

# **3D imaging applications in Earth Sciences using video data acquired from an unmanned aerial vehicle**

by

Tara McLeod

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in partial  
fulfillment of the requirements for the degree of

Master of Science

In

Earth Sciences

Carleton University

Ottawa, Ontario

©2012

Tara McLeod

## **Abstract**

For three dimensional (3D) aerial images, unmanned aerial vehicles (UAVs) are cheaper to operate and easier to fly than the typical manned craft mounted with a laser scanner. This project explores the feasibility of using 2D video images acquired with a UAV and transforming them into 3D point clouds.

The Aeryon Scout – a quad-copter micro UAV – flew two missions: the first at York University Keele campus and the second at the Canadian Wollastonite Mine Property. Neptec’s ViDAR software was used to extract 3D information from the 2D video using structure from motion.

The resulting point clouds were sparsely populated, yet captured vegetation well. They were used successfully to measure fracture orientation in rock walls. Any improvement in the video resolution would cascade through the processing and improve the overall results.

## Acknowledgements

This dissertation would not have been possible without the help and support of numerous individuals. I would like to express my gratitude to Professor Claire Samson, my research supervisor, for her guidance, enthusiasm, encouragement and useful critiques.

Thanks to Aeyron Labs for the use of the Scout, and especially to Mike Peasgood who was my initial contact and who got things started, to Tom Pike who trained me on the system and Kareem Shehata for his help in the field at York, in troubleshooting during the mine mission and for helping me decipher the log files.

Thanks to Neptec Design Group for the use of their ViDAR software and especially to Martin Labrie who showed me how the program worked and helped with troubleshooting along the way, as well as pointing me to sources about the workings of structure from motion software.

Thanks to the people at York University and all their help. I am grateful for the assistance of many people in the field at York: Prof. James Elder, Eduardo Corral Soto, Ron Tal, Bob Hou, Larry Wang, Vishal Kumar, Herman Badwal, Chao Luo, Ravi Persad, and Solomon Chan. I also wish to thank Prof. James Elder and Prof. Gunho Sohn and their labs for hosting me at York and their input and suggestions regarding data processing. I am particularly grateful to Larry Wang for the PCA program for the York data and to Bob Hou who provided several photos.

A warm thanks to Mr. Bob Vassily, President of Canadian Wollastonite, for granting us access to his property. I would also like to express my 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 Sarah Davey, M.Sc. for her help during the survey.

I would like to express great appreciation to Jason Mah, Ph.D. for providing the rockmass characterization software and for all of the time he spent adapting it for use with my data. Thanks also to Po Lai, M.Sc. candidate at Carleton University, for the use of his visualization software and for providing the smoothing algorithm for the mine data. I wish also to acknowledge John Harrison, Ph.D from McGill University for taking the time to orient me with his software.

# Table of Contents

Abstract.....	ii
Acknowledgements.....	iii
List of Tables .....	vii
List of Figures .....	viii
1 Introduction and Research Objectives.....	1
2 Unmanned Aerial Vehicles.....	7
2.1 General.....	7
2.2 Flight Regulations.....	9
2.3 Aeryon Scout.....	11
2.4 UAVs in urban settings.....	19
2.5 UAVs in Earth Sciences.....	20
3 Imaging in three dimensions.....	23
3.1 Imaging methods .....	23
3.2 Structure from Motion theory.....	24
3.3 ViDAR software .....	29
4 Image Processing .....	35
4.1 Principal Component Analysis .....	35
4.2 Classification software .....	35
4.3 Rock mass characterization software and associated programs.....	38
4.4 Visualization software.....	39
5 Missions .....	40

5.1 Familiarization.....	40
5.2 York Mission.....	40
5.2.1 Special Flight Operations Certificate .....	41
5.2.2 Survey site.....	42
5.2.3 Flights.....	42
5.3 Mine Mission.....	47
5.3.1 Special Flight Operations Certificate .....	47
5.3.2 Survey site.....	47
5.3.3 Flights.....	52
6 Results and Discussion .....	56
6.1 Test Flights .....	56
6.2 York Mission.....	60
6.2.1 Point cloud representation of different features .....	60
6.2.2 Range accuracy .....	65
6.2.3 Classification software .....	72
6.3 Mine Mission .....	74
6.3.1 Rock mass characterization software .....	74
6.3.2 Visualization software.....	82
7 Concluding remarks .....	86
References .....	89
Appendix A: Basics of structure from motion.....	94
Appendix B: Manual measurements from mine site.....	96

## List of Tables

Table 2.1: UAV classification system.....	8
Table 2.2: Specifications of the Aeryon Scout .....	12
Table 2.3: Specifications of the daylight camera .....	18
Table 5.1: Details of the York University flights.....	46
Table 5.2: Details of the mine flights .....	54
Table 6.1: Range accuracy results.....	71
Table 6.2: Results from the rock mass characterization analysis.....	78

## List of Figures

Figure 1.1: Implementation flow chart.....	2
Figure 2.1: Aeryon Scout.....	13
Figure 2.2: Scout tablet display.....	15
Figure 2.3: Scout base station.....	16
Figure 2.4: Scout daylight camera .....	18
Figure 3.1: Structure from Motion basic principles.....	26
Figure 3.2: 3D to 2D transformation with a pinhole camera.....	28
Figure 3.3: Building a point cloud in ViDAR .....	31
Figure 3.4: ViDAR algorithm diagram .....	32
Figure 3.5: KLT algorithm diagram.....	33
Figure 4.1: Normals from a ZH operator.....	37
Figure 5.1: Map of York University Keele Campus .....	43
Figure 5.2: Subsection of Keele Campus.....	44
Figure 5.3: Stereonet of manual measurements from the mine site.....	51
Figure 5.4: Satellite photo of Canadian Wollastonite Mine flight area.....	54
Figure 5.5: Scout flying at the Canadian Wollastonite Mine .....	55
Figure 6.1: Gravel mound point cloud in ViDAR .....	57
Figure 6.2: Gravel mound mesh.....	58
Figure 6.3: Gravel mound mesh with corresponding photo .....	59
Figure 6.4: Feature tracking in ViDAR on the Life Sciences building .....	62
Figure 6.5: Feature tracking in ViDAR on the computer science building.....	63

Figure 6.6: Point cloud of the computer science building..... 64

Figure 6.7: Best fit of Petrie point cloud to wire frame model..... 66

Figure 6.8: Point cloud of the west face of the Life Sciences building ..... 68

Figure 6.9: PCA results for the Life Sciences building ..... 69

Figure 6.10: Distribution of points for the Life Sciences building..... 70

Figure 6.11: Point cloud of the northeast corner of the computer science building..... 73

Figure 6.12: Mine mission flight path ..... 75

Figure 6.13: Locations of west 27 and west 29 at mine site ..... 80

Figure 6.14: PCA results for west 29 ..... 81

Figure 6.15: Visualization results on unsmoothed point cloud ..... 83

Figure 6.16: Visualization results on smoothed point cloud ..... 85

# 1 Introduction and Research Objectives

This project is a proof-of-concept linking several existing technologies – an unmanned aerial vehicle (UAV), structure from motion (SfM) software and 3D image processing tools – into a working prototype geared toward applications in the earth sciences (Figure 1.1). The overarching objective of the project is to demonstrate that the technologies can be combined effectively, and to assess their strengths and limitations for the chosen applications.

The UAV selected for the project is the Scout: a small quadrotor UAV developed by Aeryon Labs of Waterloo, Ontario that is equipped with a daytime video camera<sup>1</sup>. The video images recorded by the Scout during flight are converted into 3D images using ViDAR, a particular implementation of SfM principles programmed by Neptec Design Group of Ottawa, Ontario<sup>2</sup>. A typical 3D image from this process contains thousands to tens of thousands of points with (x, y, z) Cartesian coordinates grouped into an unstructured “point cloud”. Three 3D image processing programs were used: classification software that identifies points as singular, linear or planar (Harrison, 2009; unpublished), rock mass characterization software that measures the orientation of fractures on rock walls semi-automatically (Mah et al., 2011) and visualization software that displays the strike and dip of surfaces in a mesh derived from the points (Po Lai, personal communication).

---

<sup>1</sup> [www.aeryon.com](http://www.aeryon.com)

<sup>2</sup> [www.neptec.com](http://www.neptec.com)

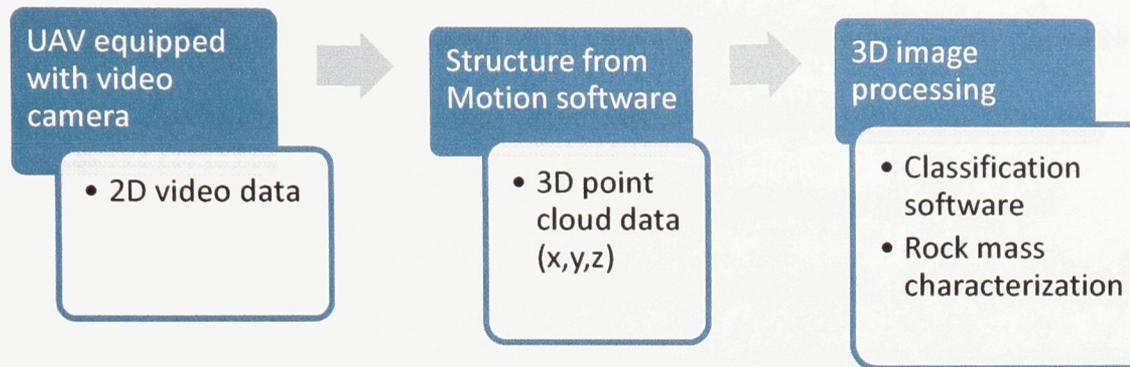


Figure 1.1: Implementation flowchart

In practice, obtaining detailed, cost-effective aerial images in three dimensions (3D) is challenging. While planes or helicopters mounted with laser scanners are typically used, the equipment, fuel, maintenance and the pilot's time are expensive. For safety reasons, these aircraft also need to maintain a certain altitude or distance from the target. The first airborne Light Detection and Ranging (lidar) systems were used primarily to acquire topographic data (Campbell and Wynne, 2011). Currently, most airborne lidar systems, which are mounted on planes and helicopters, are used for laser altimetry. It also has applications in urban areas for determining building densities, mapping urban structures and building footprints, mapping terrain and assessing flood risk (Campbell and Wynne, 2011; Crutchley and Crow, 2009); in forestry for estimating forest structure and timber volume (using multiple returns); in archeology for documenting structures under heavy vegetation and delineating features that are otherwise obscured by modern development (Crutchley and Crow, 2009); and in hydrography to monitor coastal features and map subsurface topography (green wavelengths penetrate water) (Campbell and Wynne, 2011). The resolution of the collected data depends on the intended application and is often quite crude (several meters) when the goal is delineating large targets like roads and buildings (Campbell and Wynne, 2011). Other applications need better resolution: the Environment Agency in the UK used 2m resolution for terrain mapping and for archeological investigations a 0.5m resolution works well. Resolution depends on factors such as the altitude and speed of the platform (lidar mounted on helicopters produce much higher resolution than the data acquired from fixed-wing aircraft) as well as factors relating to the lidar

system itself such as the number of pulses per second, the scan frequency and the scan angle. Higher resolutions require flights at lower elevation, larger overlaps and smaller scan angles (Crutchley and Crow, 2009). Although the technology is moving toward higher resolution, with the current systems higher resolution costs more.

Unmanned aerial vehicles, which are cheaper and easier to fly, are a promising alternative. There have been issues, however, with their reliability and with the weight limitations for the payloads they can carry. Most lidar systems designed for aerial use weigh on the order of 50kg to 100kg given that they were originally designed with a fixed-wing aircraft or helicopter as the intended platform (Johnson, 2006). Although there has been considerable advancement with lidar systems for UAVs over the last several years, there has been relatively little progress made for their use with mini or micro UAVs (the Scout is classified as a micro UAV; see Table 2.1). There are lighter weight systems that have lower ranges and 2D scanning capability (Johnson, 2006), but these too tend to exceed the Scout's payload weight (400g) and dimensional (10cm x 10cm x 10cm) restrictions. However, other quadrotor systems can be equipped with lightweight 2D laser rangefinders such as the Hokuyo URG lidar system that can build a 3D occupancy grid (Morris et al., 2010). To circumvent such weight constraints, this project explores the feasibility of using a miniature video camera mounted on a small UAV to acquire 2D images and transforming them into 3D point clouds using SfM software, rather than relying on a large UAV carrying heavy laser scanners.

In this project, both urban and mining scenes were imaged. Three dimensional models of an urban environment have a variety of applications in urban planning, urban

surveillance, virtual heritage conservation and simulations relating to the urban scene (such as training for terrorism scenarios) (Frueh and Zakhor, 2003). In practice to date, this has required both aerial and terrestrial lidar data in order to completely and accurately represent the scene: the former to capture the roof detail and the latter for the façades of the buildings. Three dimensional aerial scans have certain logistical challenges to balance achieving the desired resolution while maintaining the necessary altitude. Alternatively, models can be constructed using photogrammetry. Several approaches have been used to gain the terrestrial data: some involving importing geometry from building plans, or stop-and-go methods such as 3D imaging using a stationary laser camera mounted on a tripod or photogrammetry. Another method is to use a vehicle equipped with several lidar scanners (for example: Terrapoint's TITAN system (Hefford et al., 2009) and Optech's Streetmapper) scanning on either side of the vehicle, above and below, such that the data will be acquired gradually while driving at normal speeds on the roads (Frueh and Zakhor, 2003).

In open-pit mining operations, there is a need to go beyond visual inspection and to gather quantitative data about fractures on rock walls in a time-efficient manner. This information is required to tune blasting methods and to plan for material recovery. In addition to potential accessibility issues, there are also safety issues associated with sending personnel in proximity to potentially unstable rock walls. The remote-controlled UAV imaging system developed in this project could be a solution.

Two missions were flown for the project. The first set of flights was conducted at York University. The goals of this mission were to get experience imaging urban scenes

with both man-made structures and vegetation, to test the classification software and to assess errors. The second mission was conducted at an above-ground mine site to gain experience imaging natural scenes as well as to test the fracture orientation software.

In addition to evaluating the effectiveness of the technological combinations, specific research objectives were 1) to issue recommendations on how to fly the Scout in order to obtain high-density point cloud data, 2) to comment on the strengths and limitations of Neptec's ViDAR software including performance across various surficial textures, and 3) to determine what worked well and where improvements could be made in both the classification and rock characterization software programs.

This thesis will present information about the use and evolution of UAVs, with an emphasis on the Scout UAV. Information will be presented regarding 3D data and the variety of ways that such data can be acquired. It will discuss the theory behind SfM and the specific SfM program used as well as the image processing methods. The details of both the York University and mine missions will be presented and the results of the analyses discussed. Finally, the viability of the technology combination of this project will be addressed and recommendations regarding future efforts will be offered.

## **2 Unmanned Aerial Vehicles**

### ***2.1 General***

Historically, aerial surveys have been flown with manned aircraft. A major drawback of this method is the cost: the vehicle itself is expensive, as is the time of a highly trained pilot, as well as fuel and maintenance. Over the last several years, unmanned aerial vehicles (UAVs) are becoming a more viable option (there is some contention as to the most appropriate acronym, with “unmanned aerial system” (UAS) becoming a favoured contender, but “UAV” will be used for the purposes of this paper). There are various models of UAVs designed to fulfill various needs and the designs are perpetually and rapidly evolving. The vehicle may employ a fixed wing or rotary wing design and can be remotely controlled, semi-autonomous, autonomous or some combination thereof. At the moment, UAVs are most frequently used for military purposes, but are also used in scientific research, public safety and commercial tasks (Bento, 2008).

As of yet, there is no consensus as to the classification of UAVs, but several groups have proposed reference standards; the European Association of Unmanned Vehicle Systems (EUROUVS) set out a classification system based on maximum take-off weight, maximum flight altitude, endurance, and data link range (Table 2.1) (Bento, 2008).

Table 2.1: UAV classification (adapted from Bento, 2008)

	<b>Category (acronym)</b>	<b>Maximum take-off weight (kg)</b>	<b>Maximum flight altitude (m)</b>	<b>Endurance (hour)</b>	<b>Data link range (km)</b>
Micro/Mini UAVs	Micro (MAV)	0.10	250	1	<10
	Mini	<30	150-300	<2	<10
Tactical UAVs	Close Range (CR)	150	3,000	2-4	10-30
	Short Range (SR)	200	3,000	3-6	30-70
	Medium Range (MR)	150-500	3,000 – 5,000	6-10	70-200
	Long Range (LR)	-	5,000	6-13	200-500
	Endurance (EN)	500 – 1,500	5,000 – 8,000	12-24	>500
	Medium Altitude, Long Endurance (MALE)	1,000 – 1,500	5,000 – 8,000	24-48	>500
Strategic UAVs	High Altitude, Long Endurance (HALE)	2,500 – 12,500	15,000 – 20,000	24-48	>2000
Special Task UAVs	Lethal (LET)	250	3,000-4,000	3-4	300
	Decoys (DEC)	250	50-5,000	<4	0-500
	Stratospheric (Strato)	TBD	20,000- 30,000	>48	>2000
	Exo- stratospheric (EXO)	TBD	>30,000	TBD	TBD

## ***2.2 Flight Regulations***

An important issue in conducting flights with a UAV is obtaining a Special Flight Operations Certificate (SFOC) from Transport Canada (TC), the federal agency responsible for air safety in Canada. The certificates impose many restrictions in the interest of safety; both the system and the operator are assessed. For the purposes of the application, a distinction is made between UAVs and model aircraft with the latter being mechanically driven aircraft with a maximum take-off weight (MTOW) less than 35kg used strictly for recreational purposes – launch for any other reason designates the vehicle as a UAV. While model aircraft do not require SFOCs, UAVs do. There are three application processes that can be followed: 1) for UAVs over 35kg, 2) for UAVs with a MTOW less than 35kg that are operated within visual line of sight, below 400ft in class G (“uncontrolled”: where Air Traffic Control (ATC) does not exercise authority) airspace and without autonomous navigation (this combination allows for a simplified application) and 3) for unmanned vehicles operated recreationally that exceed the 35kg weight limit and so cannot be defined as model aircraft (Transport Canada, 2008).

There are a large number of roles to be filled when operating a UAV: the operations manager who is in control of the logistics and conduct of the operation, the ground supervisor who supervises the area of operation, the pilot, the payload operator, and observers who deal with collision avoidance and monitor the vehicle and the surroundings. It must be specified in the application how each role will be fulfilled and who will be involved and responsible (Transport Canada, 2008).

