USING VIDEO ACQUIRED FROM AN UNMANNED AERIAL VEHICLE (UAV) TO MEASURE FRACTURE ORIENTATION IN AN OPEN-PIT MINE

T. McLeod, C. Samson, M. Labrie, K. Shehata, J. Mah, P. Lai, L. Wang and J.H. Elder
Department of Earth Sciences, Carleton University, Ottawa, Ontario

This project explored the feasibility of using video images acquired with an unmanned aerial vehicle (UAV) to obtain three-dimensional (3D) point clouds using structure from motion (SfM) software. Missions were flown using an Aeryon Scout: a lightweight, vertical take-off and landing quadrotor micro UAV with a miniature video camera. The initial mission captured urban scene images that were used to assess system performance while the main mission focused on rock walls where 3D images were used to successfully measure fracture orientations. Point clouds generated from this combination of technologies were sparse, but in the future, improvements in the resolution of original video images would cascade through the processing and improve the overall results. Such a system could have a multitude of applications in the mining industry, contributing to both safety and financial considerations.

Ce projet examine la faisabilité de l'utilisation d'images vidéo acquises au moyen d'un véhicule aérien sans pilote (UAV) pour obtenir des nuages de points tridimensionnels (3D) en utilisant un logiciel de structure à partir du mouvement (SfM). Des missions ont été effectuées avec un Aeryon Scout, un micro UAV léger à décollage et atterrissage vertical propulsé par quatre rotors et équipé d'une caméra vidéo miniature. La mission initiale a permis de capter des images de scènes urbaines utilisées pour évaluer la performance du système alors que la mission principale a mis l'accent sur des parois rocheuses où des images 3D ont été utilisées pour mesurer avec succès l'orientation des fractures. Les nuages de points générés à partir de cette combinaison de technologies étaient clairsemés, mais à l'avenir des améliorations de la résolution des images vidéo originales se succéderaient dans le traitement et amélioreraient les résultats globaux. Un tel système pourrait avoir une multitude d'applications dans l'industrie minière et contribuer à des considérations en matière de coût et de sécurité.

Introduction and Research Objectives

Obtaining detailed aerial images in 3D is challenging. While planes or helicopters mounted with LiDAR (Light Detection And Ranging) are typically used, the equipment, fuel, maintenance and pilot's time can be very expensive. UAVs, which are cheaper and easier to fly and maintain, are promising alternative to traditional methods [*Bento* 2008] even though there have been some issues with reliability and weight limitations that UAVs can handle (most airborne LiDAR systems weigh on the order of 50 to 100 kg) (as in *Frueh and Zakhor* [2003]). To circumvent weight constraints, a few research teams have flown still cameras on UAVs (weight ≈ 5 kg) and applied photogrammetric principles, combined with laser altimetry data and the use of cooperative target points, to construct 3D images and digital elevation models [*Lambers et al*. 2007; *Eisenbeiss* 2008; *Haarbrink and Eisenbess* 2008].

The project described in this paper explored the feasibility of an even lighter system: a miniature video camera (weight ≈ 100g) mounted on a small UAV (take-off weight ≈ 1.3 kg) that acquired 2D images that were used to obtain 3D point clouds with SfM software, rather than relying on a large UAV carrying heavy sensors [*McLeod* 2012]. The proof-of-concept project linked several existing technologies together—a UAV, SfM software and 3D image processing tools—to demonstrate that multiple technologies can be combined effectively, and to assess their strengths and limitations for imaging geological scenes.

Aeryon Scout UAV

System Description

The UAV selected for the project was an Aeryon Scout, a vertical take-off and landing

quadrotor micro-UAV (Table 1). The system included three main components: the vehicle, the control station and the base station.

The vehicle contained four moving parts and could be quickly assembled without a need for tools. The custom battery attached to the top of the vehicle with the payload secured to the bottom. The UAV navigated using a GPS with the aid of other sensor systems including sonar that established altitude for heights of 2 to 4 m above the ground depending on the underlying surface, a pressure altimeter (for altitudes beyond the range of sonar), a temperature sensor, a three-axis magnetometer, a three-axis accelerometer, and a three-axis gyroscope. Data collected during the flights were originally stored on internal storage and were downloaded after the missions. Log files recorded altitude, speed, position (latitude and longitude), and orientation of the camera (pitch and yaw).

