal images together to create a larger data set could be used to aid analysis. Specifically, it allowed for the determination of a temperature gradient across areas of rock much larger than the field of view of the sensor. It was noted with this first data set that the sharp temperature gradient visible coincided with a step in the hanging wall profile, and not necessarily with an unstable area of rock mass. A larger data set was subsequently captured.

3.2. Data set 2.

A second data set was captured with more overlap between images and more consistent camera motion (in accordance with our simplified model). Although the sequence was captured in an ordered manner, this was not used to determine the pairing relationships during stitching. Rather, robust feature matching is used to find matches between all possible image pairs. The stitched image is then generated by iteratively choosing the image pair with the highest number of matched features that has a common connection to the current merged set. Finally, the images are merged into a common reference frame using the pair-wise transformations and weighted blending function is used to remove the seams.

It is apparent in Figure 3 that a metallic roof support (arrow shown) is visible in the image and is at a different temperature to the hanging wall rock. Since different materials have different thermal emissivity constants, this could be due to either a significantly different emissivity and similar temperature, or a real difference in temperature. As the support is metallic and the rock quartzite, and the ventilation air would preferentially cool the support as it protrudes into the airflow, it is likely that both reasons are valid. On the other hand, the blue (cooler) area at the bottom left of the image represents a cooler rock mass, and therefore indicates a potential threat where the rock mass is at risk of sudden separation. (Naturally, this was a concern to the author during data collection as he lay on his back taking images vertically upwards.) However, during the data collection it was noted that the cooler area represented an area that was protruding out of the hanging wall, and could be preferentially cooled having additional surface area exposed to the cooled ventilation air. Therefore, the blue area was not conclusively a high-risk region.

Additional data collection and verification is necessary to test this hypothesis in subsequent work.
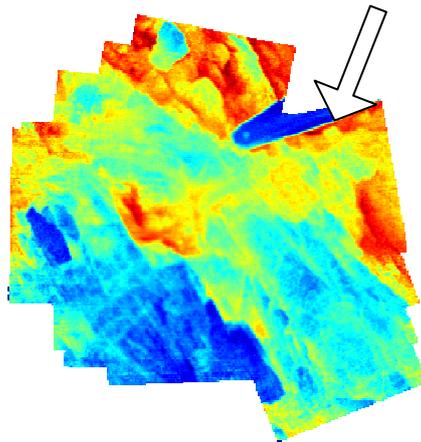
Figure 3.Stitched thermal image of mine hanging wall using 17 images,
blue=cold and red=warm (Driefontein Mine)

As a result of this data set it was concluded that three-dimensional profile information is required in order to robustly assess risk from the thermographic data. This led to the development of a multi-sensor setup geared for both thermal and 3D acquisition. This is discussed further in Section 5.

3.3. Subsequent data set (3rd)

The third data set was collected at Bafikeng Rasimone Platimum Mine (BRPM) at a shallow depth of 250m below surface approximately 8 hours after blasting, and after cleaning had been completed. Drilling had commenced. The data was collected with the multi-sensor setup comprising a FLIR A300 thermal imaging camera and two 3D sensors. The following sections discuss the 3D sensors and their use in producing combined thermal and 3D data respectively.

4. 3D MAPPING

To complement the thermal data with 3D rock structure, a logical step is to gather 3D data of the mine environment. Creating 3D maps of mines is not new [13], [14] used a laser scanner to map a mine portion and [15] successfully mapped a tunnel section from an abandoned mine.

The approach taken with this work was to use a high-end sensor (Reigl laser scanner) to generate a ground-truth data set. This demonstrated the potential of 3D data for our purpose, and resulted in an investigation into less expensive, more portable options for our multi-sensor setup. Ultimately, this led to the selection of a SwissRanger SR4000 TOF (time-of-flight) camera, and later this was compared with the Microsoft XBOX Kinect sensor. The following sections describe each sensor.