Risk management is the primary concern of the application: identifying possible risks, assessing the implications and planning ways to mitigate or control each risk. The application must present information about the system (the craft and any other components), the purpose of the operation, the security plan, personnel qualifications and participants' responsibilities during the operation. A detailed plan for the operation must be included indicating the boundaries within which the operation will be conducted, the altitudes and routes to be used and the locations and heights of obstacles. Normal UAV operations must be described from pre-flight to landings. Abnormal events and emergency procedures must also be covered: anything that could potentially go wrong and how the problem would be addressed. After the application is submitted, any additional information that is requested by the TC inspector must also be provided (Transport Canada, 2008).

Initially, SFOCs are issued for specific operations. Once it has been demonstrated that the equipment can be operated safely and that the pilot can conduct operations safely, certificates are easier to obtain and can be granted for longer periods of time and without being tied to specific sites. The Aeryon Scout (discussed in the next section) has been extensively tested and proven, and long-term authority has been granted for the developers – this made the SFOCs for this project easier to obtain. The small size of the vehicle (1.3kg) also allowed for the simplified application (application process (2) described above) and reduced the safety concerns (Transport Canada, 2008).

In both applications for this project, operations were restricted to daylight hours since video was to be taken. Meteorological restrictions were also imposed in the

interest of visibility. Weather limitations specified in the Canadian Aviation Regulations (CARs) (for minimum visual meteorological conditions for visual flight rules (VFR) flights) are the default restrictions unless a case can be made to alter the limits, which was not required for this project (Transport Canada, 2008).

### ***2.3 Aeryon Scout***

The Aeryon Scout is a vertical take-off and landing (VTOL) micro UAV equipped with four rotors, which allows it to hover (Table 2.2, Figure 2.1). The system includes three components: the mission control station, the base station and the vehicle.

The control station is a tablet with a touch interface to control the vehicle and the camera. It displays the UAV's position in the environment and the live video feed from the payload camera (Figure 2.2). The control station allows for satellite map-based navigation (as seen in Figure 2.2), but in the absence of a map, one can also navigate using the LiveMap feature which georeferences a photo taken by the UAV's camera (from a high altitude) and places it on the screen as a navigational reference, or one can navigate using the camera. The navigational portion of the display (generally the map or photo) can be adjusted and reoriented to aid navigation. There are two different ways to navigate regardless of the reference system: using preselected waypoints to define a path for the vehicle to follow, or by directing the UAV to a point on the map by touching and holding at the desired location. The tablet also has a system status display that

Table 2.2: Specifications of the Aeryon Scout

Category	Micro-UAV
TC category	UAV with MTOW < 35kg
Operational range	up to 3 km
Battery	Lithium polymer
Battery voltage	12 V
Battery weight	450 g
Operational duration	up to 25 minutes
Operational velocity	up to 50 km/h (14 m/s)
Wind limit	50 km/h (gusts to 80 km/h)
Maximum altitude	500 m
Operational weight	Approx. 1.3 kg
Payload weight	up to 400 grams
Dimensions	80 cm x 80 cm x 20 cm



Figure 2.1: Aeryon Scout (bottom photo modified from [www.aeryon.com](http://www.aeryon.com))

keeps track of various operational considerations including GPS lock, battery status, and memory status. The vehicle is also programmed with smart fault handling for cases such as loss of communication, low battery and excessive wind; the treatment of these cases can be set by the user – whether to return to home position or land in place.

The base station relays the signal from the control station to the UAV allowing for long distance communication (Figure 2.3). The link between the mission control station and the base station is a short-range Wifi connection – it limits the distance between the two devices to about 5m. A radio modem supplies the link between the base station and the vehicle allowing for an operational range of up to 3km. For the purposes of downloading the data, a Wifi link can be established between the vehicle and the tablet, but this is a short-range link (approximately 5m).

The vehicle itself has only four moving parts and can be quickly assembled without need of tools. The custom battery attaches to the top of the vehicle, and the payload is secured to the bottom. The UAV navigates using GPS, but also makes use of a range of other sensor systems including sonar (used to establish altitude for heights of 2-4 m depending on the underlying surface), a pressure altimeter (for altitudes beyond the range of sonar), a temperature sensor, 3 magnetometers, 3 accelerometers, and 3 gyroscopes. Data collected during the flights are first stored on the vehicle's flash drive (2GB) to be downloaded later onto the tablet.

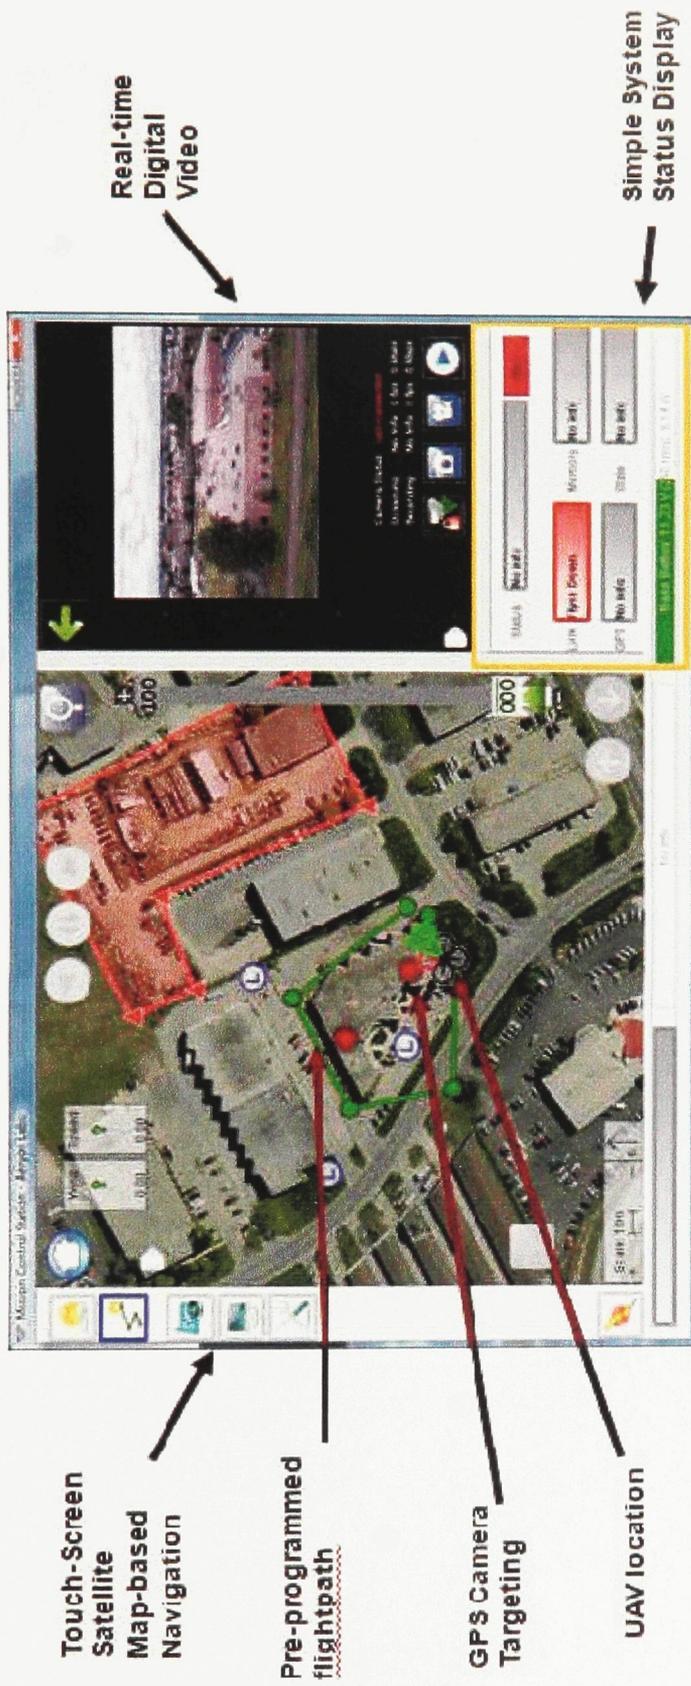


Figure 2.2: Tablet Display for the Aeryon Scout

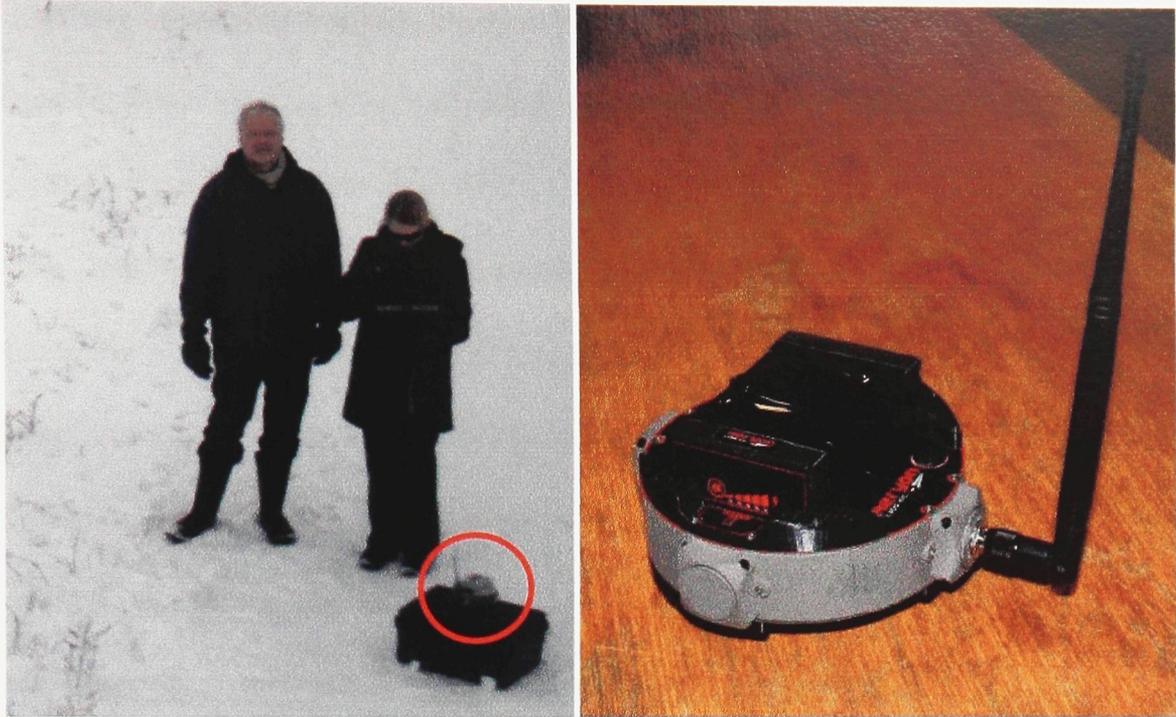


Figure 2.3: Photo of Tom Pike and Tara McLeod in the field; photo taken by the Scout in flight; the base station is circled (left). The base station used in the project is shown on the right.

There are several available payloads for the system, all plug-and-play, but the one used for the project was the daytime imaging camera (Figure 2.4). This camera is gimbaled (supported such that it will remain stable when the vehicle shifts) and weighs only 112g. The characteristics of the video images taken by the camera for this project are listed in Table 2.3. Other camera payloads offer different resolutions and fields of view, and also compensate differently for the various motions that the system may encounter in the air. The payload that was used is typically equipped with a dome to protect the camera, but, in this case, the dome was removed so that it could not cause any potential distortion to the video. The camera can be targeted on a specific point on the map, by setting the pitch and yaw of the camera or by selecting the point of interest in the field of view. The former two approaches were used as required depending on the feature being captured.

The Scout begins recording video once it is ready to take-off but imaging can be started and stopped at the push of a button. The data are usually recorded as compressed video files, but Aeryon engineers made adjustments so that less compressed data could be recovered, which are more compatible with the ViDAR software. Log files are also created during the flights recording altitude, speed, position (latitude and longitude), orientation of the camera (pitch and yaw) among other factors.

The Scout can and has been employed in many commercial applications: for image acquisition in a disaster area, search and rescue, traffic surveillance, forensic investigations and emergency response among others.

Table 2.3 : Characteristics of video images taken by the Scout for this project

Horizontal field of view (HFOV)	37°
Vertical field of view (VFOV)	29°
Output data format	*.m4v (less compressed), *.mpg
Resolution (pixels)	640x480 (total 307,200 pixels)
Frame rate	12 frames per second (fps) (for this resolution)



Figure 2.4: Original daylight camera designed for the Scout and the model used for the project (less the dome cover) (figure from [www.aeryon.com](http://www.aeryon.com)).

## ***2.4 UAVs in urban settings***

There are many arguments in favour of employing UAVs in urban surveillance, however, current flight regulations are a major hindrance. Different countries have different airspace restrictions: in Canada, an SFOC must be obtained from TC; in the United States, a Certificate of Authorization (COA) from the Federal Aviation Administration (FAA) is required but can only be obtained by government agencies and impose many additional restrictions; there are fewer restrictions in the UK (Aeryon, 2012; Sofge, 2009). Some manufacturers are pushing to have smaller UAVs classified as model craft to eliminate the need for a permit (Sofge, 2009).

UAVs would be very useful for police, particularly for smaller forces. The use of manned aircraft is usually limited to larger agencies and their application is limited again by the number of aircraft, the availability of pilots and the high cost. UAVs have the potential to multiply the force at a substantially lower cost than manned aircraft (Murphy and Cycon, 1998). Police in Merseyside County, Britain are testing a UAV equipped with a camera. The 20 minute lifespan of the battery for the model limits its use to specific missions such as monitoring traffic jams or recording video evidence rather than long-term patrols (Sofge, 2009). Demonstrations at the McKenna Military Operation in Urban Terrain (MOUT) training site have displayed the usefulness of a UAV in surveillance and reconnaissance operations, in delivering non-lethal payloads (smoke canisters, steel spikes for tires and leaflets) and search and tracking of man-sized targets. UAVs could also be employed for traffic surveillance, for neighbourhood patrol,

as communication relays and as chemical sensors. VTOL craft also have the advantage of being able to take-off and land in small areas and can place sensors for long-term surveillance (Murphy and Cycon, 1998). UAVs could also be useful in fire management, investigation and damage assessment to track people and hotspots during the blaze and to investigate the scene afterward without endangering personnel and to collect evidence without contaminating the scene (Aeryon, 2012).

Uses in urban environments are not limited to police and other government agency programs. UAVs can be used in environmental sensing (as at smoke stacks), as inspection systems providing a safe and easy way to take a closer look, in aerial photography (potentially in real estate or promotional events) and private security (Aeryon, 2012).

## ***2.5 UAVs in Earth Sciences***

UAVs have many applications in earth sciences including high-altitude atmospheric sampling, monitoring wildfires (Wegener et al., 2004), glaciological studies (Runge et al., 2007; Whitehead, 2010), pipeline monitoring, digital terrain mapping (Barnard, 2008), and magnetic and electromagnetic geophysical surveys (Barnard, 2008; Samson et al., 2010; Anderson and Pita, 2005). With the lower operating costs surveys can be flown with greater frequency and several UAVs can be flown simultaneously for accuracy, efficiency or fault tolerance. Additional benefits of UAVs are endurance, the ability to fly at night, less disturbance to certain measured fields and the ability to fly lower and

slower which allows, in many instances, for better data quality than manned flights (Barnard, 2008; Caron, 2011; Caron et al., 2011).

Using UAVs also permits surveys in remote areas. For example, they are particularly useful in climate studies in polar regions where they can make regular measurements to monitor mass-balance of glaciers, retreat rates and changing flow dynamics. There are several other remote areas that scientists would like to survey that are more than 1000 nautical miles (more than 1800 km) from landing sites for wheeled aircraft. Accessing such locations from the landing sites requires endurance greater than 10 hours or the ability to refuel in remote areas (Runge et al., 2007) – both problems are alleviated through the use of UAVs which can be launched in more remote areas or from the deck of a science vessel to help access the desired area (Wegener, et al., 2004). UAVs can be employed on short notice to deal with rapid changes and have the capacity to make multiple flights in a single field season (Whitehead, 2010).

The current optimal altitude for geophysical surveys is approximately 80m which pushes the safe limits for manned aircraft. There is also a call for surveys at lower altitudes of 30 to 50m (Anderson and Pita, 2005). UAVs can fly lower than manned aircraft can safely manage and more slowly thus producing at least helicopter quality data, but at a lower cost. Also, in the case of magnetic and gravity surveys, UAVs offer less disturbance to the measured parameters, although the equipment is closer to electrical noise in a UAV (Barnard, 2008). The GeoSurv II from Carleton University (Samson et al., 2010; Caron, 2011; Caron et al., 2011), the Venturer from Universal

Wings/Stratus Aeronautics<sup>3</sup>, and the GeoRanger from Fugro Airborne Surveys<sup>4</sup> were developed to execute geophysical surveys. The GeoRanger system was developed to complement Fugro's existing fleet of manned aircraft for magnetic surveys, providing alternatives for data acquisition and can operate in harsh or remote areas (including offshore) and at low altitudes.

---

<sup>3</sup> <http://stratus-aero.com/>

<sup>4</sup> <http://www.fugroairborne.com/>

## 3 Imaging in three dimensions

### 3.1 Imaging methods

Three dimensional images usually begin as point clouds: collections of unstructured points in 3D coordinate space. Different data collection systems derive this information differently and produce point clouds of varying densities and resolutions.

Systems for obtaining three dimensional (3D) images are one of two basic types: triangulation systems and time-of-flight (ToF) systems (the latter of which is more commonly referred to as lidar). Triangulation-based laser cameras calculate distance from the geometry of the projected and reflected laser beams: the reflected beam is focused on a linear detector array using mirrors and the specific location of the beam on that array allows for the determination of distance. The ToF measurements can be made directly by determining the two-way travel time of reflected laser pulses, or indirectly by determining their phase shift. Triangulation allows for higher precision at close to mid-range (< 5m), and reaches the limits of its effectiveness at ranges of 10 to 20m. ToF lasers can be applied at far greater ranges, even up to hundreds of kilometers, as long as the reflected signal is strong enough for detection. The accuracy of the ToF systems is relatively consistent, on the order of centimeters to tens of centimeters, and does not vary with the range to the target – precision is largely limited by the spot size of the pulse (English et al., 2005).

Both approaches can be further subdivided into methods with different strengths and weaknesses in terms of resolution, accuracy and precision. There are both single point and slit triangulation scanners. The former has higher resolution and accuracy,

but the latter is the more widely used model since it is simpler and faster than single point scanners, but still provides range accuracy on the order of tens of micrometers (up to hundreds of micrometers at the farther end of the useable range ( $\approx 5$  m)) (Blais, 2004). ToF systems can use single pulse or pulse trains which use the time to determine the range, or amplitude or frequency modulated constant waves, which both use phase shift. These last two are faster than the triangulation systems, also operate at shorter ranges, and provide millimeter accuracy (English et al., 2005).