The control station was handled with a touch interface tablet to control the vehicle and the camera. The left of the screen displayed the spatial position of the UAV on a map or georeferenced image to help with navigation; the right side of the screen contained a live video feed from the payload camera and status information. There were two ways to navigate the UAV: (1) using preselected waypoints to define a path for the vehicle to follow, or (2) by directing the UAV to a point on the map by touching and holding at the desired location. The vehicle was also programmed with smart fault handling for cases such as loss of communication, low battery or excessive wind.

Table 1: Specifications of the Aeryon Scout

Dimensions	80 cm x 80 cm x 20 cm
Battery	Lithium polymer
Battery voltage	12 V
Battery weight	450 g
Operational range	≤ 3 km
Flight endurance	≤ 25 min
Operational velocity	≤ 14 m/s (50 km/h)
Maximum take-off weight	1.3 kg
Payload weight	≤ 400 g
Wind limit	50 km/h (gusts to 80 km/h)
Maximum altitude	500 m

The base station relayed the signal from the control station to the UAV thereby allowing for long distance communication. The link between the control station and the base station was a short-range Wi-Fi connection, limiting the distance between the two devices to about 50 m (or about 5 m in noisy environments). A radio modem supplied the link between the base station and the vehicle allowing for an operational range of up to 3 km.

The payload used for the project was the day-time imaging camera. This vibration-dampened camera weighed 112 g and captured video images at a resolution of 640 x 480 pixels, at 12 frames per second with a field of view of 37.5° (horizontally) by 28.6° (vertically).

Flight Regulations

An important issue in conducting flights with a UAV in Canada is obtaining a Special Flight Operations Certificate (SFOC) from Transport Canada, the federal agency responsible for air traffic safety. For this project, using a UAV with a maximum take-off weight less than 35 kg that is operated within visual line of sight, below 400 ft in class G uncontrolled airspace and without autonomous navigation, allowed for a simplified application [*Transport Canada* 2008].

Structure From Motion Software

SfM software extracts 3D information from 2D images using corresponding feature points captured on multiple views. The particular implementation of SfM employed in this project used the Harris feature detector [*Harris and Stephens* 1988], which searches for locations in the image where there are two dominant and different edge directions in the local neighborhood of a point. It typically looks for points in the image that have a neighborhood region with high contrast and that are easy to identify from frame to frame: the criterion for a pixel to qualify as a feature point is a textured patch with high intensity variation in both x and y directions. The feature detector first estimates the intensity variation at each pixel inside the patch and constructs a covariance matrix. A pixel is then selected as a feature if both eigenvalues of the covariance matrix are higher than a certain threshold. For this work, the patch size chosen was 21 x 21 pixels.

Feature points are matched between frames using an iterative pyramidal implementation of the Kanade-Lucas-Tomasi (KLT) tracking algorithm

[*Shi and Tomasi* 1994]. The image at each pyramidal level is constructed to have half the resolution of the previous level. The original video image had a resolution of 640 x 480 pixels, therefore the resolution at the next levels were 320 x 240 pixels, 160 x 120 pixels and so on. Most cases are satisfied with a maximum of four levels. The analysis begins at the lower resolution to determine the amount of displacement between the frames and those results are used as the initial estimate of displacement in the next level up. The process continues until highest resolution is reached. Using this method, only the residual difference between the image matches needs to be computed at each level after the estimated correction is applied. The result is the location in the current image of a feature point in the previous image, which can also be expressed as the motion depicted in Figure 1.

In order to further validate matches between images, a model is used to relate corresponding feature points. The model used in the SfM implementation is the fundamental matrix (F), which relates corresponding points between two images using homogenous coordinates (x, y, 1) by describing a line (epipolar line) in the other image on which the corresponding point must lie [*Hartley and Zisserman* 2000]. For a match to be valid, this relation must hold. This relation comes from the fact that any 3D point is one vertex of a triangle in space linking itself to both camera focal points. That same triangle intersects the image planes of both cameras to form, in each image, the epipolar line. Therefore, the corresponding point of a projection in one image has to be on the epipolar line of the other image based on this triangular relation.

This implementation estimates the fundamental matrix using the RANdom Sample Consensus (RANSAC) algorithm, which works by initially using a small random subsample of data points to estimate the parameters defining the model and then determining other points that are consistent with the model, within a predefined threshold tolerance [*Fischler and Bolles* 1981]—we employed a threshold of 3 pixels from the epipolar line in the image. The model is then re-estimated using these new inlier points, and this process is repeated until convergence. The algorithm is run multiple times from different initial samples, and the model with the maximum number of inliers is selected.