4.1. Riegl data set

A ground truth data set was collected using the Riegl LMS-Z390i 3D laser scanner. It is large, heavy, cumbersome, and scans slowly. While an impressive amount of data was collected, the process of transporting the sensor to the mine, and moving it to new

locations in the narrow stope environment made this a once-off activity. Newer alternatives that are more compact, such as the Faro focus 3D 20 and 120 scanners are available but they are still prohibitively expensive, large and slow for mobile applications.
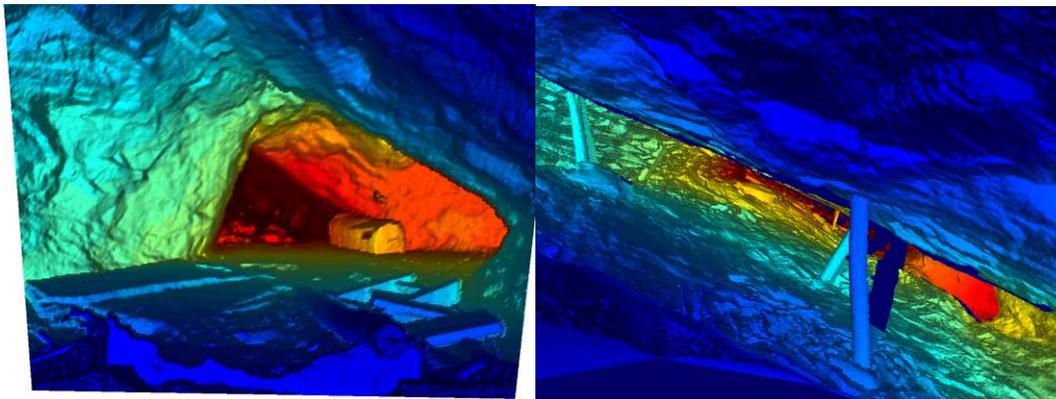

Figure 4: large scale Riegl laser scanner data set representations.

Tunnel and stope data was collected, and the visualisations in Figure 4 show the best level of detail that can be achieved with current technology. Images are textured by depth – blue for close and red for far away. From the data it is seen that there is indeed sufficient resolution for determining whether an area of wall is protruding into the ventilation air. Therefore, further sensors were evaluated for a mobile sensor head.

4.2. TOF data Set

A 5 m range MESA Imaging SwissRanger SR4000 was used for the next data set collection at a shallow (600m) disused gold stope. The distance ambiguity problem (aliasing) [16] typically experienced with this type of sensor proved to be a significant issue for processing. In a tunnel environment this is a persistent problem since extreme points in the direction of the tunnel often exceed the range of the camera, resulting in aliased measurements. Similarly, in a stope environment the hanging wall (ceiling) and foot wall (floor) stretch continuously beyond the sensor range causing the same problem. Although this can be addressed by amplitude filtering, it adds unwanted post-processing and is not completely robust. As a consequence, another 3D sensor option was sought towards achieving a faster mapping solution for thermal data fusion.

Figure 5 shows a 3D model generated from the TOF data. A median filter is used to reduce noise and ICP (iterative closest point) is used to register the 3D view of each range image. Following this, ray-carving is used to generate a triangular mesh model using Marching Tetrahedra polygonisation [17]. As the surface is traversed, points are projected into each range image and those that project to points in front of the closest range image are removed. This effectively carves away parts of the volume that are inconsistent with the set of range images. The model is post-processed with an iterative smoothing technique [18].
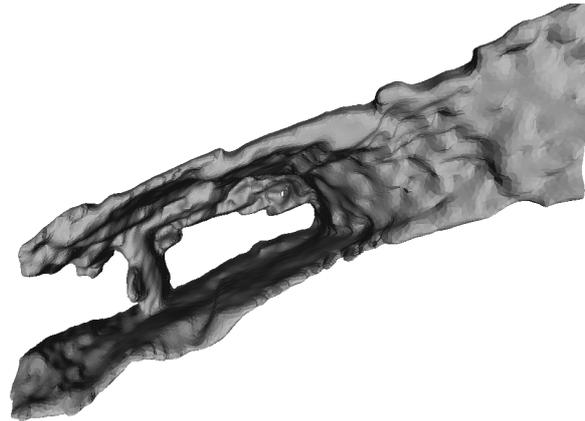
Figure 5. SR4000 data visualisation using a ray carving technique.(Gold Reef City Mine)