Another method of obtaining 3D data is using stereovision. This method uses stereoscopic parallax – the difference in the appearance of an object due to differing perspectives, with the difference being greater for nearer objects. Stereo pairs of images are used to select points to be used in the construction of a point cloud (photogrammetry represents the full development of this technique) (Campbell and Wynne, 2011). One weakness of such a system is that the number and location of measured points is determined by the features that the system can track, rather than a scanning pattern, which makes the distribution unpredictable, ergo fewer or sparser features that can be successfully identified in subsequent images leads to a sparser point cloud (English et al., 2005).

### ***3.2 Structure from Motion theory***

Many different properties of a scene impart 3D information for a biological visual system: contours, shading, size, texture, shadow, binocular stereo and motion. Many of these are based on experience, but the stereo and motion cues are the most important

for depth perception because they require fewer assumptions than the other cues, which is also better for computer vision. Moving through an environment, projected points move relative to one another. This is motion parallax (basically stereoscopic parallax but with less disparity between images and possibly more than a single transformation between images), which is what SfM seeks to exploit (Hogervorst, 1966; Trucco and Verri, 1998).

The field of computer vision has spent years working on the problem of extracting structure from motion (SfM). The problem represents an attempt to reverse the standard 3D-to-2D image-formation process (such as the 3D world creating 2D images on a person's retinas). The creation of 2D video from a 3D scene involves the coordinates of the features in 3D space, the relative motion in 3D space between the camera and the scene and the camera's internal geometry; the reversal attempts to reconstruct these factors from the captured video (Figure 3.1). This requires either exploiting additional information about the scene (useful in the case of geometric models) or using corresponding points in multiple views and triangulating to find the 3D coordinates (Robertson & Cipolla, 2010). The program used in this project uses the latter approach.

The fundamental basis for SfM was established in 1913 when Kruppa proved that from 2 views of five 3D points, one could determine the positions of the points in 3D space (less the scale factor) as well as the relative position and orientation of the camera (Robertson & Cipolla, 2010). The field of computer vision has made substantial strides in software that exploits this point.

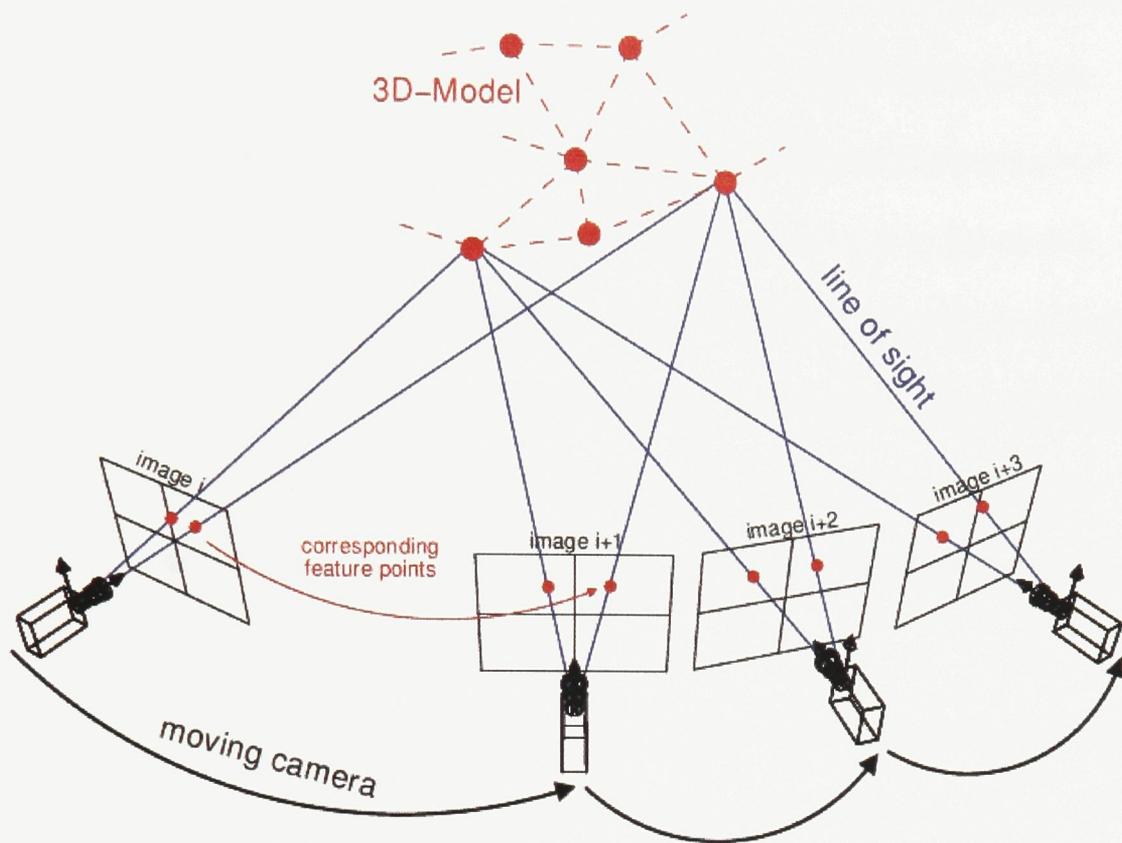


Figure 3.1: Basic principle of Structure from Motion (from Ackerman, 2010)

In the pinhole projection model (a good approximation for the behavior of most cameras), the relationship between the 3D point and the corresponding 2D point involves three basic transformations: 1) a rigid body 3D to 3D transformation that relates the point in the world coordinate system to the same point in the camera coordinate system (and thus describes the camera pose); 2) the 3D to 2D transformation that relates the coordinates in the camera coordinate system (3D) to the coordinates in the camera image plane (2D); and 3) the 2D to 2D transformation that relates the coordinates in the camera image plane to pixel coordinates (Figure 3.2) (Robertson & Cipolla, 2010).

The first issue with a SfM program is to solve the correspondence problem: correctly identifying the same point in successive video frames. The solution can take an interactive approach (i.e. manually identifying like points), which has the obvious disadvantage of being time consuming, or an automatic approach. The latter has proven to be a difficult task since no two pixels are exactly alike from one scene to the next. It would be impossible to compare each pixel in a picture with each pixel in the following picture, and impractical given that not all points are equally suited to matching. Instead, feature matching detects salient points of interest – the zones around such points display intensity variations that make the points easier to identify (Robertson & Cipolla, 2010) (see appendix A for more information).

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 and the ability of the program to extract good features to be tracked in each frame - feature points are more

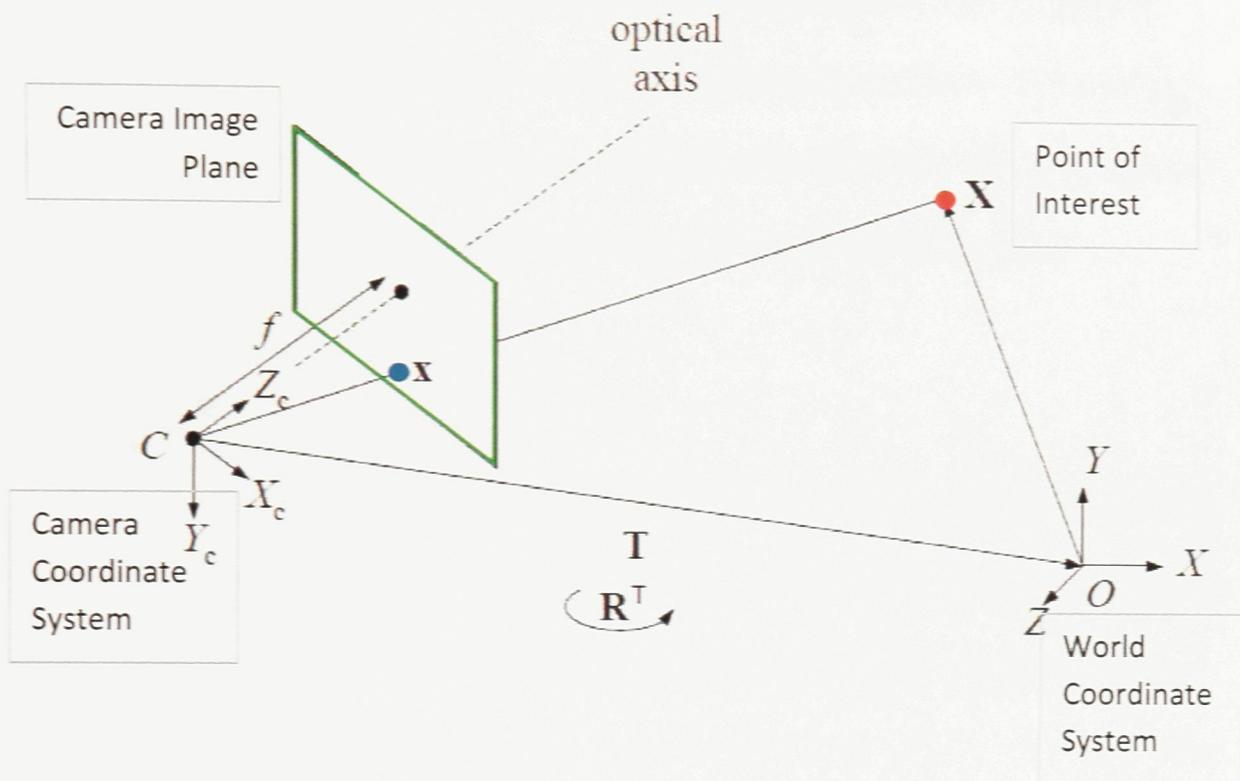


Figure 3.2: Pinhole projection (in blue) of a 3D point (in red) onto a camera image plane (green parallelogram). “R” and “T” are the rigid body transformation between the world coordinate system and the camera coordinate system,  $f$  is the focal length. In a real camera system, the camera image plane would be behind the camera and inverted (modified from Robertson and Cipolla, 2010).

abundant and easier to track in scenes with more texture (a rock wall is an example of a good scene providing plenty of contrasts) than more uniform surfaces (a man-made cement wall would be difficult to reconstruct). Also, although the relative distances between the various features in the final cloud are internally consistent, SfM programs cannot offer a scale factor for a scene, and although intensity is key in identifying points and is used in constructing the point cloud, ViDAR does not provide intensity information about the scene.

### ***3.3 ViDAR software***

The particular SfM software used in this project is the ViDAR software developed by Neptec Design Group. The traditional SfM approach uses image pairs, but ViDAR always draws its information from a trio of images in order to improve the robustness of the correspondence between feature points. Once corresponding feature points have been tracked through three frames, the location of the 3D point in space is resolved using triangulation – the intersection of imaginary lines projecting from the camera to discrete feature points in 3D space. For triangulation to work there must be a certain level of disparity between the video frames that are used. As such, when the camera is stationary and/or panning, the lines from the consecutive frames are parallel, providing infinite, inconclusive intersections. Ideally, the 3D point lies at the intersection of the projection lines through the measured pixel positions (recall from §3.2 the camera image plane and the conversion to pixel coordinates). However, due to measurement noise, the lines may not intersect. As such, the goal becomes to minimize the sum

of squared errors between the measured pixel position and a predicted pixel position. This will offer a more accurate solution than relying solely on noisy observations (Robertson & Cipolla, 2010).

ViDAR allows the user to select a number of frames to skip in the analysis to speed up processing and avoid having to analyze and reject too many frames due to insufficient displacement. The software also automatically detects when a frame has achieved the required disparity from the previous one and adds it to the 3D estimation. In this project, every other frame was used despite the low velocity of the vehicle in the field in order to avoid skipping good, usable frames. The 3D point cloud output by ViDAR is built incrementally with the addition of each successive video frame included in the analysis (Figure 3.3).

Figure 3.4 is a detailed flowchart of feature detection and tracking in ViDAR. The saved “tracked points” in Figure 3.5 (bottom of the first column) is the input at the circles in Figure 3.4 (except for the first input image which would be from the first set of saved “feature points”). In ViDAR, feature points are detected and matched between frames using an iterative pyramidal implementation of the Kanade-Lucas-Tomasi (KLT) tracking algorithm (Bouquet, 2001) (Figure 3.5). Since the standard KLT algorithm is only valid when the pixel displacement is small, a pyramidal system is used. The image at each pyramidal level is constructed to have half the resolution of the previous level, so if the original image,  $I^0$ , has a resolution of 640x480, then the resolution of the image at the next level,  $I^1$ , would be 320x240,  $I^2$  would have a resolution of 160x120 and so on.

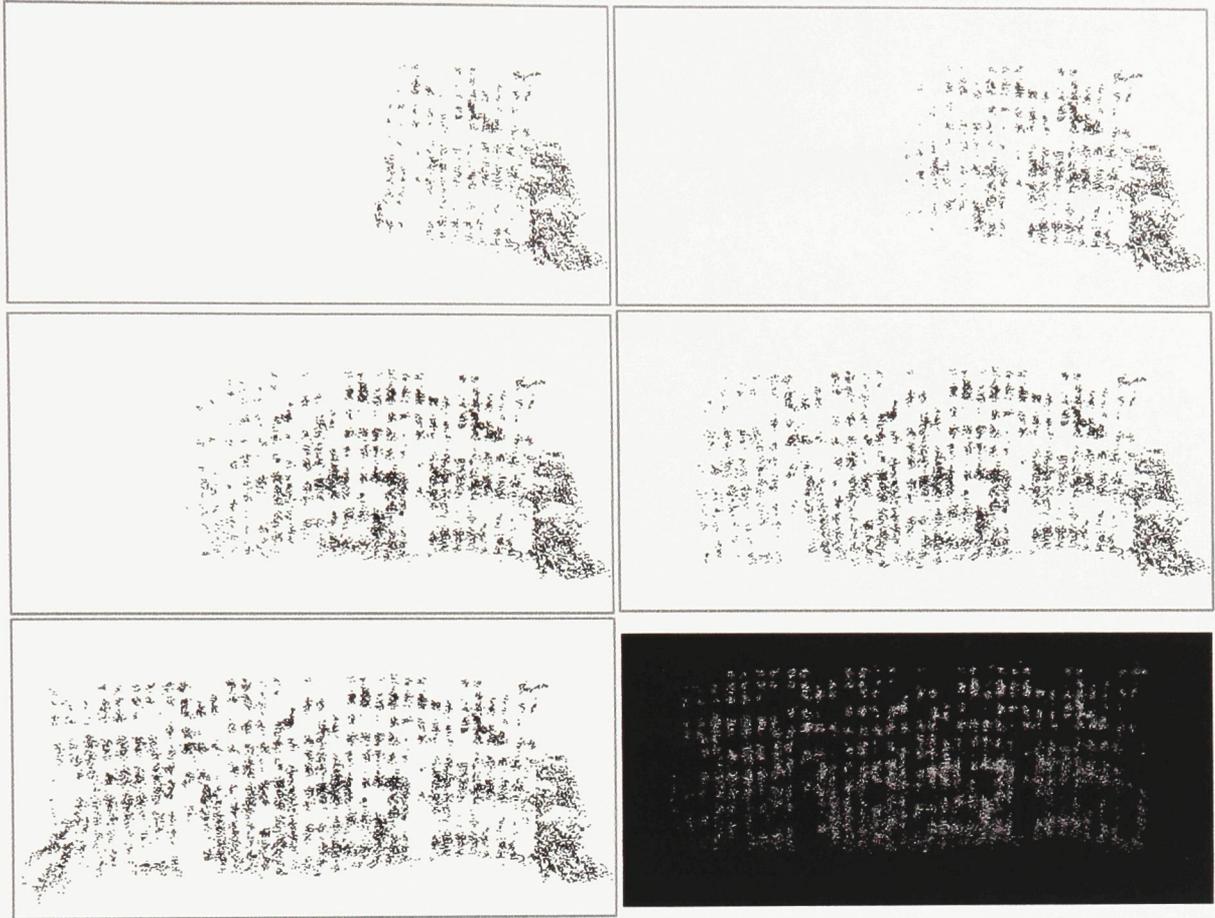


Figure 3.3: Example of building a point cloud in ViDAR using the west face of the Life Science building. The top left figure shows the point cloud using 4 camera images, then 8, 12, 16 and 19 images. The final point cloud was constructed from 20 images and had 9737 points. The bottom right figure shows the point cloud with the ground removed (7612 points).

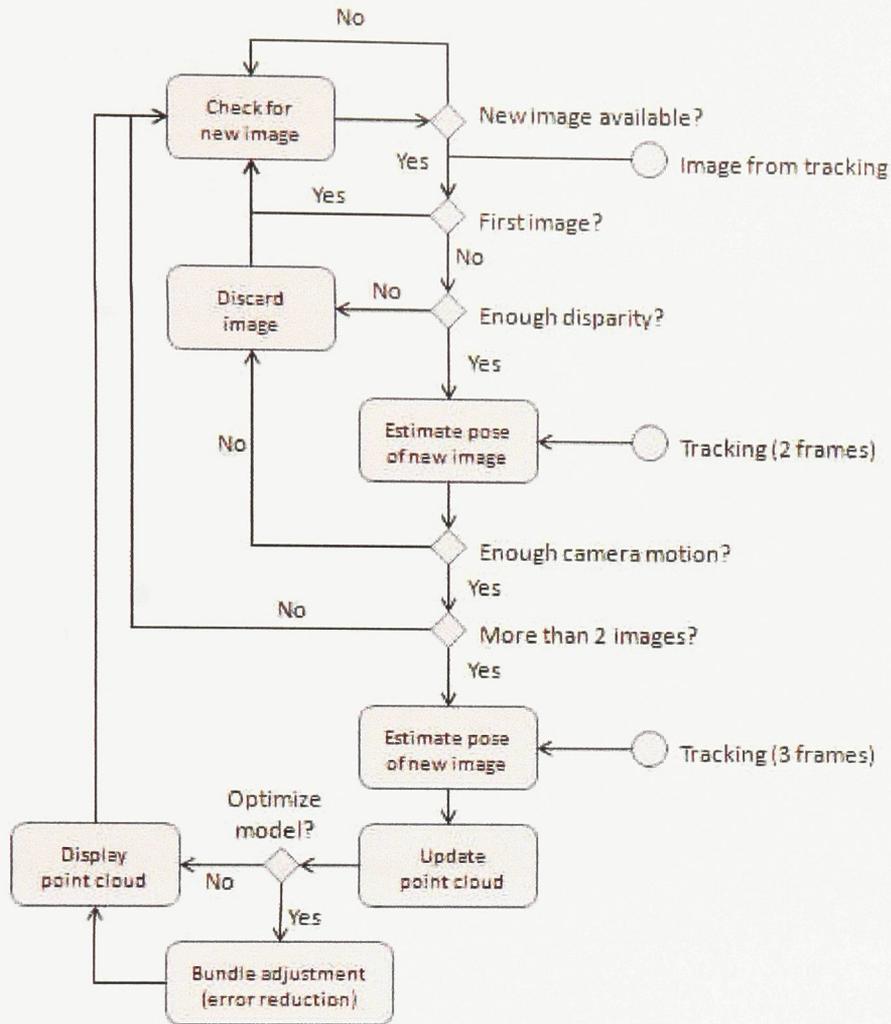


Figure 3.4: ViDAR algorithm. The circles represent the three input frames and the KLT tracking represented in Figure 3.5 (Source: Neptec Design Group).