The traditional SfM approach uses image pairs, but this implementation draws its information from a minimum of three frames in order to improve the robustness of the correspondence between feature points. Once corresponding feature points have been tracked through three frames, the software automatically detects when a frame has achieved the minimum required translation for accurate estimation. In most cases, every other frame was used despite the low velocity of the vehicle in the field (typically on the order of 2 m/s) to avoid skipping good, usable frames and maximize point cloud concentration. The criterion for a tracked frame to be selected and kept is based on the average triangulation angle from the two projected rays for all matches between the current frame and the previous one. A high angle threshold will result in points that are more accurate, based on the fact that the error on the range decreases as the triangulation angle increases. For instance, the worst case would be two parallel rays (stationary and/or panning camera). The system, in that case, is not well conditioned and the location of the intersection (3D point) can be anywhere on those rays, resulting in a very large error on the range. This work uses a minimum angle of 3 degrees for a frame to be used in building the model.

For each new frame, the matches are added to the global model and the locations of the 3D points in space are resolved using the new features and the ones from the previous tracked frames. A 3D point in the point cloud requires a minimum of three images as stated above, but can also be estimated based on more images if it is tracked in more than three frames. The point cloud is built incrementally. Each new frame adds new points (from the last trio of frames) and updates previous points by using an extra projection in the triangulation process.

The performance of SfM is limited by several fundamental factors. The number of points in the output point cloud is determined by the quality of the raw video data, motion blur and the ability of the software to extract salient features that are tracked consistently in successive frames. Also, although the relative distances between the various features in the final cloud are internally consistent, the scene can be recovered only up to a scale factor. On the other hand, the scale factor could easily be estimated using a single distance measurement, for example, by using a time of flight laser to measure the distance from the camera to a particular point in the scene.

System Performance

A mission was conducted in an urban environment (York University's Keele Campus in Toronto, Canada) to help establish system performance. In urban scenes, buildings tend to have surfaces of relatively uniform intensity: banks of windows, large panels, homogenous brickwork and so on, which are not conducive to feature extraction. There were

some buildings, however, that were more compatible with SfM—those with many smaller windows or contrasting panels in the façade, such as the Life Sciences building at York University (Figures 1 and 2). By comparison, natural scenes can offer more variation in intensity in a small area which is far more conducive to good results with SfM; rock faces, vegetation, and so on, do not tend to have the smooth, regular surfaces of man-made structures and as such, natural scenes are captured far more readily.

The point cloud of the relatively flat west façade of the Life Sciences building, with an average point cloud concentration of 0.0015 point/cm^2, provided an opportunity to estimate range accuracy (Figure 2). The windows of the first level of the building are set back approximately 30 cm compared to the upper levels so were removed for the analysis to eliminate any potential effect that points from the first level may have caused. A plane of best fit in the least-squares sense was determined for the data using principal component analysis (PCA) and the distance from that plane was calculated. For a range of 25-30 m (as delineated by the vertical axis of figure 2, bottom left), the analysis indicated that 68% (one standard deviation) of the points fell within 44 cm of this plane.