4.3. XBOX Kinect data set

During the evolution of the project a new sensor (the Kinect) was released, and has quickly become a hit in the robot community due to its low cost. The Microsoft XBOX Kinect has been successfully used in indoor mapping [19] using [20] and thus it was a natural progression to test the system underground. In contrast to TOF, Kinect uses a near-IR projector and camera to scan the environment. A colour camera is also built into the unit, but for 3D scanning, using just the IR camera means that no external light source is required.
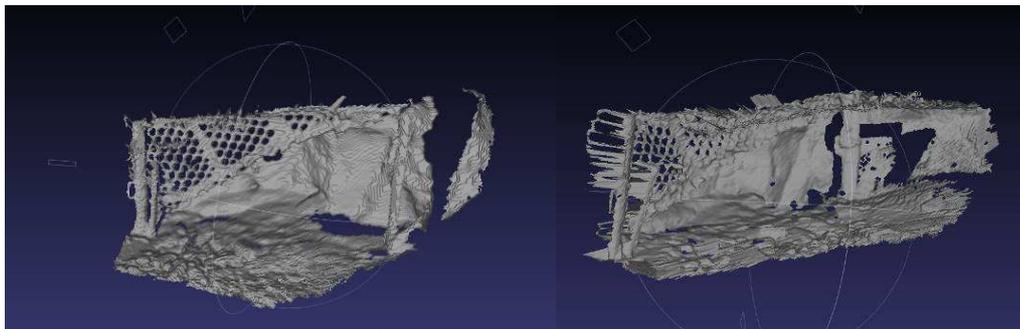
Figure 6. Two perspectives of the Kinect stope data set. (BRPM Mine)

The sensor was tested in an office environment and then mounted together with the other sensors during acquisition at BRPM (details given below). Figure 6 shows two screen shots of different perspectives from the BRPM Kinect data set. As with the TOF data, ICP is used to register the sequence and obtain a global reference frame for each range image. However, instead of using ray-carving, which can be slow and memory intensive, a synthetic combined range image was generated allowing a triangular mesh to be quickly generated using the image connectivity. This method was found to be preferable as it is less processor intensive and lends itself well to our thermal image stitching problem. Occluded areas where the surface cannot be seen from the various viewpoints is clearly seen by the empty sections. Merging more scans from different viewpoints can be used to create a full model, and will be undertaken in subsequent work.

5. COMBINATION OF THERMAL AND 3D DATA FOR VISUALISATION

Preliminary results showed that although the uniformity of the rock temperature produces limited features, it is possible to stitch images together to form a single large image. The combined image extends the field-of-view of the thermal imaging camera and facilitates analysis of potentially unstable areas.

5.1. The sensor head.

Our multi-sensor setup thus consisted of the three aforementioned sensors, namely the FLIR A300, the SR4000, and the Kinect, which were all rigidly mounted on a tripod. Relative calibration between the sensors was achieved by collecting correspondences through the use of a calibration object.

One of the difficulties of calibrating 3D and thermal cameras is that there is no common viewing spectrum. In one scenario, matching points were manually selected using the bottom of a beer crate. This worked well since the holes in the grid pattern create a depth difference as well as a thermal difference due to the varying background. However, manual calibration is slow so a more automated approach was sought in the form of a tennis ball mounted on a stick. A piece of backing card was mounted further back on the stick so that the ball could be clearly seen from its background. Thus, the sphere-like structure of the ball could be found in the range image, and by flexing the ball several times its glowing counterpart could be detected in the thermal image. Finally, by waving the ball around the scene multiple correspondences were automatically detected and used to calibrate the sensors. Figure 7 below shows images of the sensor head and calibration object.


Figure 7: Multi-sensor setup with SR4000 TOF camera, FLIR A300 IR camera and XBOX Kinect rigidly mounted on a tripod (left), and tennis ball calibration object(right).

During data capture, scans were taken from a single viewing position with the sensor head rotated about a pivot on a tripod. This provided a stable scanning path that facilitated ICP registration since the motion was less irregular.