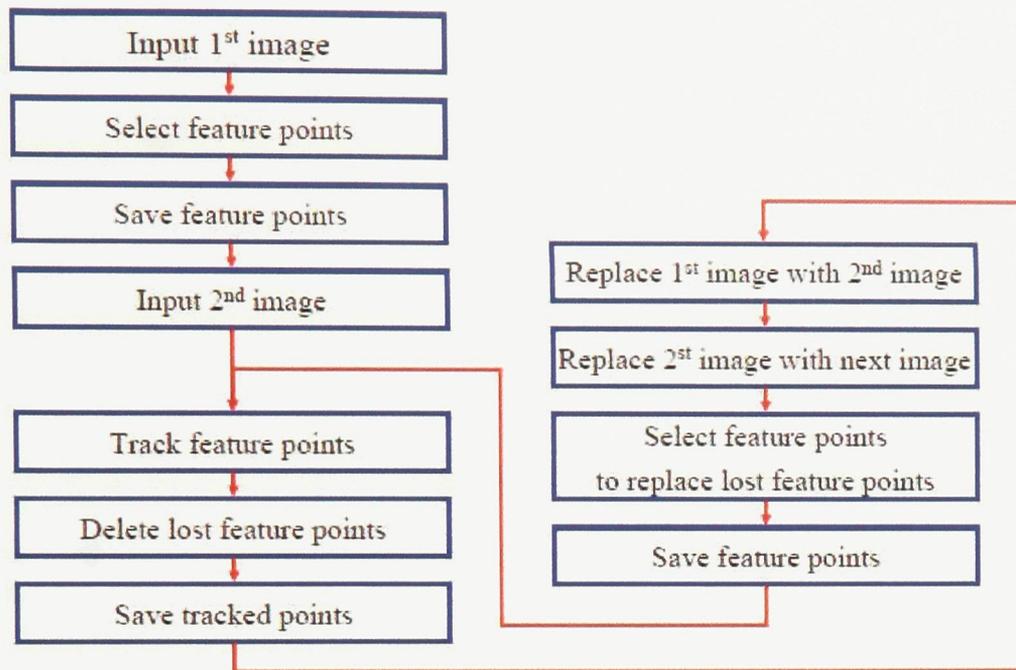


Figure 3.5: KLT tracking algorithm. This example uses two frames for tracking points whereas ViDAR uses three frames (from Suhr, 2009)

The “height” of the pyramid is picked heuristically to properly accommodate the optical flow (amount of motion), but most cases are satisfied with a maximum of 4 levels. The analysis begins at the lowest resolution (deepest pyramidal level) to determine the amount of displacement between the frames and those results are used as a starting estimate of the displacement in the next level up. The estimate is refined with the higher resolution image and the new optical flow vector is the starting point for the next level up and so on with each successive level working up the pyramid ( $I^x, I^{x-1}, \dots, I^1, I^0$ ). This process continues until the highest resolution is reached ( $I^0$ ). Using this method, the only value that needs to be computed at each level is the residual difference between estimates which is calculated after the estimated correction is applied (Bouguet, 2001). A RANdom SAmples Consensus (RANSAC) algorithm addresses any false matches by estimating the fundamental matrix between two images (the fundamental matrix relates corresponding points). RANSAC is a resampling algorithm. Rather than using a large portion of the data to formulate a solution and removing the outliers, RANSAC works by using the minimum number of data points required to estimate the model parameters and then determining the number of points that fall within a predefined error (Derpanis, 2010).

The final stage of SfM processing involves using a bundle adjustment optimization to refine the model (Figure 3.4). Once the previous processes return an estimate of the 3D points, bundle adjustment seeks to minimize the errors (Robertson & Cipolla, 2010). The post-bundle-adjustment product is then the final best estimate of the 3D scene.

## **4 Image Processing**

### ***4.1 Principal Component Analysis***

Principal Component Analysis (PCA) was used on both sets of data: to evaluate the planes in the urban data and in the rock mass characterization software for the mine data.

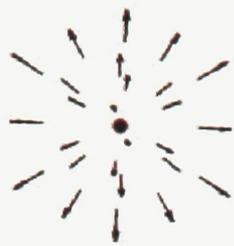
PCA works by determining the eigenvalues and eigenvectors of the covariance matrix associated with the input data. The eigenvectors are used to define a new reference system characterized by the patterns in the data allowing the data to be expressed in terms of these axes rather than the standard x, y, and z axes (all eigenvectors are orthogonal). The eigenvalues determine the significance of each detected pattern. For example, in 2D data, one eigenvector would define a line of best fit (this would have the higher eigenvalue) and the second eigenvector would be perpendicular to the first, preserving the dimensionality of the data. In 3D data from a plane, the two eigenvectors with the larger eigenvalues define the plane and the third would be normal to that plane (Smith, 2002).

### ***4.2 Classification software***

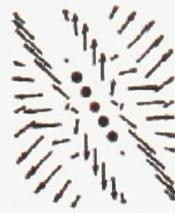
The classification software developed by John Harrison (personal communication) at McGill University was designed to distinguish between points that are singular, linear or planar. The first step is to voxelize the data, that is, to subdivide the full 3D space into smaller cubes each of which will be defined as occupied by a point (or points) from the point cloud or unoccupied. In defining the size of the voxels, the goal is to define a voxel size that is not so small as to require excessive processing time, but also not so large as

to compromise the density of the data. The classification program uses an approximation of the Zucker-Hummel (ZH) operator, but the principle is the same: the ZH operator first smoothes (blurs) the voxels to reduce noise and detail, and then computes normals in 3D space (Figure 4.1). Each point in the point cloud contributes to the calculation of normals in adjacent voxels within a specified radius; each point can produce many normals. Once the normals have been determined, principal component analysis (PCA – see next section) can be used to combine the many values into a single normal for a point.

The PCA works with the list of normals for the points within the specified radius of a point of interest and determines the three orthogonal axes that define the frequency of the orientations of the normals ( $p_1, p_2, p_3$ ):  $p_1$  corresponding with the most common direction and  $p_3$  corresponding with the least common direction. Eigenvalues are also produced for each of the three directions. If the normals are primarily oriented in one direction, the first eigenvalue (for  $p_1$ ) will be the strongest, indicating a plane. However, if normals are equally common in all directions, the eigenvalues will all be approximately the same and the point will be classified as singular. For a line, two of the eigenvalues will be stronger than the last (Figure 4.1). At the completion of the analysis, all the points in the point cloud are classified as singular, linear or planar (they can also be classified as “other” where their association is unclear).



(a) Single Point



(b) Line



(c) Plane

Figure 4.1: Normals obtained using a ZH operator on simple point clouds (Harrison, 2009; unpublished)

### ***4.3 Rock mass characterization software and associated programs***

The rock mass characterization software developed by Mah et al. (2011) was used to compute fracture orientation at several locations on the rock walls at the Canadian Wollastonite Mine. The user first selects the location of interest in the 3D point cloud and defines a subset of data by choosing an appropriate radius around that location. A best-fitting plane is computed for the subset of data using PCA. The orientation of the normal vector to that plane can be decomposed into two components: strike (orientation with respect to north) and dip (inclination angle from the horizontal).

Due to some practical difficulties, additional programs needed to be used. None of the visualization programs displayed the  $(x, y, z)$  coordinates of a point in the 3D space – a necessary input for the rock mass characterization program. The solution was to write a program in MATLAB which displayed the point cloud projected onto the three different Cartesian planes (the  $x$ - $y$ ,  $x$ - $z$ , and  $y$ - $z$  planes) since the original point cloud itself proved too large to easily maneuver. The low concentration of the point cloud and associated difficulty in identifying features necessitated the use of an additional visualization program to track the desired point through the different orientations and thus make a reliable estimate of the coordinates of the centre of the face. The results of each successive analysis (strike/dip values and figures) were used to narrow down the best  $(x, y)$  coordinates for the central point and the most effective radius in order that as much of the surface is captured without including any edges.

#### ***4.4 Visualization software***

The visualization software constructs a polygonal mesh from the point cloud using the points as vertices of triangles. The program then determines the normals for each triangle and from these it reconstructs the strike and dip of each triangle and displays the results using HSV (hue, saturation, value) (developed by Po Lai, M.Sc. candidate at Carleton University).

Given the large scatter of the data in this study, a smoothing algorithm was included in the processing that uses a geometric Laplacian operator developed by Karni and Gotsman (2000). In Laplacian smoothing, the vertices are incrementally moved to the final smoothed position. This approach does not add or remove any vertices and only information from the nearest vertices is used in determining the smoothing for a given vertex. The weighted sum of the immediately adjacent points is calculated to set the new position of a given vertex. The differences between various Laplacian methods lie in how the weights are determined (Bray, 2004).

## **5 Missions**

### ***5.1 Familiarization***

Before any of the missions could be conducted, I had to take a short two day course to familiarize myself with the Scout and its operation. The training was in late January 2011 in Waterloo, Ontario at Aeryon Labs and their permitted field sites (Aeryon has been granted long-term SFOCs for training and experimental purposes). While in the field, I flew two different flight patterns in an attempt to test a few potential data collection methods on different targets. Both efforts took advantage of the waypoint mode of the Scout: one was flying a loop around a small target (a snow-covered gravel mound) and the other was flying parallel strips in and near a parking lot with trees and a baseball diamond. Video was collected for both patterns, but processing afterward was impeded by snow obscuring many of the features (see §6.1).

### ***5.2 York Mission***

The first flights in the project were to be conducted in an urban setting – York University’s Keele campus was chosen due to the availability of a 3D model of the campus that has been under development for the last few years through the efforts of those working in the Computer Vision Department<sup>5</sup>. The goal of the mission was to obtain point cloud data that could be compared to the existing models generated by lidar and perhaps complement them. Scenes with a juxtaposition of man-made

---

<sup>5</sup> <http://icampus.apps01.yorku.ca/>

structures and natural features were also collected for use with the classification software that was designed to distinguish between the two; the software was originally developed for lidar data.

### **5.2.1 Special Flight Operations Certificate**

Two difficulties in obtaining the SFOC for York were that the campus is within 3 nautical miles (NM) (5.56km) of an airfield (Downsview Airport is approximately 1.4NM (2.59km) from the campus) and that the operations area was populated. There are always concerns in operating UAVs regarding the risks to any manned craft in the area (hence the greater concern near an airfield) as well as to people on the ground, in cars or even in buildings. The concerns regarding the proximity to the airport were alleviated by restricting the maximum altitude to less than the height of the tallest building on the campus (19 storeys at approximately 3m per storey provided a maximum height of about 60m), which did not negatively impact operations. It was the concern about risks to people that required the most accommodation. The team needed to ensure that there were never any people within a 10m lateral radius of the UAV's position whether it was on the ground or airborne (the approximate maximum radius debris would travel in the event of a crash). As such, part of the team was dedicated to keeping the projected flight path clear of pedestrians. These restrictions also limited the speed of the UAV to approximately walking pace (about 5km/h) so that those responsible for keeping the way clear would not be overtaken by the craft.

The SFOC application was submitted on 29 June 2011 and granted on 7 July 2011 for a time window between 7 July 2011 and 22 July 2011.

### **5.2.2 Survey site**

The York University Keele campus covers about 1.5 km<sup>2</sup> and consists of some 86 buildings (Figure 5.1). The missions were concentrated in the areas that would fulfill our multiple objectives: areas that had been well and thoroughly imaged, areas where only basic imaging had been completed (basic structure without details of the façades), areas that had not been imaged, and areas that had the juxtaposition of natural and manmade objects. The building façades we surveyed included brickwork, concrete faces and banks of windows.

### **5.2.3 Flights**

Beside safety, there were several other considerations in the flight planning. The altitude at which the flights were conducted had to take into account the angle of incidence of the camera, the heights of the buildings surrounding the target and, in the case of the roofs, the availability of reference points.

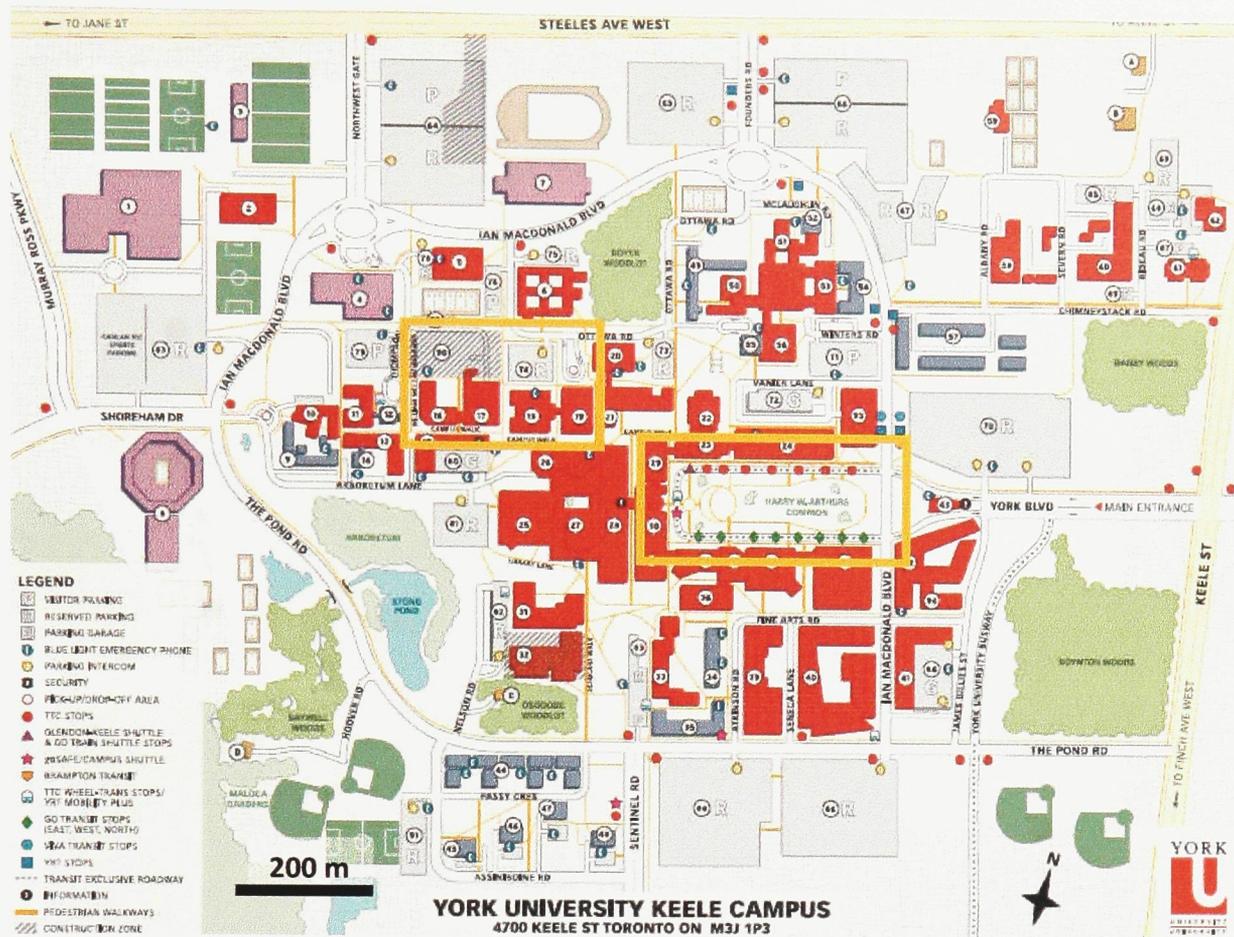


Figure 5.1: Map of York University Keele campus with boxed areas indicating flight areas.

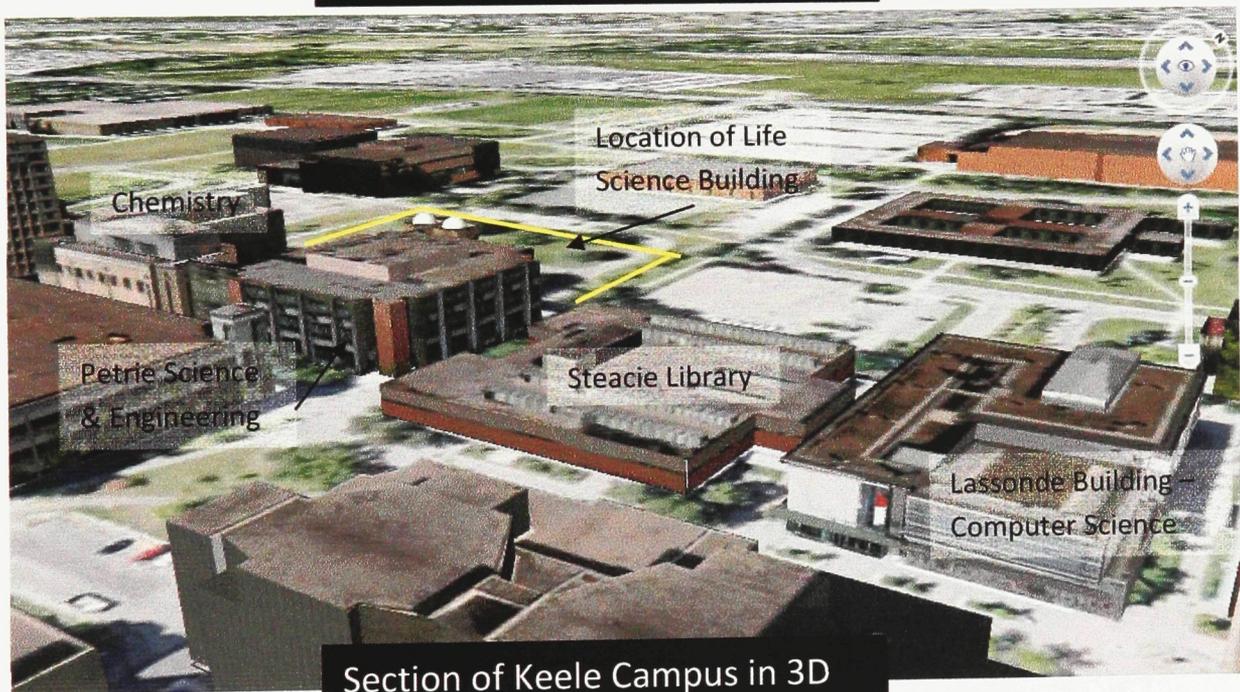
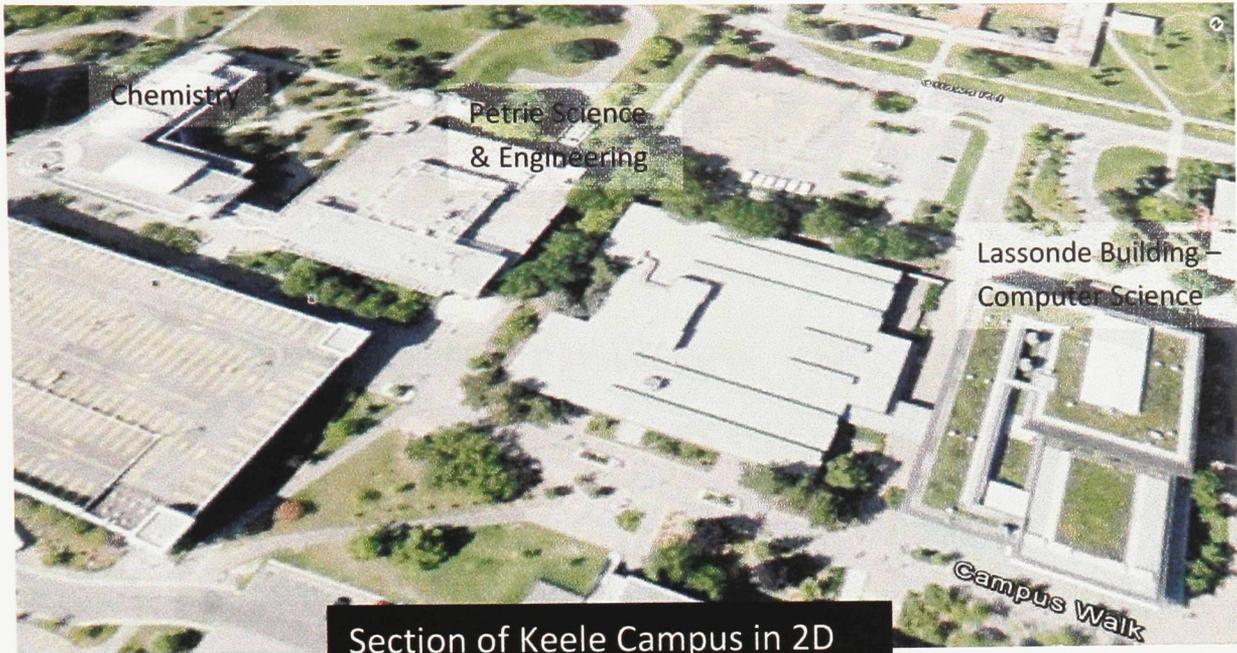


Figure 5.2: Area of York University's Keele campus corresponding to the smaller box in Figure 5.1.