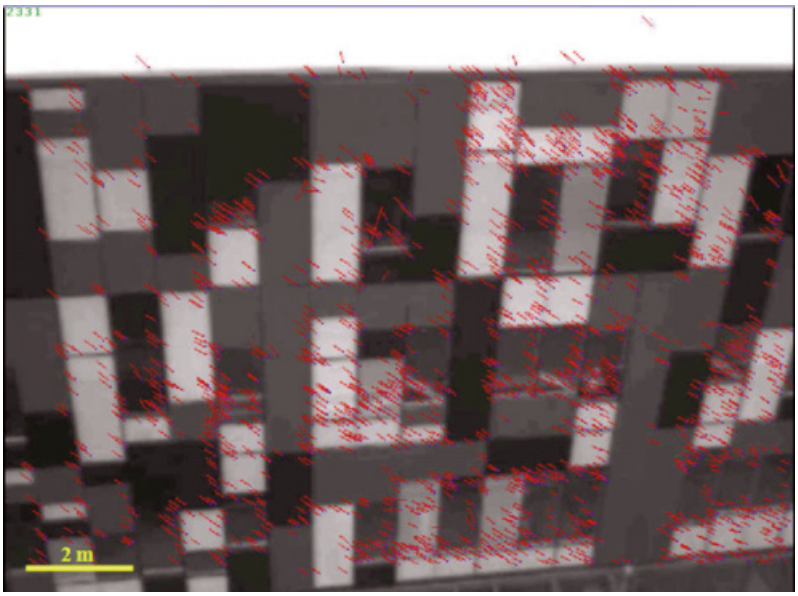

Figure 1: Video still of the west façade of the Life Sciences Building at York University showing tracked features (blue dots) and their motion (red lines) calculated by the KLT algorithm. The UAV was flying north (right to left) at an altitude of approximately 25 m and a speed of 2 m/s. There are 2,331 tracked points in this frame, concentrated along the perimeters of panels where there is strong contrast. The scale bar is approximate. A photo of the façade is included in Figure 2.

Imaging Mining Scenes

Open-Pit Mine Survey

The Canadian Wollastonite Mine in Seeley's Bay, Ontario [*Grammatikopoulos et al*. 2003], which is currently in the early stages of development, boasted an exploration trench that provided our mining site. The goals of the mission were (1) to enhance the visual perception of the 3D images in post-processing and (2) to determine if software designed to automatically estimate joint orientation [*Mah et al*. 2012] could be successfully applied to the point clouds generated using SfM.

The trench where the flights were performed (September 9, 2011) was oriented approximately north-south, bordered on the west and east by walls of wollastonite approximately 7 to 8 m at the highest points. Each wall was about 190 m long. All of the flights were flown at low altitude (less than ≈15 m) with the camera dipping about 10° in order to ensure near-perpendicular incidence to the wall. The densest point clouds resulted from video images acquired closest to the rock face (distance <8 m) and that covered the same small area repeatedly over a short time, moving back and forth, as the UAV was brought closer to the rock face.

The walls featured several fracture systems, resulting in a stepped appearance, and there was also vegetation growing on some of the benches and crevices. A total of 86 manual measurements of strike and dip were taken in order to validate the results of analysis of the point clouds: 42 on the east wall and 44 on the west wall. The measurements defined three joint sets (strike/dip in degrees): 329/85 (joint set #1), 183/8 (joint set #2) and 32/78 (joint set #3). The site was also measured with a tape measure because SfM cannot provide scale.

Visualization Software

The visualization software drapes point clouds with polygonal meshes using a 2.5D Delaunay triangulation approach [*Lee and Schachter* 1980] where the (x,y) coordinates are used for triangulation and the z coordinate is used as the height value of the vertex. The software also implements a mesh-smoothing algorithm known as the geometric Laplacian which was used to reduce scatter and improve visualization. In each successive iteration of the geometric Laplacian algorithm, the location of each vertex in the triangulation is adjusted from its location based on a weighted average of the positions of the neighboring points one edge away.

The strike and dip orientation of the triangular elements were determined by using the normal vec-

tor of the triangles and then mapped to the Hue Saturation Lightness (HSL) colour model [*Boivin et al*. 2013]. In this colour model, the 'hue' component determines colour and the 'saturation' component determines the intensity of the colour. Since hue ranges from 0° to 360°, it is a good choice for mapping strike. Dip (0° to 90°) can be expressed using saturation which is a number between 0 and 1 with 0 as grey and 1 as the colour itself. The lightness component was kept at a constant as it represents how much light is reflecting off the surface. Once the strike and dip have been converted to the HSL colour scheme, they are converted into an RGB colour model for viewing (Figure 3). Joint sets #1, #2 and #3 corresponds to patches of magenta/red, cyan/blue, and orange/yellow, respectively. The georeferenced surfaces have joint orientations that differ from the manual measurements by approximately 10°.

Joint Orientation

Joint orientation is an important parameter related to the strength of rock masses. *Mah et al*. [2012] have developed a methodology for estimating joint orientation directly from point cloud data. A best-fit plane is computed through a subset of points within a specified radius using PCA, and strike and dip values are derived directly from the orientation of the plane (Figure 4). Their software, however, was originally designed for point clouds with a very high concentration (≈ 40 point/cm^2) and cleaner rock surfaces devoid of vegetation. At the Canadian Wollastonite Mine, surfaces were noisier and the coverage was sparse, therefore the radius had to be increased as high as 75 cm so that enough points were included in the analysis (Table 2). The low point cloud concentration (0.01-0.10 points/cm^2) made it difficult to produce a reliable result while not including any nearby edges which would add an element of noise and skew the results. The log files collected by the Scout included information on the yaw and pitch of the camera used to georeference the results.