5.2. Data fusion

Given the relative calibration between the thermal and 3D sensors, the correspondence between 3D points and their temperatures can be established by projecting into the thermal image. This enables texturing the model with the thermal data. One caveat: since the 3D camera streams approximately 10 times faster than the thermal imaging camera, it is important to ensure good synchronisation so that the thermal texture is correctly

aligned. Here, we used a simple time stamping approach to obtain a global reference during capture. During texturing, the 3D and thermal images with the smallest time difference are selected (or rejected if a maximum threshold is reached). Figure 8 shows an example where a single Kinect range image is textured with its closest (temporally) thermal image.
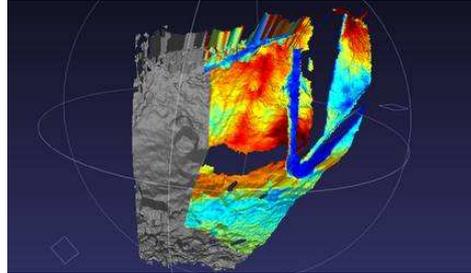
Figure 8. Single frame alignment fusing XBOX kinect and FLIR A300 IR camera.

Extending the above case to multiple scans with multiple thermal images requires that the 3D scans be registered in a common reference frame. Using ICP, we obtained an alignment for the BRPM data set and used this to generate a combined model as described in the previous section (3D mapping). Since each thermal image can be linked to its closest range image, the set of thermal images can be addressed in the common reference frame. Therefore, by projecting points from the combined model into the set of thermal images the combined model can be textured from multiple thermal images. We use a greedy approach where points are assigned a temperature from the most recent thermal image. Once assigned, the point is removed from the processing queue and only the remaining points are textured in the following iteration. Once all points have been processed or all thermal images have been exhausted the process terminates. A number of other merging strategies could be applied, such as weighted average. However, the FLIR camera produces extremely repeatable measurements and further filtering was not required.
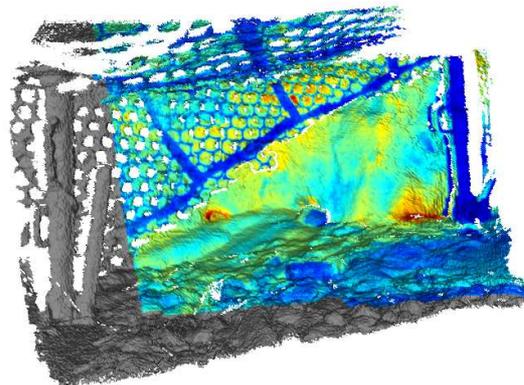

Figure 9. Merged 3D model textured with multiple thermal images (BRPM Mine).

Combining thermal and 3D data in this manner therefore provides sufficient information for a machine intelligence algorithm to determine potentially hazardous areas of the mine wall. The next phase of work will involve developing such an algorithm through the collection of multiple data sets that can be labelled by an expert. This will enable us to use the 3D thermal model to estimate the likelihood of a cooler area of rock being loose as opposed to protruding, and therefore having the potential to separate from the host

rock. An annotated visualisation can then be generated that can be used by miners to identify rocks that pose a danger.

5.3. Visualisation

Several visualisation techniques were explored so that humans can interact with data. Rendering a 3D textured model as shown in the previous figure is one method of visualisation. However, it may be better to present thermal and 3D structure as separate 2D images so that the distinction between 3D structure and temperatures can be readily seen. Figure 10 below shows an example where the textured model is reprojected into a virtual camera creating a stitched thermal image similar to those that were generated for the initial data set. However, here we are not constrained by our original simplified assumptions of near planar scenes because actual 3D points are known.
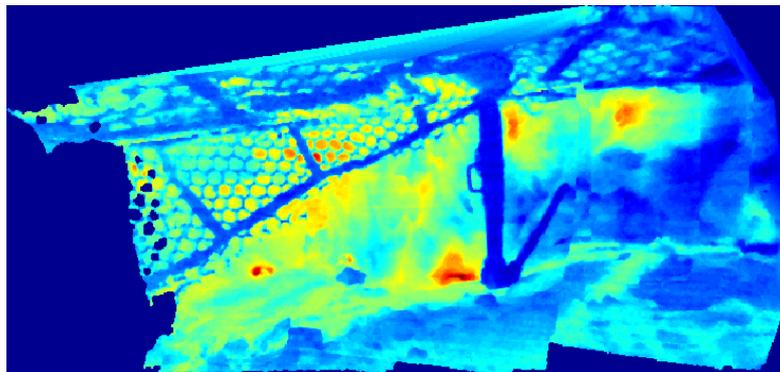
Figure 10: Wide angle visualisation - planar projection into a synthetic view (BRPM Mine).