Given that ViDAR is supposed to work best when the angle of incidence to the target is high, a relatively low altitude was required when imaging building façades. It was also prudent to capture as much of the height of the façade in one pass as possible to minimize the potential errors with the later processing. This required that the UAV maintain a distance of several meters from the target. However, the adjacent buildings also had to be taken into account, requiring a higher flight altitude, which resulted in a lower angle of incidence: the taller the adjacent building to be over-flown, the lower the angle of incidence in capturing the face. As a result of these considerations, the flights to image the façades of the selected buildings were conducted at altitudes of about 20 to 40 m, adjusting the angle of the camera such that the full height of the building (or as much as possible) was in frame. To capture the roofs of the buildings, the maximum allowable altitude was exploited (50 to 60 m) in order that reference points could be maintained given the general featurelessness of most of the roof surface of each building used. For these scans, it was possible to provide the desired high angle of incidence – the camera was angled plumb to the roof.

All the flights at York (Table 5.1) were conducted with the aid of a georeferenced map of the campus, albeit slightly out of date (the more recent construction was not represented, including one of the buildings used in the study, the life science building) (Figure 5.2). This made maneuvering around the buildings and predicting encounters with obstacles (buildings) a much simpler task. Also, the UAV was guided manually due to the need to make adjustments to the flight path for obstacles, to keep the target in frame and in case a pause was required to clear the way.

Table 5.1 : York Mission Statistics

<b>Date</b>	12 July 2011	13 July 2011	14 July 2011
<b>Weather</b>	Daily high: 33.1°C Max wind: 23km/h WNW Gusts: 26km/h Clouds: 6%	Daily high: 26.8°C Max wind: 26km/h NNW Gusts: 31km/h Clouds: 11%	Daily high: 26.9°C Max wind: 6km/h NNE Gusts: 6km/h Clouds: 2%
<b>Number of flights</b>	9	15	4
<b>Length of video recordings (min:sec)</b>	30:28	23:35	19:01

## **5.3 Mine Mission**

For the more geologically oriented aspects of the project, we sought a site that would provide clean jointing in an above-ground, open-pit mine type setting. The Canadian Wollastonite Mine<sup>6</sup> in Seeley's Bay, Ontario (near Kingston) – which is currently in the early stages of development – boasted an exploration trench which suited our purposes. The goal with the data derived from the flights was to determine if software designed to automatically determine joint orientation (Mah et al., 2011; Mah et al., 2012) could be successfully applied to the resulting point clouds.

### **5.3.1 Special Flight Operations Certificate**

The restrictions for flying at this site were somewhat less stringent, largely due to the fact that the only people present were those few directly involved in the operation and that the nearest airfield was the Ganonoque Aerodrome 3.5 NM (6.48km) away. This significantly reduced the safety concerns with respect to people and air traffic.

The SFOC application was submitted on 25 July 2011 and granted on 23 August 2011 for a time window between 23 August 2011 and 1 October 2011.

### **5.3.2 Survey Site**

Wollastonite ( $\text{CaSiO}_3$ ) is a calcium inosilicate mineral containing approximately equal parts of  $\text{CaO}$  and  $\text{SiO}_2$ . It is formed by subjecting impure limestone or dolostone to high

---

<sup>6</sup> [www.canadianwollastonite.com](http://www.canadianwollastonite.com)

temperature and pressure. It is usually white, although it may have a cream, grey or pale green tint and it rarely occurs on its own.

Most of the world's wollastonite comes from eastern North America, with the Grenville Province at the centre of attention. The St. Lawrence Wollastonite Deposit, most of which is on the Canadian Wollastonite Property, is a large wollastonite skarn with approximately 41.3% wollastonite, but up to 80% locally (Grammatikopoulos, et al, 2003). The deposit is located within a horseshoe shaped band of quartzite that opens to the east. The area consists of strongly folded, interbanded calc-silicates and silicate skarn layers. The metasedimentary rocks strike northeast and dip to the southeast and the plunge of the large-scale folds changes from northeast to southwest (Canadian Wollastonite, 2008; Grammatikopoulos, et al, 2003). There are two major map units in the deposit: wollastonite-dominated (>40% wollastonite) and clinopyroxene-dominated (>40% clinopyroxene) skarns. Other minor units include wollastonite-clinopyroxene skarn, feldspathic skarn, and mixed skarn. The main skarn units are all interlayered with each other and there are sharp contacts between the various units. The wollastonite skarn is only found in the inner portion of the fold. The layers of the unit range from a few centimetres to several metres in width and can be pale grey or pale green depending on the clinopyroxene content. The wollastonite is mostly elongate or prismatic, but there are instances of fibrous or radiating habit (Grammatikopoulos, et al, 2003).

The protolith of the wollastonite was interbanded sandstone and impure limestone that was deposited in an inter-arc basin c. 1200-1300Ma. The Grenville orogeny (c.

1200-1000Ma) was responsible for the granulite facies metamorphism in the area. The development of wollastonite is ascribed to metasomatism by magmatogenic fluids. The Grenville orogeny boasted multiple phases of plutonism from c. 1200 to 950Ma, keeping the temperatures in the area high for the duration. The metasomatism to create wollastonite occurred late in the orogeny when there was a reduction in pressure, but still a high temperature from these intrusive pulses (c. 1100-1000Ma) (Canadian Wollastonite, 2008). The hydrothermal fluid was water dominant, silica-rich and CO<sub>2</sub>-poor (Grammatikopoulos, et al, 2003).

Wollastonite is used in several industrial applications, but has six main applications: in ceramics, construction products, paints and coatings, metallurgy, plastics and friction materials. It is used in ceramics including bricks, pipes, tiles, porcelain and advanced ceramics for greater strength, shorter drying times, reduced expansion when heating and improved surface integrity and appearance. In construction products (concrete, mortar, caulking and sealants to name a few), wollastonite improves weather resistance, resistance to mildew and helps prevent surface and structural damage. For various paints and coatings, wollastonite increases resistance to weathering, corrosion, cracking and chipping and provides better durability and colour fastness. In metallurgy, particularly in the steel industry, wollastonite improves the continuous casting process offering more uniform heat distribution, better absorption of non-metallic inclusions, and other benefits. In plastics, the mineral offers improved strength, stiffness, stability, durability and surface quality, and also reduces susceptibility to heat distortion.

Wollastonite also reinforces fibers and fillers in friction materials replacing asbestos, glass fibers and synthetic materials (Canadian Wollastonite, 2008).

The trench where the flights were performed was oriented approximately north-south, bordered on the west and east by walls of wollastonite about 7-8m high at the highest points. The walls featured several fracture systems, resulting in a stepped appearance, and there was also some vegetation growing from some of the benches and crevices (see Appendix B for a complete photographic record of the survey site). Each wall was about 190m long.

Manual measurements of strike and dip were taken of joint faces in order to validate the results of analysis of the point clouds. A total of 86 manual measurements were taken on the two walls: 42 fractures on the east wall and 44 on the west wall (see Appendix B). Also, two different users took measurements. The joint surfaces were quite rough, with an almost ridge-like pattern on the subvertical faces which complicated measurements. Additionally, many of the sub-horizontal surfaces needed to be cleared of debris or vegetation before they could be evaluated. The measurements defined three joint sets with orientations of (strike/dip in degrees): 329/85 (joint set #1), 183/8 (joint set #2) and 32/78 (joint set #3) (Figure 5.3).

Since SfM programs cannot provide scale, the site was also measured with a tape measure (see Appendix B) to establish scale for later processing.

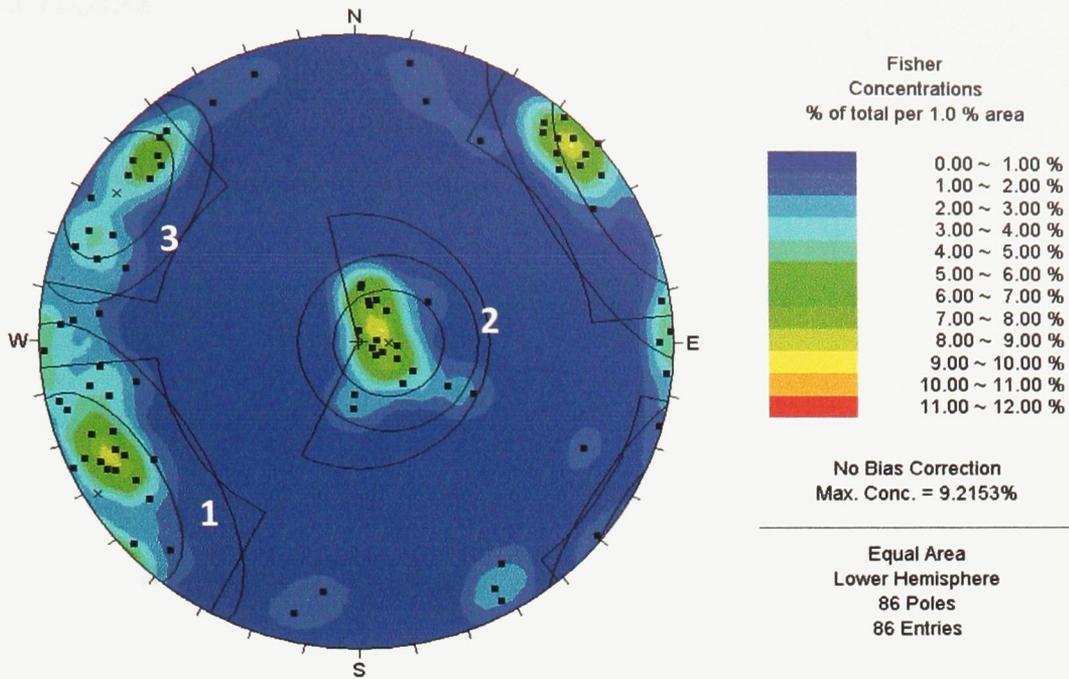


Figure 5.3: Stereonet showing the manual measurements sectioned into joint sets (with a few outliers); concentric circles around the average orientation are for one and two standard deviations.

### 5.3.3 Flights

Flights were conducted in September 2011 to survey the walls (Table 5.2, Figures 5.4 and 5.5). Although restrictions imposed by the SFOC were less stringent than for the Keele campus, there were a different set of considerations for the flights at the mine that were unfortunately complicated by a series of mechanical difficulties.

There were two main difficulties: a set of parameters in the software had not been programmed and there were issues with the battery connection. For the first, I was guided through the process for setting the parameters by Aeryon engineers over the phone. The second problem was caused by a malfunctioning electrical contact pin. The Scout uses spring loaded electrical contact pins for the contact between the body of the vehicle and the battery. The malfunction was causing an error and forcing the vehicle to land shortly after taking off for each flight in the morning. We attempted first to clean the pin, but this did not correct the problem so the spring-loaded part of the pin was cut off. Some of the tolerances were also adjusted so that the flyer would not be forced to land. This allowed us to make three flights before several of the other pins also malfunctioned and the Scout died a valiant death. Unfortunately, concern about the error reoccurring prevented a climb to a higher altitude to get a photo of the area for navigation (the LiveMap feature mentioned in §2.3) and forced navigation without the benefit of a map. This deficit made navigation more challenging, as the pilot had to keep attention on the tablet in order to ensure that the desired video was being collected and therefore could not personally judge the proximity to the other wall. The

“ground supervisor” (Claire Samson) followed the UAV and relayed information about the vehicle’s position in the trench to the pilot via radio.

All of the flights were flown at low altitudes (less than approximately 15m) with the camera at a low angle (dipping about 10°) in order to ensure a high angle of incidence to the wall and thereby, hopefully, acquire data that would be effectively processed with ViDAR. However, given the narrow spacing between the two walls of the trench (approximately 14m) and given that the flights were conducted at altitudes less than the height of the walls, the pilot had to manage the video collection goals without colliding with the opposite wall. With each flight, the pilot attempted to capture as much of the height of the wall as possible, but to stay close enough to the wall to maintain a high level of detail. The flights also employed very low velocities (approximately 5 km/h) due to a combination of mechanical difficulties and collision concerns. Adjustments to avoid collision with the opposing wall, despite disrupting the intended flight pattern, actually provided some of the best data of the flights.

Table 5.2 : Mine Mission Statistics

<b>Date</b>	9 September 2011
<b>Weather</b>	Daily high: 26° Max wind: 20 km/h SW Gusts: 29 km/h Clouds: 18%
<b>Number of flights</b>	12
<b>Length of video recordings (min:sec)</b>	8:57



Figure 5.4: Satellite photo of the flight area at the Canadian Wollastonite Mine Property showing the trench and the east and west walls (from Google Earth).



Figure 5.5: Aeryon Scout in flight at the Canadian Wollastonite Mine Property (altitude approximately 25m).

## 6 Results and Discussion

### 6.1 Test Flights

The original intention with the data for the project was to obtain a triangular mesh from the point cloud data. While the video from the parking lot failed to produce a well-populated point cloud due to snow obscuring most features, the video of the gravel mound did result in a point cloud where the main features could be recognized. The flight path had been set as a circle around the mound (Figure 6.1) with the camera set to focus on the centre of the mound (one of the options for targeting the camera) and the target was imaged from several heights with a velocity of 10m/s. Figure 6.1 shows the locations of the 53 video frames selected to build the point cloud, tracing the flight path, and the point cloud itself in map view as well as from the side. Despite the snow obscuring many of the features and the low intensity contrast (everything in shades of grey), ViDAR was able to identify and track several points in what was not covered by snow and capture the basic shape.

The point cloud was then imported into Polyworks<sup>7</sup> visualization software without having been otherwise processed. Polyworks draped a triangular mesh over the point cloud and the result is shown in Figures 6.2 and 6.3. The mesh is very irregular due to the noise, the sparsity of points and the fact that the large snow-covered areas of the mound were not represented in the point cloud.