Six sites were conducive to analysis; at four of them the data was smoothed using the geometric Laplacian before analysis (Figure 3, Table 2). The dip of the sub-vertical surfaces agreed with the manual measurements within 3°. There was only one measurement done on a sub-horizontal surface because such surfaces are notoriously difficult to capture at grazing angles, leading to low point cloud concentration [*Sturzenegger and Stead* 2009; *Mah et al*. 2011]. In addition, results at that location could have been negatively affected by accumulated debris. At sites 27 and 29 on the west wall where the point cloud concentration was the highest, the angular difference between the normal vectors of the manual measurements and the results from the point cloud analysis is at most 2°.

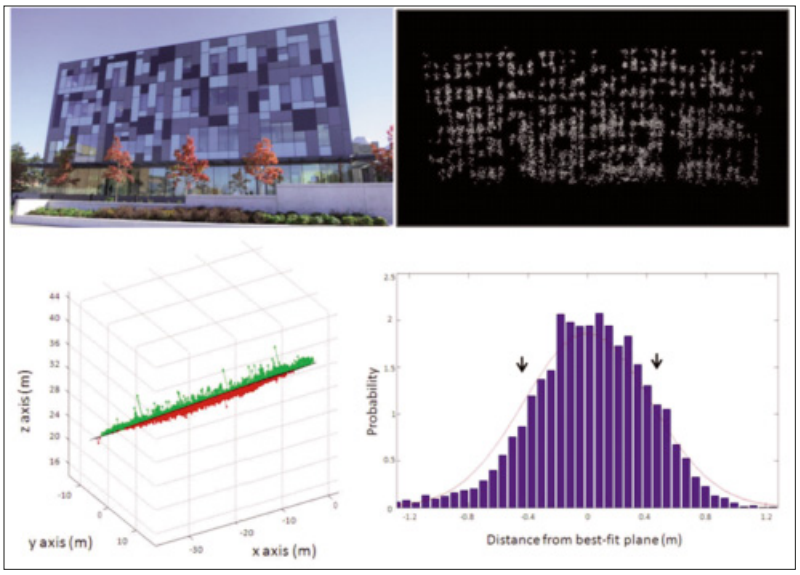

Figure 2: Photograph (top left) of the west façade of the York University Life Sciences building (courtesy of Bob Hou) which is approximately 39.5 m wide and 12 m high. Corresponding point cloud in cross-sectional view (top right) of the west façade. The point cloud was generated from 20 frames captured from a distance of ~25 m at a velocity of ~7.5 m/s, and includes 7,521 points. PCA output for full face (bottom left) showing green points above and red points below the best-fit plane, and corresponding Gaussian distribution of errors (bottom right); arrows indicate one standard deviation.

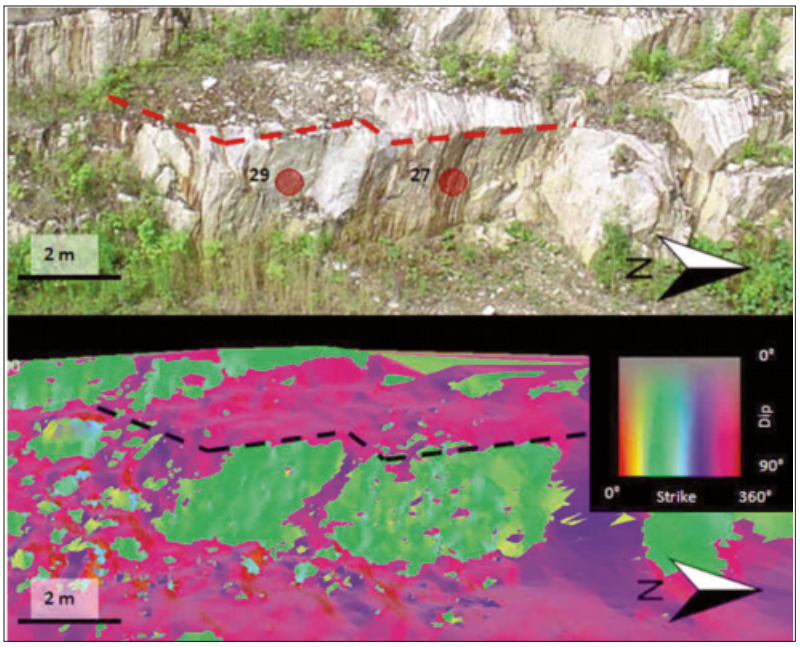

Figure 3: Digital photograph of a portion of the west wall (top) and the corresponding dip and strike overlay (bottom). Joint orientation measurements done at sites 27 and 29 are reported in Table 2.