As the model becomes larger it becomes more difficult to manage connectivity of the mesh. This is especially problematic when iteratively updating the model with new scans of a previously modelled region. One option is to simply plot the points. However, this creates a sparse structure that is difficult for visualisation. Another alternative is to assign a small patch (surface element or surfel) to each point and allow the patch to adapt to new measurements. Essentially, this provides the flexibility of a point-cloud but with surface characteristics of a mesh. This is demonstrated in Figure 11 below where an iterative model was generated using surfels [21]. Surfels are initialised when points occupying empty regions are added and updated us each range image is processed. Orientations are determined by estimating surface normals from the local pixel neighbourhood in the range images. A scale is assigned depending on the distance to the sensor. This allows model resolution to adapt appropriately when the camera is moved closer to a surface. Each surfel is modelled by a hexagon rendered with four triangles. In the figure, the surfels are coloured by distance from the sensor.
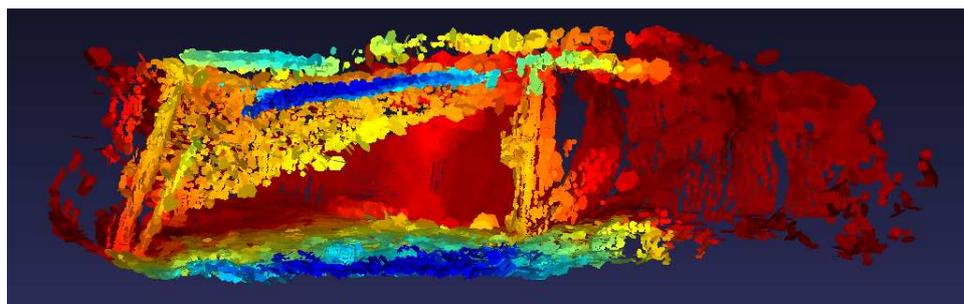
Figure 11. 3D model of BRPM data set constructed using surfels (BRPM Mine).

6. CONCLUSIONS

We propose that functional technologies exist for enabling a robot to navigate in the stope environment (30m X 3m X 1m) of South African Hard Rock mines using a dTOF beacon system for sensor localisation, and an array of sensors to map the environment. Using thermal imagery and 3D structure we have shown that a reasonable approximation of the mine can be created as a 3D model. It remains to be shown whether this is sufficiently accurate or consistent for machine intelligence algorithms to reliably generate hanging wall risk criteria.

The initial work with only an IR camera to evaluate hanging wall risk was insufficient for complete analysis as rock mass protruding into the stope will also be cooled by the passing ventilation air. This indicated that 3D topographical information would be needed to assess the rock mass stability. Additionally, the lack of features in the thermal data made stitching the data together unreliable. By rigidly mounting the thermal and 3D sensors, it was possible to register the 3D and thermal so that the resulting 3D surface could be thermally textured. This in turn could be used to determine rock mass stability. We have shown that once a risk criterion is achieved, the visualisation of that risk can be achieved in a number of ways. Synthetic view and surfel representation are two likely candidates for use in the mining application.

7. RECOMMENDATIONS

The success of the fusing of thermal data and 3 dimensional structure provides a positive foundation for a mine wall stability assessment tool. The development of robust machine intelligence algorithms to determine a suitable approximation of stability from the thermal and structural information is the next step in development. The data discussed above was captured from a single point-of-view (POV) using only rotation and elevation to vary the sensor viewpoint. This produced stable measurements and consistent features that aided registration.

The next step is to combine multiple POV data sets to build a more complete stope model, and then progress to moving sensors using loop closure to limit the accumulation of incremental errors. Visualisation of this information must be presented in such a way so as to be useful to a miner, and this should be explored further taking into account potential use in the unfavourable mining environment.

8. ACKNOWLEDGEMENTS

9. REFERENCES

[1] Sandvik. Automine. *2011(May 2),* Available: http://www.miningandconstruction.sandvik.com/.
[2] J. J. Green, P. Bosscha, L. Candy, K. Hlophe, S. Coetzee and S. Brink. Can a robot improve mine safety? Presented at CAD/CAM, Robotics and Factories of the Future.
[3] J. Green and D. Vogt. A robot miner for low grade narrow tabular ore bodies: The potential and the challenge. Presented at 3rd Robotics and Mechatronics Symposium (ROBMECH 2009).