---

<sup>7</sup> [www.innovmetric.com](http://www.innovmetric.com)

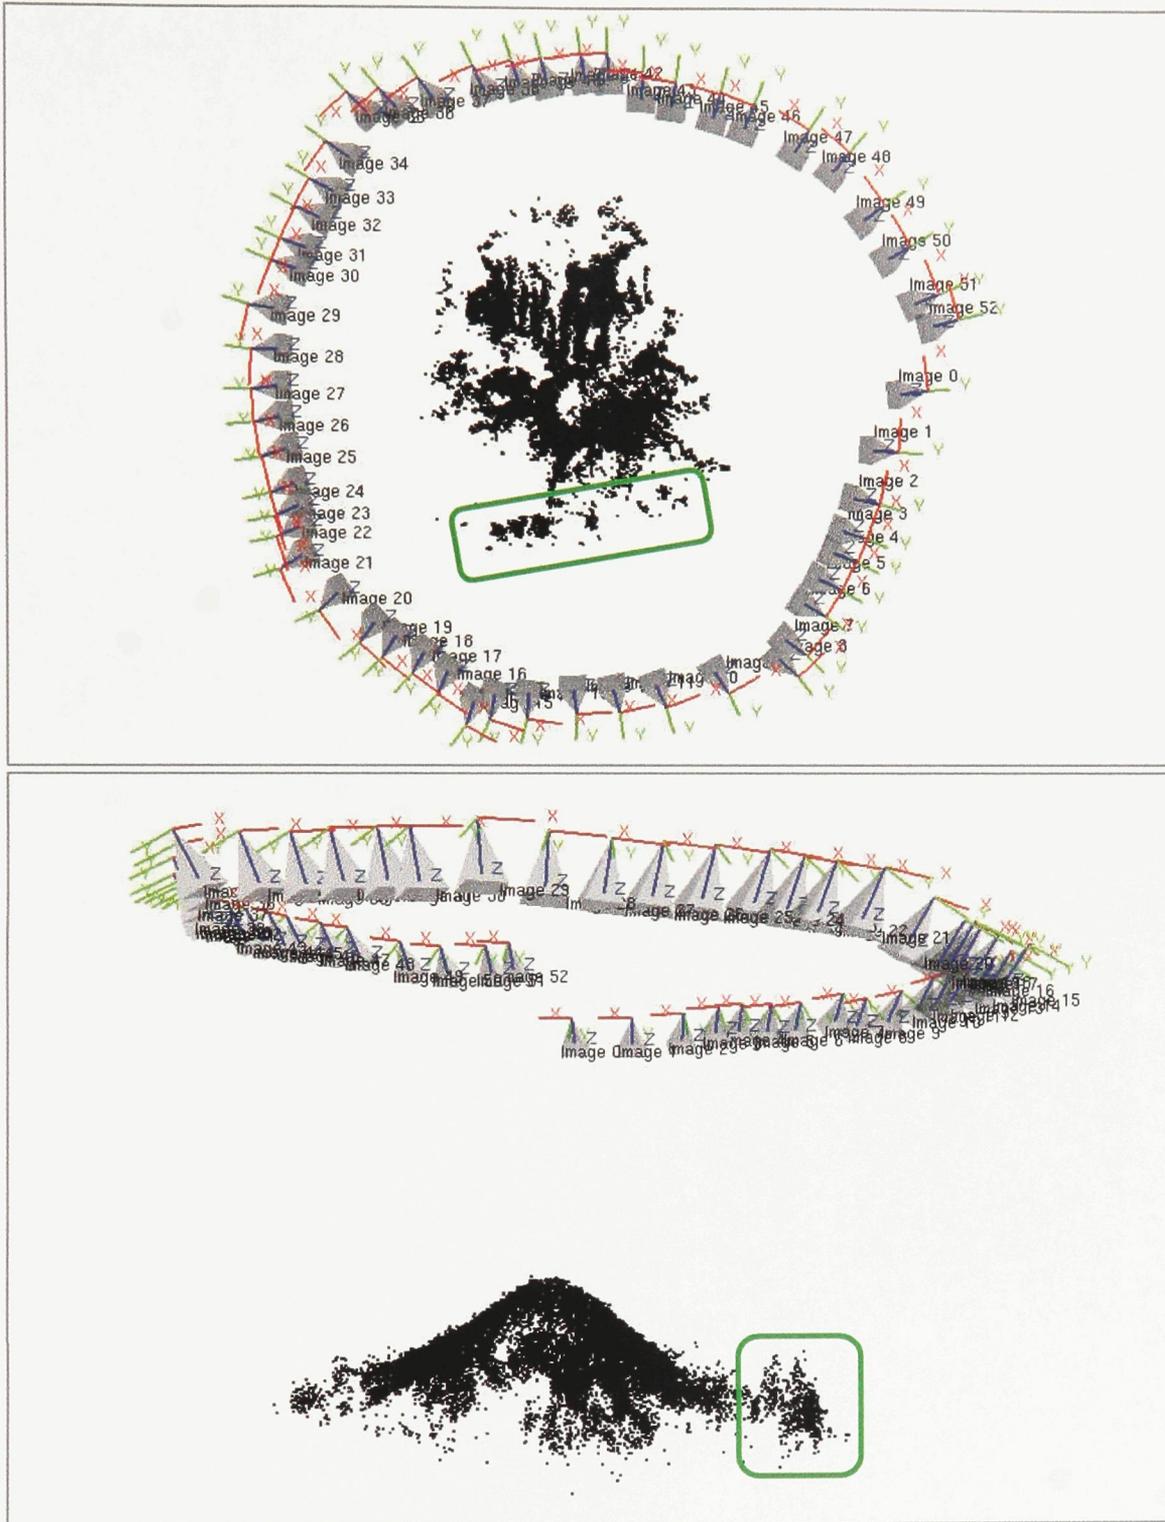


Figure 6.1: Gravel mound from the test flights in map view and the calculated locations of the 53 images used to construct the point cloud. There are 22639 points in the point cloud. Also evident are some trees near the mound (green rectangle). The mound is approximately 20m in diameter.

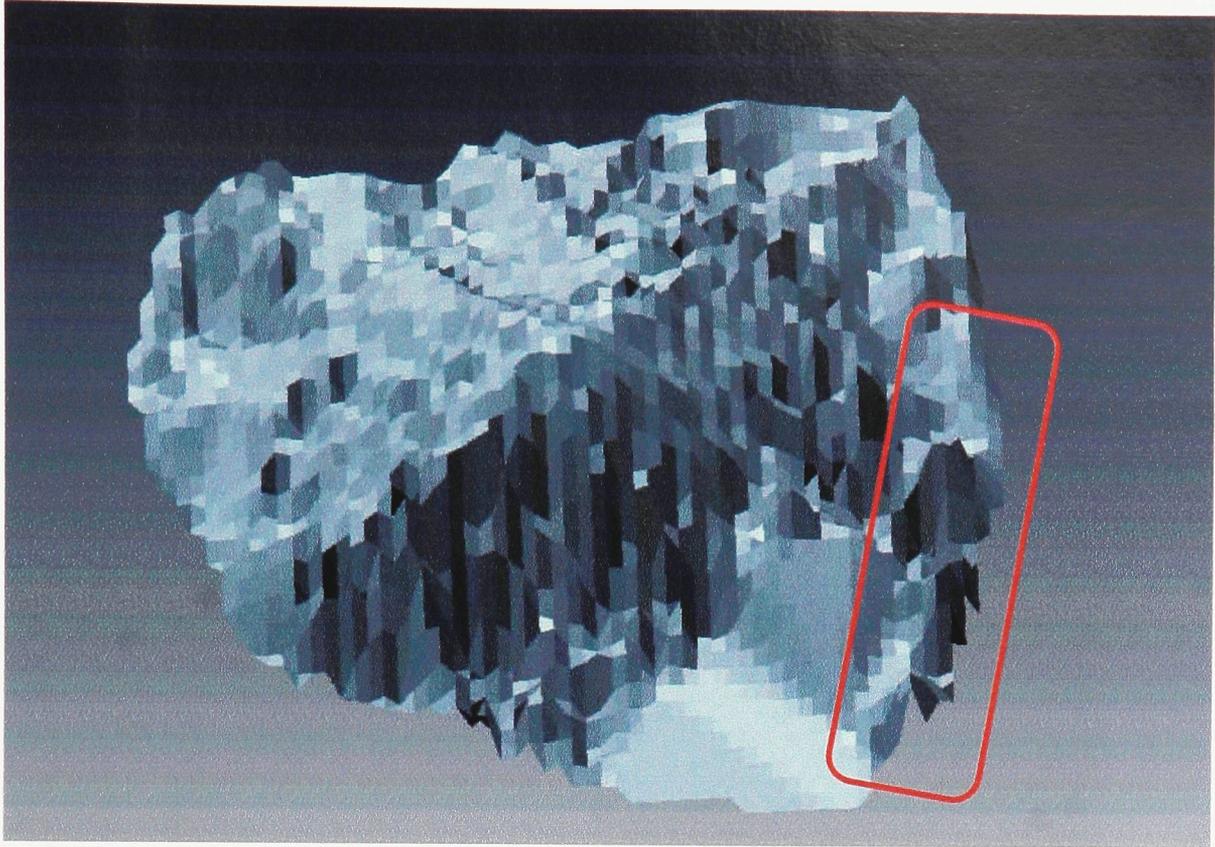


Figure 6.2: Polyworks mesh of the above point cloud in map view with the trees indicated (red box).

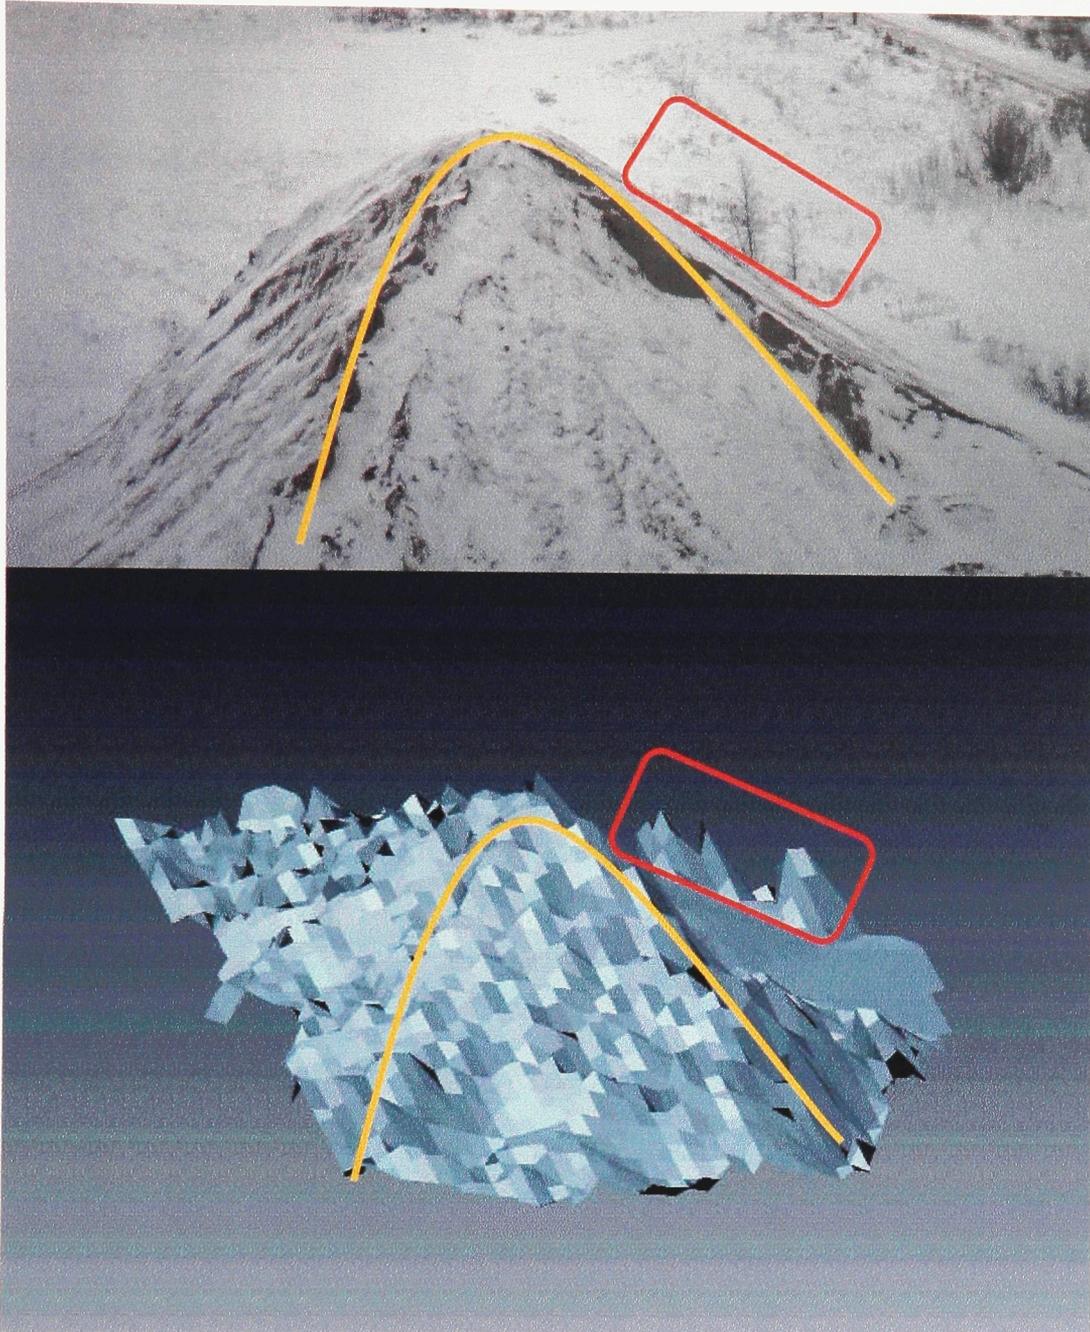


Figure 6.3: Camera still from the UAV and the Polyworks mesh in the corresponding orientation as indicated by the trees (red box) and the curve of the face (yellow curve).

## **6.2 York Mission**

### **6.2.1 Point cloud representation of different features**

The first issue to arise was in applying SfM processing at the relatively low 640x480 resolution of the video. Features in the video could be difficult to discern by visual inspection – a task even more difficult for the software – and many of the frames appeared more pixelated than the expected resolution. Since the density of a point cloud derived using SfM is dependent on the number of features that can be tracked between the frames, the low resolution made the production of even a recognizable point cloud very difficult.

In urban scenes, buildings tend to have surfaces of relatively uniform intensity: banks of windows, large panels, homogenous brickwork and so on. Coupling such homogenous surfaces with low resolution video made feature tracking difficult. Although ViDAR was set to track a maximum of 100,000 points per frame, frames that primarily featured buildings rarely reached more than a few thousand tracked points, even in the cases that worked comparatively well. The feature tracking on the western face of the Life Sciences building only had man-made features to draw from and ViDAR generally tracked approximately 2000 to 4000 points per frame (Figures 3.3 and 6.4). The many corners of panels, where different intensities are juxtaposed and create strong contrasts, are the reason why the program could track as many features as it did and also why the façades of so few other buildings were as well represented. In contrast, the example in Figure 6.5 has 7190 tracked points, most of which are in the

vegetation. The feature tracking on the buildings between the two examples differs greatly. Despite the comparatively few points tracked on the Life Science building, all of those feature points are part of the façade, whereas on the computer science building, there are only a few dozen feature points being tracked on the building façade. By comparison, vegetation can offer variation in intensity in a small area: different positions of leaves will reflect light differently and vegetation does not tend to have the smooth, regular surfaces of man-made structures. As such, vegetation is captured far more readily and preferentially populates the point clouds.

The point clouds of the buildings primarily consisted of corners and to some extent, edges, since these features often exhibited strong intensity contrasts. In many cases, however, the few isolated feature points that were registered to the point cloud were not sufficient to create a recognizable scene and most of the point clouds produced seemed a noisy collection of points. As an example, the point cloud of the southeast corner of the computer science building is displayed in Figure 6.6 (Figure 6.5 shows the feature tracking on the eastern façade). Most of the feature points on the building that were being tracked in Figure 6.5 failed to track through enough frames to be included in the point cloud and as such, the seams between the panels delineated by feature points in Figure 6.5 are not seen in the point cloud in Figure 6.6; instead, the point cloud is dominated by vegetation.



Figure 6.4: Still of the Life Sciences building west face showing tracked features (blue dots) and their calculated motion (red lines). The UAV was flying north (right to left) at an altitude of approximately 25m. There are 2331 tracked points in this frame concentrated along the perimeters of panels where there is a strong intensity contrast. The scale bar is approximate.

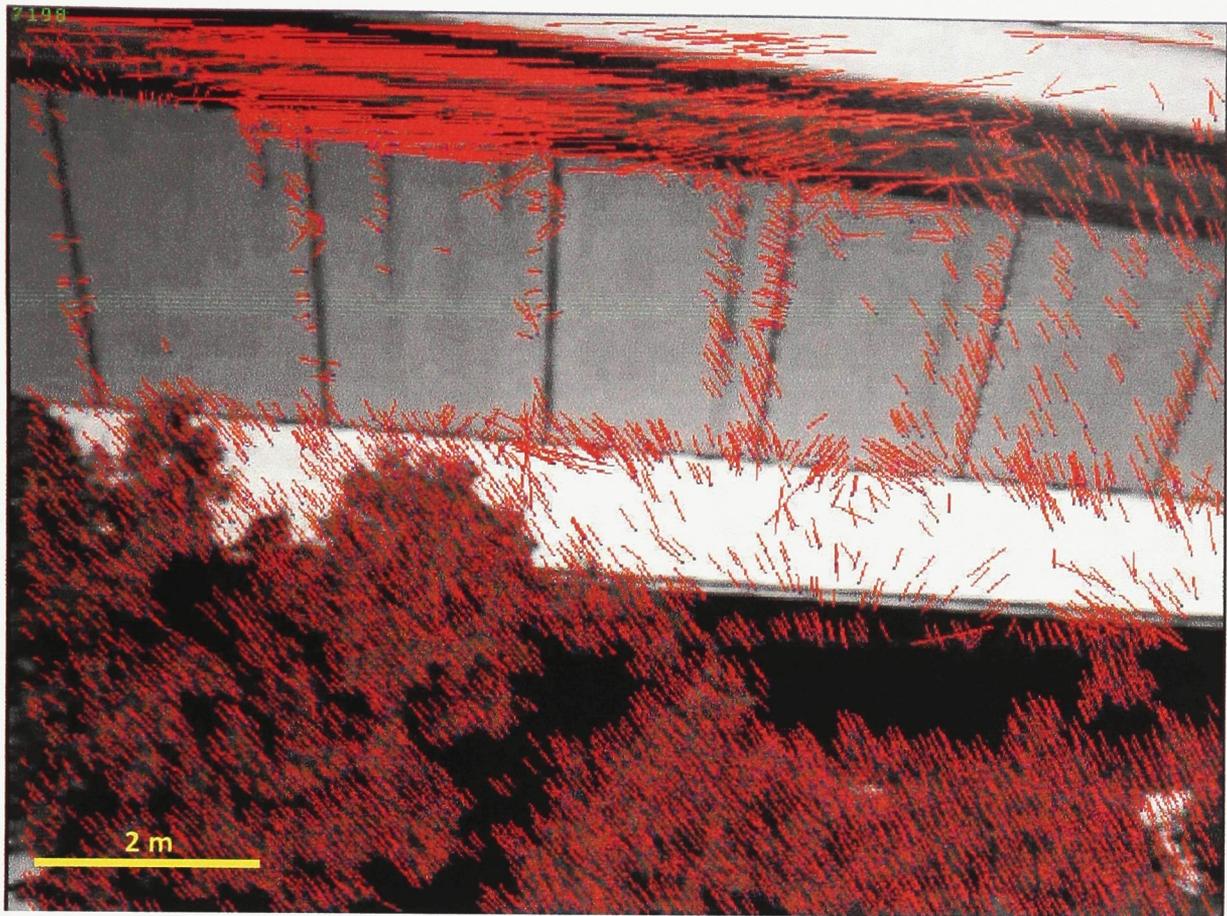


Figure 6.5: Computer Science building at York University near the SE corner. The blue dots are the tracked features and the red lines are how those features are calculated to have moved since the previously analysed frame. For the most part, corresponding feature points were properly identified, leading to a clear pattern in the optical flow (red lines approximately the same length and orientation). However, there were some instances where points between frames were incorrectly identified as being the same. This led to several of the lines at the top edge of the building having anomalous length and orientation. The UAV was flying south (right to left) at an altitude of approximately 20m. Most of the 7190 points being tracked are in the vegetation whereas the panels of the building are only sparsely represented. As in Figure 6.4, most of the tracked points on the building are along the edges of panels where there is a strong intensity contrast. The scale bar is approximate.



Figure 6.6: Point cloud (top) for the SE corner of the computer science building with green lines indicating the approximate edges of the building façade. The scale bar is approximate. On the bottom is a photo of the SE corner of the computer science building (courtesy of Bob Hou, York University).

There were some buildings, however, that were slightly more compatible with SfM – those with many smaller windows or panels in the façade: the Chemistry and the Life Sciences buildings. One of the few point clouds with recognizable features was of part of the Chemistry building and part of the courtyard that lies between the chemistry building and the Petrie building (Figure 5.2). This point cloud was registered to a wire-frame model of the Chemistry building by Larry Wang – a graduate student at York University working with Professor Gunho Sohn. He manually selected four corresponding points between the wire-frame model and the point cloud and performed a minimum fit alignment (Figure 6.7).