Discussion and Concluding Remarks

The project demonstrated that the technologies under consideration could be effectively combined and that the Scout UAV is a highly maneuverable way to capture a rich variety of video scenes. The main challenge was in applying SfM processing methods at the relatively low 640 x 480 pixel resolution which, in turn, resulted in an overall lower point cloud concentration. Points on a well-represented flat surface provided an opportunity to quantify the range accuracy which was on the order of a half meter at a distance of 25 to 30 m. However, natural scenes with high intensity variation produced higher point cloud concentrations than man-made surfaces of urban scenes (0.01-0.10 point/cm^2 versus 0.002 point/cm^2).

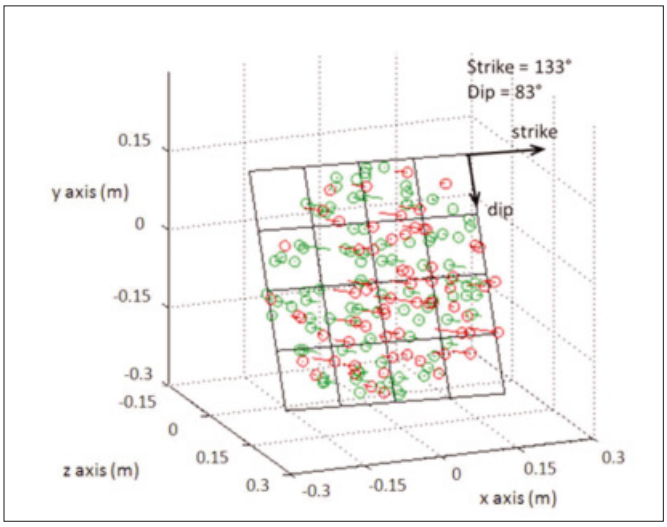

Figure 4: PCA results at site 29 on the west wall. Red and green points are below and above the best-fit plane, respectively.

The point cloud concentrations and range accuracy achieved in practice constrained the image processing algorithms that could be used to process the 3D data generated by the prototype. Encouragingly, rock mass characterization and visualization software designed for point cloud concentrations at several orders of magnitude higher performed well and led to reasonable estimates of fracture orientation.

The technology is advancing rapidly and the new Aeryon SkyRanger can record 30 frames per second at 1080p. With higher resolution video, results will continue to improve and further progress towards a robust system lie in developing a quantitative understanding between the UAV flight parameters (distance to target, velocity, incidence angle, etc.) and the characteristics of the 3D point clouds generated by SfM.

A working system of the type explored in this project would present an opportunity for many applications. Within the mining community, it could offer a safe alternative for surveying potentially hazardous rock faces from a safe distance and allow for inexpensive repeatability of aerial scans as well as the option of gaining aerial data on a smaller scale. It could also allow a mining operation to create and maintain an archive of regular 3D models of a site at different phases during its development. Such a system could also be used to create estimates of rock volume produced after a blast to expedite its removal and potentially to determine and optimize fragmentation (the size of the rocks

Table 2: Comparison between manual joint orientation measurements and results from the point cloud analysis

		Manual compass measurements			Point cloud analysis						
Site	Joint set number	General Orientation	Strike (°)	Dip (°)	Smoothing	Point cloud concentration (point/cm^2)	PCA radius (cm)	Number of points in PCA	Strike (°)	Dip (°)	Angular Difference (°)
West 27	2	subvertical	136	85	no	0.09	22.5	138	136	84	1
West 29	2	subvertical	134	81	no	0.10	22.5	160	133	83	2
East 9	1	subvertical	283	80	yes	0.01	75.0	207	244	77	38
East 9 (parallel face)	1	subvertical	283	80	yes	0.02	75.0	291	264	82	19
East 10	2	subvertical	186	88	yes	0.01	75.0	141	153	89	33
East 13	3	subhorizontal	132	11	yes	0.01	75.0	199	189	20	17

produced by blasting). In addition to its value in mining, there could also be geological applications such as determining slope stability.