[4] Loni Prinsloo. (2011, 6 JAnuary 2011). South african mine deaths down 24% in 2010. *2011(May 3),* pp. 2. Available: http://www.miningweekly.com/article/south-african-mine-deaths-down-24-in-2010-2011-01-06.

[5] K. Hlophe. GPS-deprived localisation for underground mines. Presented at 3rd Biennial CSIR Conference. Science Real and Relevant. Available: http://researchspace.csir.co.za/dspace/bitstream/10204/4225/1/Hlophe_2010_P.pdf.

[6] K. Hlophe, G. Ferreira and J. J. Green. A posture estimation system for underground mine vehicles. Presented at CAD/CAM, Robotics and Factories of the Future 2010.

[7] FLIR, "FLIR commercial vision systems, 'avoiding accidents with mining vehicles', Application Stories. 2008," vol. 2010, pp. 2, 2008.

[8] V. A. Kononov. (2000, September 2000). Pre-feasibility investigation of infrared thermography for the identification of loose hanging wall and impending falls of ground. Safety in Mines Research Advisory Committee. South Africa. Available: http://researchspace.csir.co.za/dspace/handle/10204/1811.

[9] D. Vogt, V. Z. Brink, S. Donovan, G. Ferreira, J. Haarhoff, G. Harper, R. Stewart and M. Van Schoor. Mining research for enhanced competitiveness. Presented at Science Real and Relevant: 2nd CSIR Biennial Conference.

[10] D. Vogt, V. Z. Brink and S. Schutte. New technology for real-time in-stope safety management. Presented at Hard Rock Safe Safety Conference 2009. Available: http://researchspace.csir.co.za/dspace/handle/10204/3680.

[11] D. Vogt, V. Z. Brink, S. Brink, M. Price and B. Kagezi. New technology for improving entry examination, thereby managing the rockfall risk in south african gold and platinum mines. Presented at CSIR 3rd Biennial Conference 2010. Science Real and Relevant. Available: http://researchspace.csir.co.za/dspace/bitstream/10204/4255/1/Vogt_2010.pdf.

[12] H. Bay, T. Tuytelaars and L. Van Gool. SURF: Speeded-up robust features. Presented at 9th European Conference on Computer Vision. Available: http://www.mendeley.com/research/surf-speededup-robust-features/.

[13] C. Baker, A. Morris, D. Ferguson, S. Thayer, W. Whittaker, Z. Omohundro, C. Reverte, W. Whittaker, D. Haehnel and S. Thrun. A campaign in autonomous mine mapping. Presented at IEEE International Conference on Robotics and Automation.

[14] D. Huber and N. Vandapel. Automatic 3D underground mine mapping. Presented at Field and Service Robotics 2003.

[15] A. Nüchter, H. Surmann, K. Lingemann, J. Hertzberg and S. Thrun. 6D SLAM with an application in autonomous mine mapping. Presented at IEEE International Conference on Robotics and Automation.

[16] D. Droeschel, D. Holz and S. Behnke. Multi-frequency phase unwrapping for time-of-flight cameras. Presented at Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on. Available: http://www.ais.uni-bonn.de/papers/IROS-2010-Droeschel.pdf.

[17] J. Bloomenthal. (1994, An implicit surface polygonizer. *Graphics Gems IV 1*pp. 324–349.

[18] G. Taubin. Curve and surface smoothing without shrinkage. Presented at Iccv.

[19] MIT, " Visual Odometry For GPS-Denied Flight And Mapping Using A Kinect," vol. 2011, March 2011, 2011.

[20] P. Henry, M. Krainin, E. Herbst, X. Ren and D. Fox. RGB-D mapping: Using depth cameras for dense 3D modeling of indoor environments. Presented at The 12th International Symposium on Experimental Robotics (ISER).

[21] H. Pfister, M. Zwicker, J. Van Baar and M. Gross. Surfels: Surface elements as rendering primitives. Presented at Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques 2000.