### **6.2.2 Range accuracy**

Despite the challenge in obtaining point clouds of planes with enough concentration for further analysis, two faces of the life science building produced point clouds with recognizable planes (south and west), with one face (west) being captured in two separate videos (the west façade had an average point concentration of  $0.00152$  points/cm<sup>2</sup>). As such there were fundamentally three planes for analysis to determine how accurately the point cloud reconstructs a plane (slight processing differences yielded three different versions of the west façade from the first life sciences video). The analysis was done using PCA (§4.1); a plane of best fit was determined for the data and the deviation from that plane was calculated (Figures 6.8, 6.9, and 6.10). Table 6.1 shows the results of the analyses on one example of the south face and four different

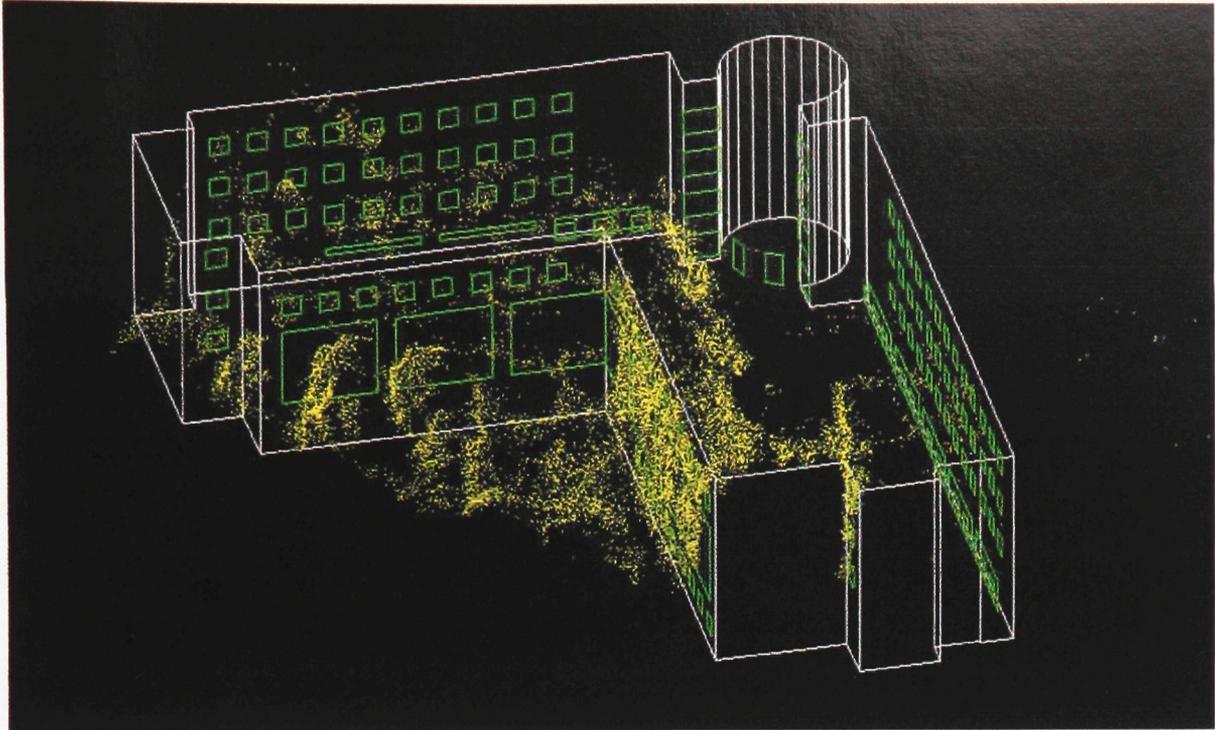


Figure 6.7: The point cloud of part of the Chemistry building with the trees in the courtyard (curved features in the foreground) registered against the wire model of the Chemistry building, courtesy of Larry Wang, York University.

point clouds of the west face. In each case, the vehicle was about 25 m from the façade (the altitude and angle of the camera was estimated by visual inspection, from which followed a rough estimation of the distance from the building).

The analysis assumes that the faces of the building would be perfect planes, which is an oversimplification. Measurements were taken of the Life Sciences building with a tape measure to establish the scale of the building as well as any relief of the surface. The panels on the walls of the Life Sciences building do not provide a perfectly flat surface, but supply a variation in relief of approximately 10cm. The windows of the first storey are set back approximately 30cm from the upper storeys; to eliminate any potential effect, the analysis was redone after removing points from the first storey. However, the noise in the data was such that the two analyses were not significantly different and although there was a minor change, it was not consistently negative or positive.

The most accurate results (i.e. the smallest standard deviation) were from the point clouds of the west face from the first life science video (file: lifesci1-wface3) (Figure 6.8). The error on the positions of the points in this case, for a range of approximately 25m, indicate that 68% of the points fall within 44cm; to include 95% of the data, the distance from the plane of best fit is extended to 88cm.

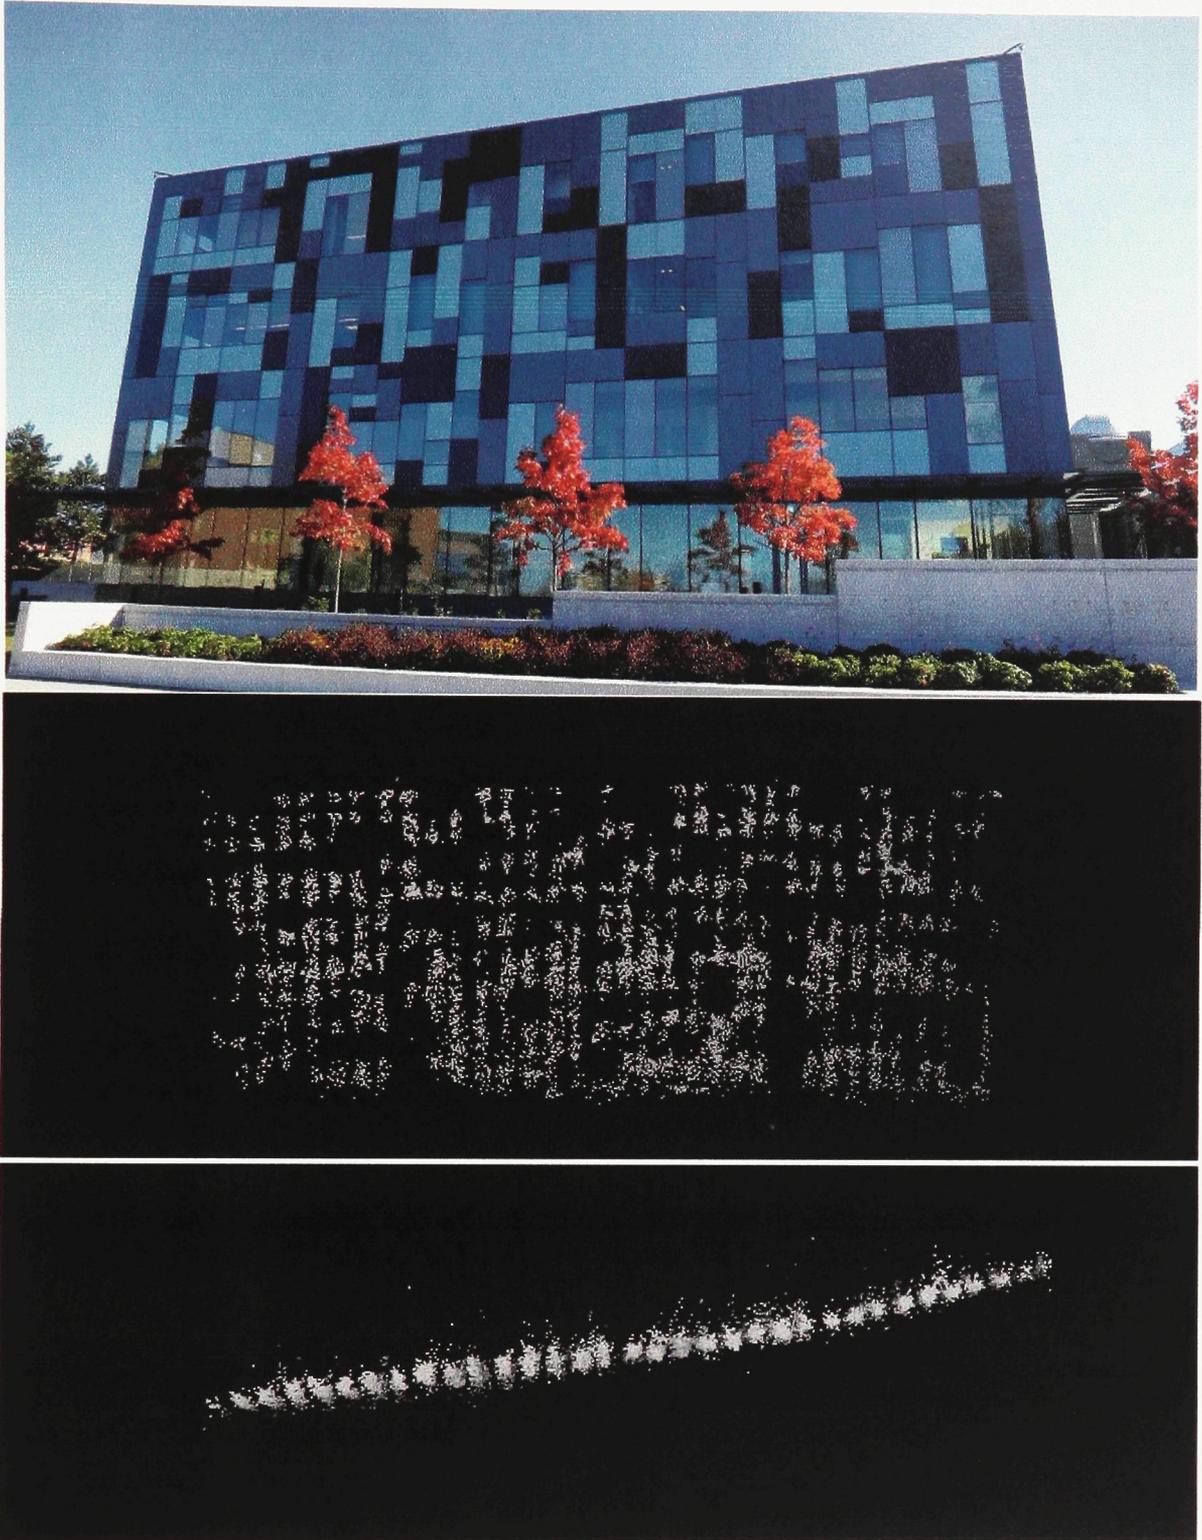


Figure 6.8: Photo (top) of the west face of the Life Sciences Building (courtesy of Bob Hou, York University) and the point cloud of the west face (lifesci1-wface3) in cross-sectional (middle) and map (bottom). The west wall is 39.5m long. There are 7521 points in the cloud.

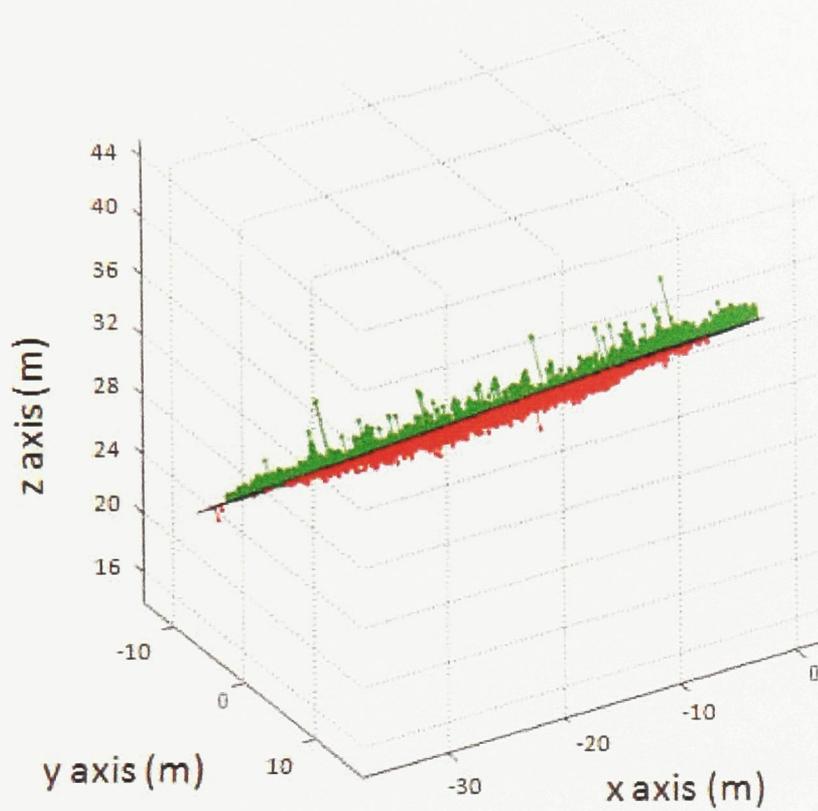


Figure 6.9: PCA output for the west face of life science building (lifesci1-wface3) showing green points above and red points below the best-fit plane (bottom).

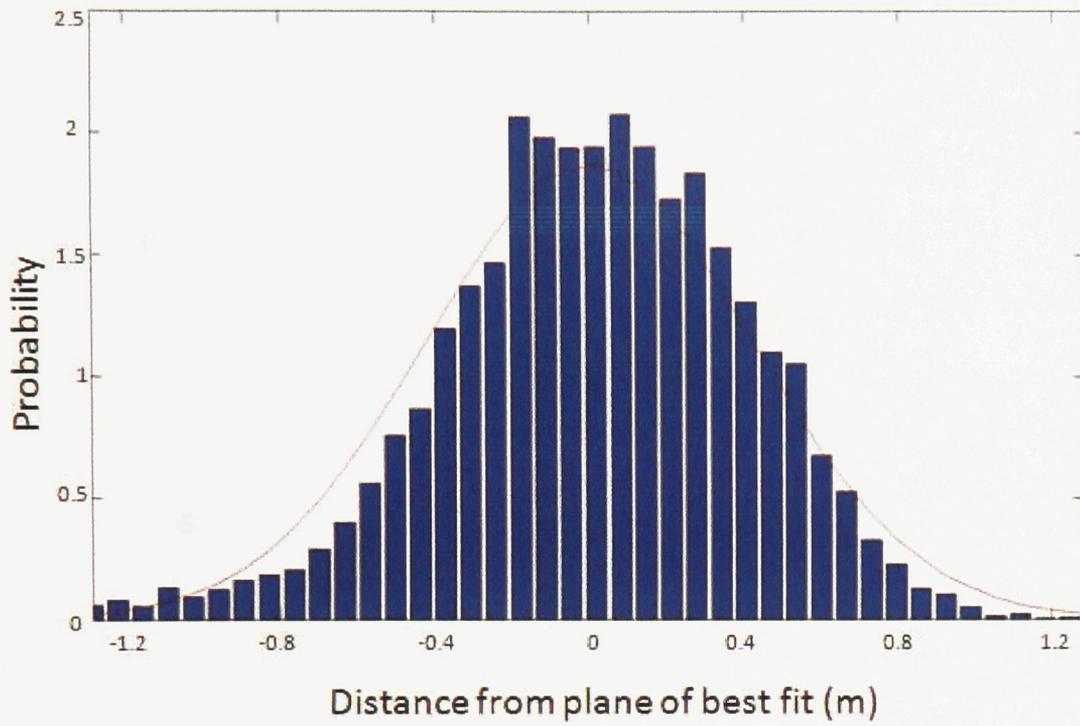


Figure 6.10: Error and Gaussian distribution within  $3\sigma$  for the west face of the life science building (lifesci1-wface3)

Table 6.1: Range accuracy of points from the planar façades of the life science building. Files ending in “b” indicate that the first storey had been removed from the point cloud.

<b>File</b>	<b>Number of points in point cloud</b>	<b>One standard deviation, <math>\sigma</math> (includes 68% of data) (m)</b>	<b>Two standard deviations, <math>2\sigma</math> (includes 95% of data) (m)</b>	<b>Three standard deviations, <math>3\sigma</math> (includes 99.7% of data) (m)</b>
lifesci1-sface1	13720	0.85	1.71	2.56
lifesci1-sface1b	12115	0.85	1.70	2.55
lifesci1-wface1	7149	0.52	1.05	1.57
lifesci1-wface1b	6448	0.52	1.04	1.56
lifesci1-wface2	7104	0.46	0.93	1.39
lifesci1-wface2b	6109	0.47	0.94	1.41
lifesci1-wface3	7521	0.44	0.88	1.31
lifesci1-wface3b	6332	0.44	0.88	1.32
lifesci2-wface1	10904	0.54	1.07	1.61
lifesci2-wface1b	8378	0.52	1.04	1.56

### **6.2.3 Classification software**

Given the poor representation of planes in the urban data, there were very few scenes that displayed both a man-made plane and vegetation. One of the few places where a surface was captured was the roof of the computer science building during one of the scans of the northeast corner. The reason this plane was captured is because it has sections of grass on the roof.

Initial analysis had both the roof and the surfaces of the trees next to the building classified as being part of planes. Analyses were performed at increasing radii (from 3 to 17 voxels, or approximately 0.75 to 3m) in an attempt to improve the identification by including more of a planar surface, but this was unsuccessful. Increasing the radius for the analysis did change the points that were identified as being part of a plane, but the change was not significant and did not exclude the trees from the results. Although the normals for the points on the roof did become progressively better aligned at higher radii, fewer of the points on the roof were included at higher radii due to the inclusion of more empty space past the edge of the point cloud.