Acknowledgements

This project was supported by GEOIDE and the Canadian Network of Centres of Excellence in Geomatics, via a Phase IV project awarded to Dr. James Elder (principal investigator). We are grateful for the assistance of many people from the groups of Dr. James Elder and Dr. Gunho Sohn during the mission at the Keele campus: Eduardo Corral Soto, Ron Tal, Bob Hou, Vishal Kumar, Herman Badwal, Chao Luo, Ravi Persad, and Solomon Chan. We wish a warm thanks to Mr. Bob Vassily, President of Canadian Wollastonite, for granting us access to his property. We would also like to express our appreciation to Prof. Steve McKinnon of the Department of Mining at Queen's University for leading us during a reconnaissance tour of the mine, and to Ms. Sarah Davey, M.Sc., for her help during the mine mission.

References

Bento, M. 2008. Unmanned Aerial Vehicles: An Overview. *Inside GNSS*, **3**(1): 54-61.

Boivin, A., P. Lai, C. Samson, E. Cloutis, S. Holladay, and F.A. Monteiro Santos. 2013. Electromagnetic induction sounding and 3D laser imaging in support of a Mars methane analogue mission. *Planetary and Space Science*, 82-83: 27-33.

Eisenbess, H. 2008. The autonomous mini helicopter: a powerful platform for mobile mapping. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Vol. XXXVII, Part B1, 977-984.

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

Frueh, C. and A. Zakhor. 2003. Constructing 3D City Models by Merging Ground-Based and Airborne Views. *Proceedings of the 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2: 562-569.

Grammatikopoulos, T.A., A.H. Clark, and B. Vasily. 2003. The St. Lawrence deposit, Seeley's Bay, SE Ontario—A major wollastonite skarn in the Frontenac Terrane. *CIM Bulletin* 96 (1070): 47-54.

Haarbrink, R.B. and H. Eisenbess. 2008. Accurate DSM production from unmanned helicopter systems. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Vol. XXXVII, Part B1, 1259-1264.

Harris, C. and M. Stephens. 1988. A Combined Corner and Edge Detection. In *Proceedings of the Fourth Alvey Vision Conference*. 147-151.

Hartley, R. and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge Univ. Press, Cambridge, UK, 2000.

Lambers, K., H. Eisenbeiss, M. Sauerbier, D. Kupferschmidt, T. Gaisecker, S. Sotoodeh, T. Hanusch. 2007. Combining photogrammetry and laser scanning for the recording and modeling of the late intermediate period site of Pinchango Alto, Palpa, Pery. *J. of Archeological Science,* 34(10): 1702-1712.

Lee, D.T. and B.J. Schachter. 1980. Two algorithms for constructing a Delaunay triangulation. *Int. J. Parallel Programming,* 9(3): 219-242.

Mah, J., C. Samson, and S. McKinnon. 2011. 3D laser imaging for joint orientation analysis. *Int. J. of Rock Mechanics and Mining Sciences* 48(6): 932-941.

Mah, J., C. Samson, S. McKinnon, and D. Thibodeau. 2012. 3D laser imaging for surface roughness analysis. *Int. J. of Rock Mechanics and Mining Sciences,* 58: 111-117.

McLeod, T. 2012. 3D imaging applications in earth sciences using video data acquired from an unmanned air vehicle. Master of Science thesis 2012, Carleton University.

Robertson, D.P., and R. Cipolla. 2009. Structure from Motion. In Varga, M. (Ed.), *Practical Image Processing and Computer Vision*, John Wiley.

Shi, J. and C. Tomasi. 1994. Good features to track. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (CVPR94).

Sturzenegger, M, and D. Stead. 2009. Close-range terrestrial digital photogrammetry and terrestrial laser scanning for discontinuity characterization on rock cuts. *Eng. Geol*.106: 163-182.

Transport Canada. 2008. Staff Instruction: The review and processing of an application for a Special Flight Operations Certificate for the Operation of an Unmanned Air Vehicle (UAV) System. Ottawa: General Aviation.

MS rec'd 13/03/28
Revised MS rec'd 13/07/2

Authors

Tara McLeod
mcleod.tk@gmail.com

Tara McLeod has an undergraduate degree and a master's degree in Earth Sciences from Carleton University. She was introduced to research through an undergraduate research award from the National Science and Engineering Research Council of Canada in the summer of 2010 working with 3D images of meteorites. Building on that experience, her master's project focused on 3D imaging in natural and urban scenes using structure from motion techniques. Her interests include hydrogeology, planetary geology and geophysics.

Claire Samson
Claire.Samson@carleton.ca

Martin Labrie

Jason Mah

Kareem Shehata

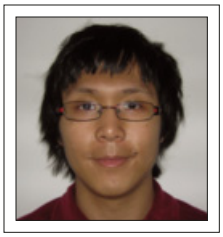

Po Lai

Langyue Wang

James Elder

Claire Samson has an undergraduate degree in engineering physics from Laval University, an M.Sc. in geological sciences from McGill and a Ph.D. in physics from the University of Toronto. She has a wide range of experience, both in industry and academia. In 1991-92, she was a Research Associate at Cambridge University in England. From 1993 to 1999, she worked for the Shell Oil group in the Netherlands. In 2000, she joined Neptec Design Group, a high-tech company specializing in vision systems for space applications. She joined Carleton University in 2003. In both her research and fieldwork, Claire is keen to address practical problems and to involve industry partners. Her current interests include 3D imaging, airborne geophysics and planetary geology.

Martin Labrie has a master's degree in computer science from Laval University. He specializes in computer vision and 3D processing. He also has a strong background in shape from motion algorithms. Martin worked at Defense Research and Development Canada from 2002 to 2005. He joined Neptec Design Group, a high-tech company specializing in vision systems and 3D sensors for space applications, in August 2005. Martin has a wealth of experience in stereo-vision systems, camera calibration and 3D processing. He also has the background to implement vision algorithms in hardware such as FPGA.

Jason Mah has an undergraduate degree in Mechanical Engineering from the University of Waterloo and a Ph.D. in Earth Sciences from Carleton University. He has a multi-disciplinary background from both industry and academia. Since 2003, he has worked for Neptec Design group, a company specializing in vision systems and rover technologies. In this capacity, he specialized in thermal design and environmental testing of space systems. His Ph.D. focused on 3D laser imaging for underground rock characterization. In the winter of 2012, he was the course instructor for Engineering Geoscience at Carleton University. His current research interests include 3D laser imaging, geophysics, and rock mechanics.

Kareem Shehata is the lead software developer at Aeryon Labs, and is responsible for the software design of both the Aeryon Scout and the Aeryon SkyRanger. His work with robotics dates back to his studies at the University of Waterloo, where he took part in student teams while completing a B.A.Sc. in Electrical Engineering. Kareem has over 10 years of experience in industry, including working with Achronix/Honeywell Aerospace on electronics and software for the JSF, and ATS Automation Tooling Systems for several advanced automation projects such as the BioProcessors automated cell experimentation system.

Po Lai has an undergraduate degree in computer science from Carleton University. During his undergraduate degree, Po worked at several software companies and gained valuable experience in the industry. He was first exposed to research through an undergraduate summer research award from the National Sciences and Engineering Research Council of Canada in 2008. Po is in the final stages of completing a master's in computer science at Carleton University and is planning to pursue doctoral studies in computer vision at the University of Ottawa. Po is passionate about solving problems in earth sciences through the use of computational frameworks. His research interests include computer vision, computational geometry, and applications of computer science to geology.

Langyue (Larry) Wang, working as a Web Mapping Developer for the Weather Network, is also a part-time Ph.D. student at GeoICT Laboratory of York University studying Geomatics Engineering. He Received his M.Sc in Geoinformatics from the International Institute for Geo-Information Science and Earth Observation (ITC), the Netherlands. He has over 13 years of working experience in GIS, remote sensing and mapping field. He also worked as a GIS Consultant for a Japanese high-tech company and in leadership roles for a Chinese governmental mapping organization. His current research interest is to combine geomatics technology and computer vision algorithms to bring real-time dynamics into 3D virtual world to provide augmented location based intelligence.

James Elder received his Ph.D. in Electrical Engineering from McGill University in 1996. He is currently a member of the Centre for Vision Research and a Professor in the Department of Electrical Engineering and Computer Science and the Department of Psychology at York University. His research seeks to improve machine vision systems through a better understanding of visual processing in biological systems. Dr. Elder's research has won a number of awards and honours, including the Premier's Research Excellence Award and the Young Investigator Award from the Canadian Image Processing and Pattern Recognition Society. His 3DTown research has recently been featured in the National Post and on CBC radio and television. ❑