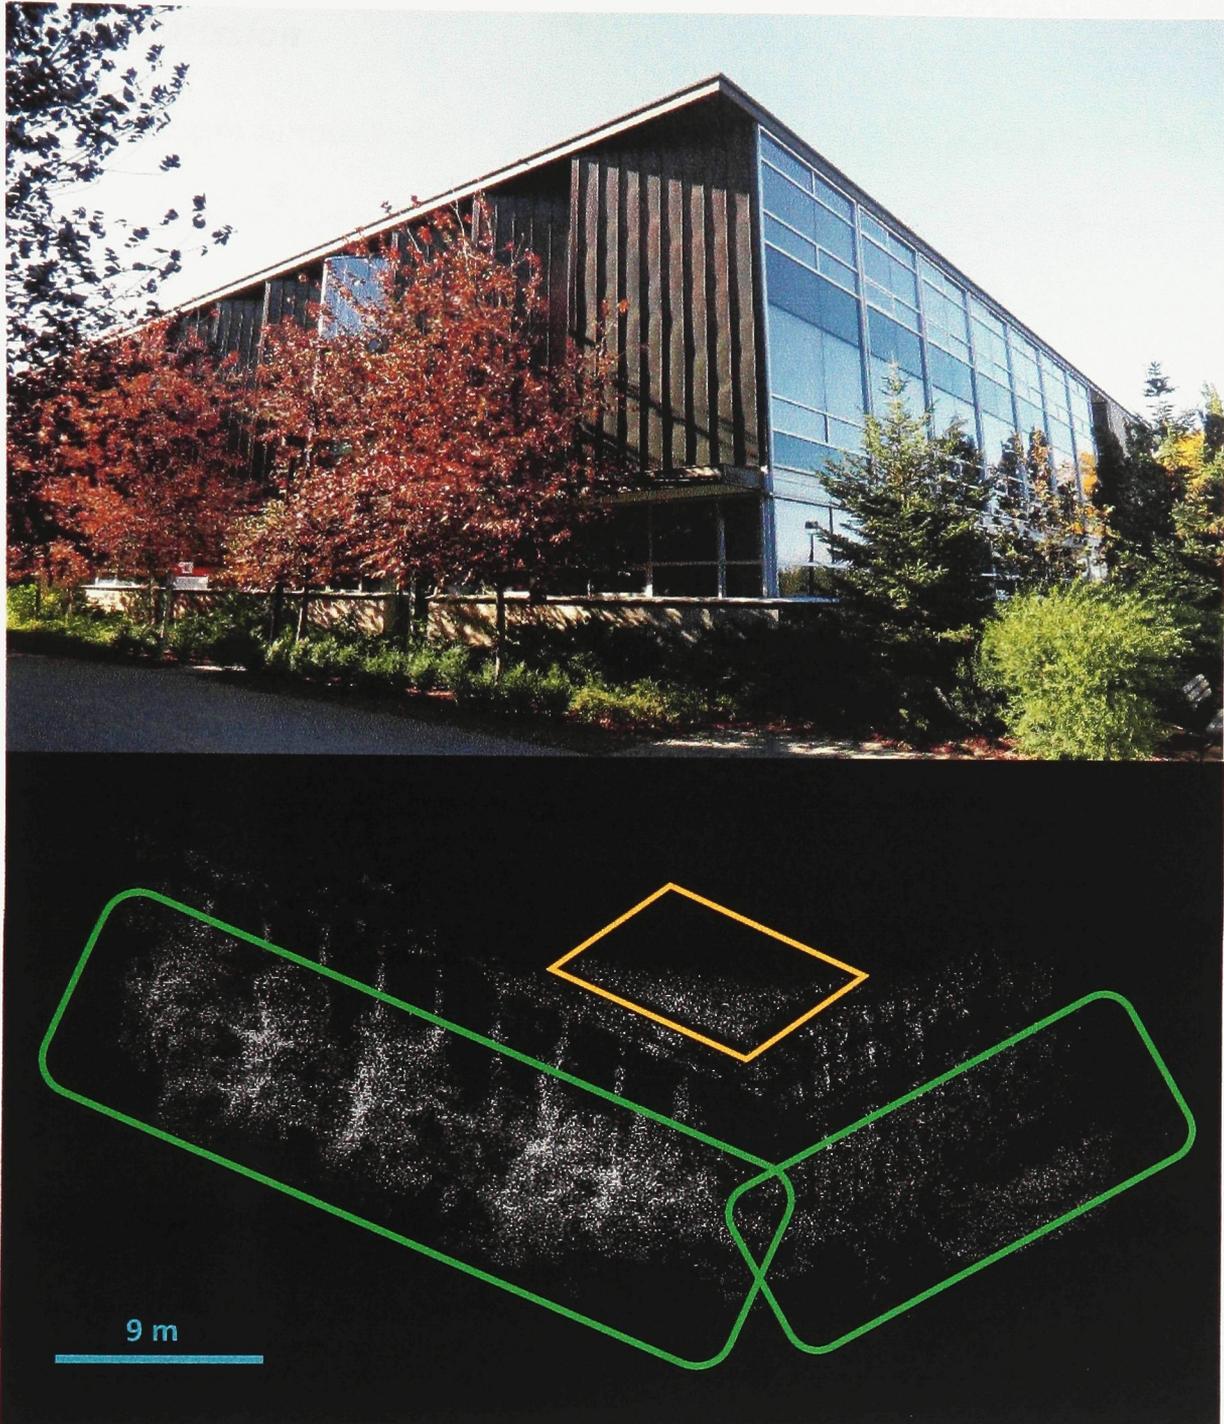


Figure 6.11: Photo (top) (courtesy of Bob Hou, York University) and point cloud (bottom) of the north east corner of the computer science building showing the roof surface (yellow box) and the trees adjacent to the building (green boxes).

## **6.3 Mine Mission**

The problems surrounding the low resolution of the original video images were not limited to the urban scenes. At the mine as well, the video needed to be taken very close to the rock face in order to capture enough features to build a point cloud with recognizable features. So even though the rock faces did offer significantly more texture and intensity variation on a surface, the video resolution was too low to produce high-density point clouds.

### **6.3.1 Rock mass characterization software**

The rock mass characterization software used in this study (Mah et al., 2011) was originally developed for point cloud data with a very high point concentration ( $\approx 40$  points/cm<sup>2</sup>) and cleaner rock surfaces devoid of vegetation. At the Canadian Wollastonite Mine site, surfaces were noisier and the coverage was sparse. The highest point cloud concentration resulted from (1) video images taken the closest to the rock walls (distance to the rock wall  $< 8$  m) and (2) from video that covered the same small area repeatedly over a short time – moving back and forth across the space as the UAV was brought closer to the rock face (Figure 6.12). The only segment of the video that displayed both these qualities was from the west wall. There were other segments of video that were taken closer to the rock face and while the resolution was better than for videos from farther away, they did not produce point concentrations comparable to the video demonstrating the multiple short passes.

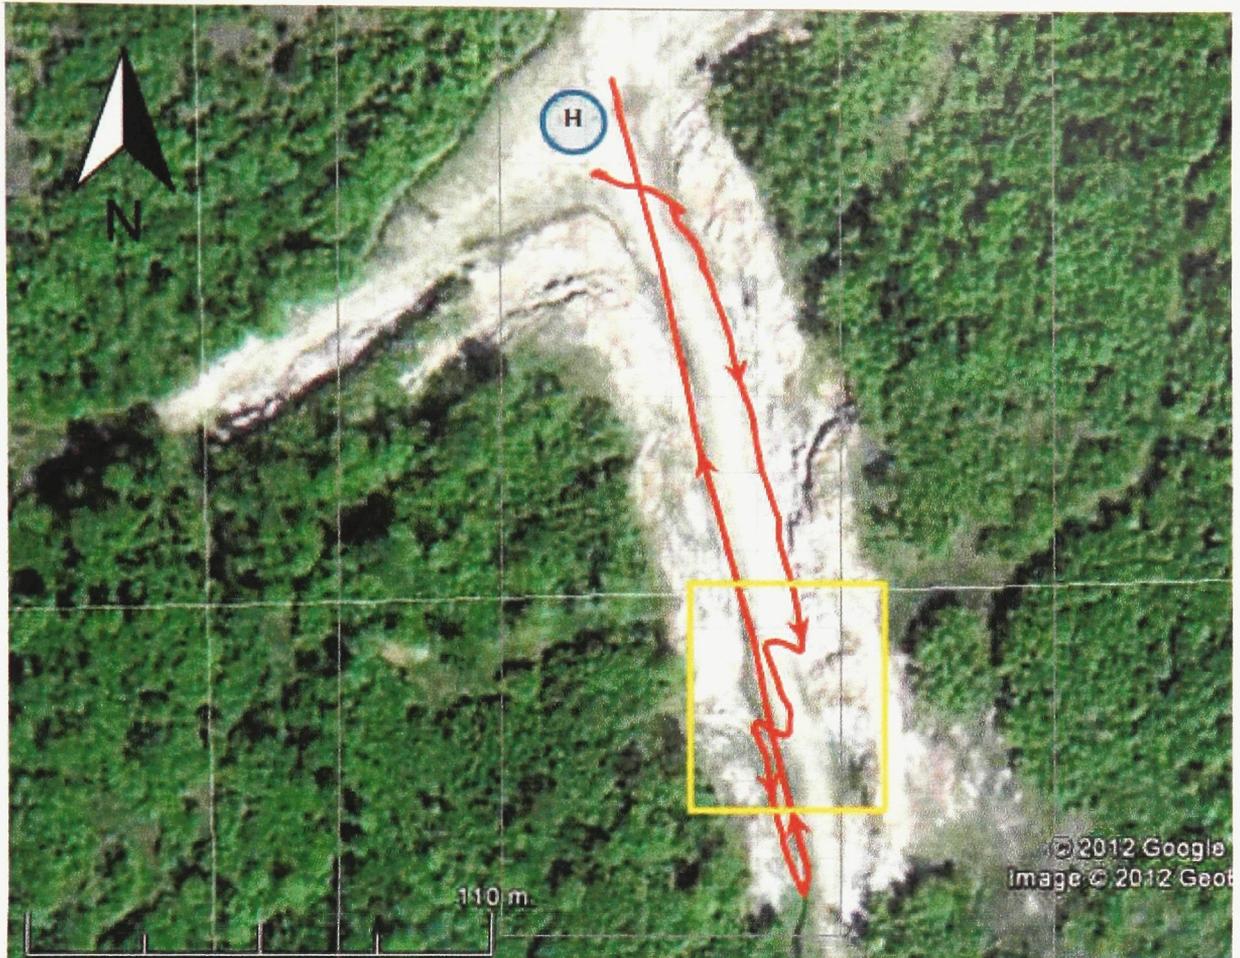


Figure 6.12: Flight path for the west wall (red line) with approximate location of "home" (blue circle). The yellow bounding box indicates the location for measurements 27 and 29. The altitude for the flight was set at approximately 7.5m.

In the Mah et al. (2011) study, the radius chosen for PCA was between 5 and 80 mm which corresponded to a subset of approximately 32 to 8,243 points. In this case, with a much smaller point cloud concentration, the radius had to be increased so that enough points were selected to compute the best-fitting plane (tens of centimeters). This sparsity inherently excluded any smaller surfaces from consideration. The size of the face to be considered was approximated from the length measurements taken in the field and a radius was selected that would encompass most of the face. Although several surfaces are in excess of  $1\text{m}^2$ , there is a necessary balance to be maintained between sampling enough of the surface to include sufficient points to produce a reliable result while not including any nearby edges which would add an element of noise and skew the results. This matter was complicated by the fact that there was no easy way to determine the coordinates of a specific point in the point clouds and the other programs used to address this issue provided only a rough estimate. Another challenge is that ViDAR does not provide scale for the point clouds it produces. Therefore, we had to use the measurements made in the field (see appendix B) and correlate them with the point cloud in order to determine the scale. This information was required to define an appropriate radius for the PCA analysis in the rock mass characterization software.

The strike and dip of the camera are necessary to georeference the point cloud to north and to thus properly determine the orientation of the rock face. Solving the issue of determining these factors during the period of data collection was relatively straightforward. The log files collected by the Scout during the flights included

information on the yaw and pitch of the camera (used to determine strike and dip, respectively). This process was made easier by the fact that the camera orientation was not intentionally adjusted after data collection began, although there were some minor variations likely attributable to turbulence and gimbaled reactions. This was done intentionally in the field and proved to be an asset since the log files do not indicate when the camera was recording and data were not collected continuously from takeoff. However, once the video was started, recording continued until landing which narrowed down the approximate time by working backward. For one flight, this method worked exactly as intended with very little variation on strike or dip of the camera: strike mostly stayed within  $10^\circ$  of the average value and the dip was fixed. For another flight, it was the variation in the strike of the camera that allowed the timing for the video segment to be determined. For the third flight, the GPS position data in the log files (latitude and longitude) were used to more precisely determine the part of the flight in which the relevant video was collected.

Due to the low point concentration of the point clouds, only 6 locations were conducive to analysis with the rock characterization software (Table 6.2). With the introduction of the smoothing algorithms, some analyses were done on smoothed point clouds with the hope that the reduction in noise would allow for more accurate results. The smoothing also made it easier to pick out features and determine where to set the area for PCA analysis. All of the results from the east wall were determined from the smoothed point clouds; results from the two locations at the west wall were obtained without using smoothing.

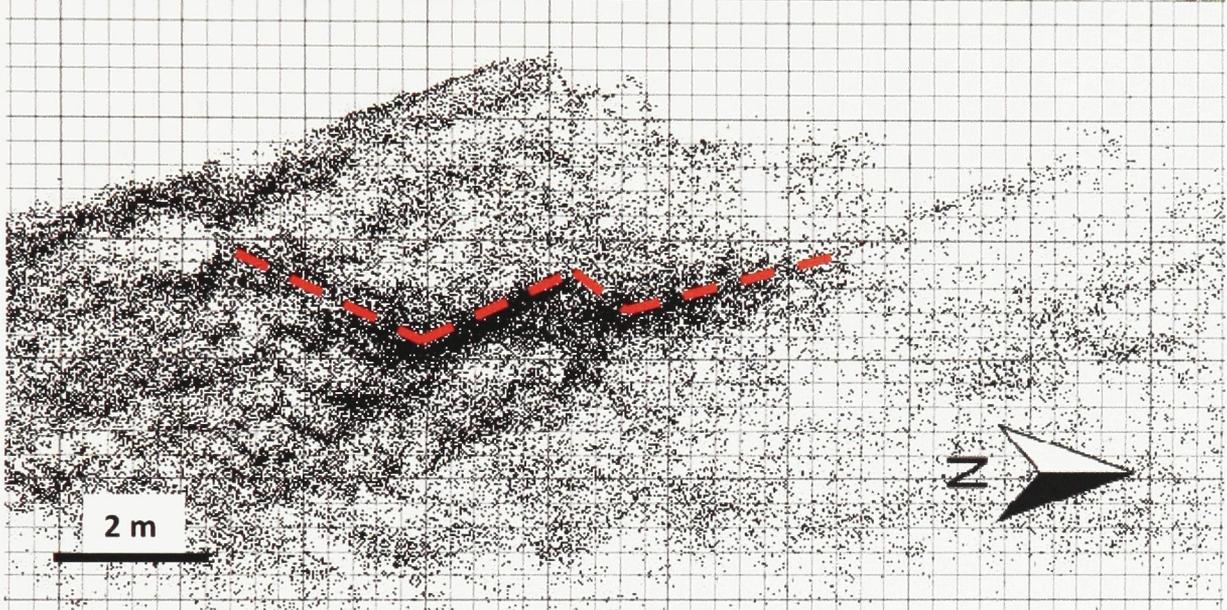
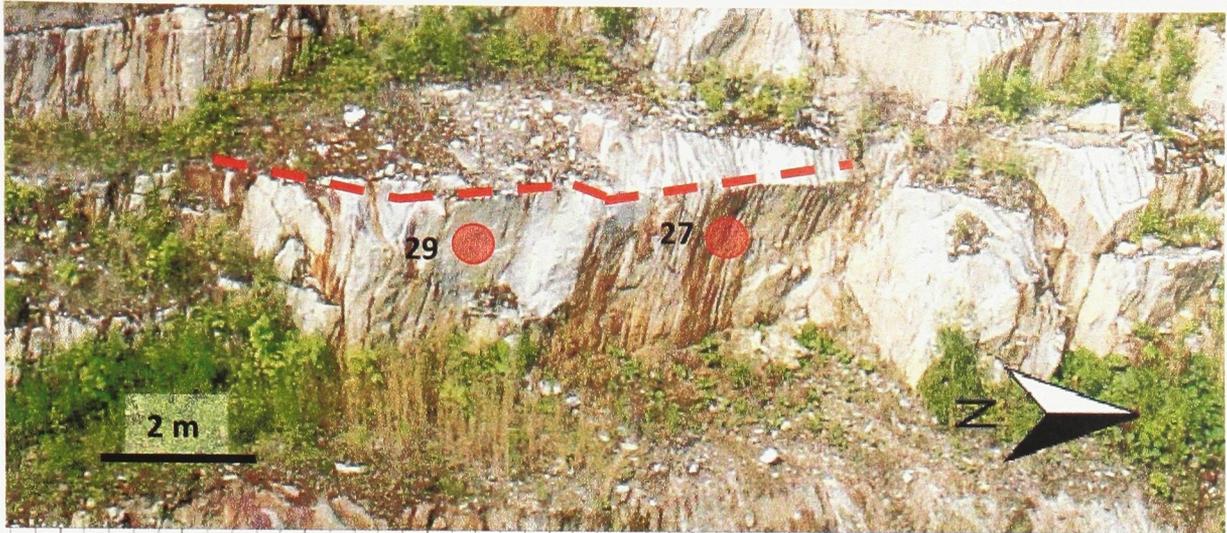
Table 6.2: Comparison between joint orientation measurements and results from the rock mass characterization software

Measurement number	Manual measurement				Software measurement							
	Joint set number	Strike (°)	Dip (°)	General Orientation	smoothing	Strike (°)	Dip (°)	PCA radius (pt clid units)	PCA radius (approx.) (cm)	Number of points in PCA	Point cloud concentration (points/cm <sup>2</sup> )	Angular difference
West 27	2	136	85	subvertical	no	136	84	0.15	22.5	138	0.09	1
West 29	2	134	81	subvertical	no	133	83	0.15	22.5	160	0.10	2.23
East 9	1	283	80	subvertical	yes	244	77	0.65	75	207	0.01	38.3
East 9 (parallel)	1	283	80	subvertical	yes	264	82	0.4	75	291	0.02	18.87
East 10	2	186	88	subvertical	yes	153	89	0.35	75	141	0.01	33
East 13	3	132	11	subhorizontal	yes	189	20	0.4	75	199	0.01	16.67

The dip of the subvertical surfaces (measurements west 27, west 29, east 9, east 9 (parallel), and east 10) agreed with the manual measurements within 3°. There is only one measurement (east 13) done on a subhorizontal surface. Subhorizontal surfaces are notoriously difficult to capture. They are typically imaged at grazing angles, leading to low point cloud concentration. In addition, results for the subhorizontal surface could have been negatively affected by the debris that collected on the subhorizontal surface at that location.

Figure 6.13 shows locations #27 and #29 on the west wall where the point cloud concentration was the highest (Table 6.2). At these locations, the analyses led to strike/dip measurements of 136/84 and 133/83 respectively giving a 1° difference in dip for #27 and a 2° difference in dip for #29 (Table 6.2). This was the best agreement between the manual and software measurements.

The method for these data involves considerable trial and error: for the (x, y) coordinates of the central point of analysis and for the radius itself. There was also a great deal of error propagation in the analysis. First the surfaces at the mine were usually quite rough and uneven and did not offer many smooth surfaces. The instrument noise has previously been mentioned as well as the low density of points. The quality of the surfaces combined with the sparsity of the points made it very difficult to find surfaces that had a consistent orientation and were large enough to contain a sufficient number of points. This is likely the reason why the analysed surfaces



point cloud of both figures. The

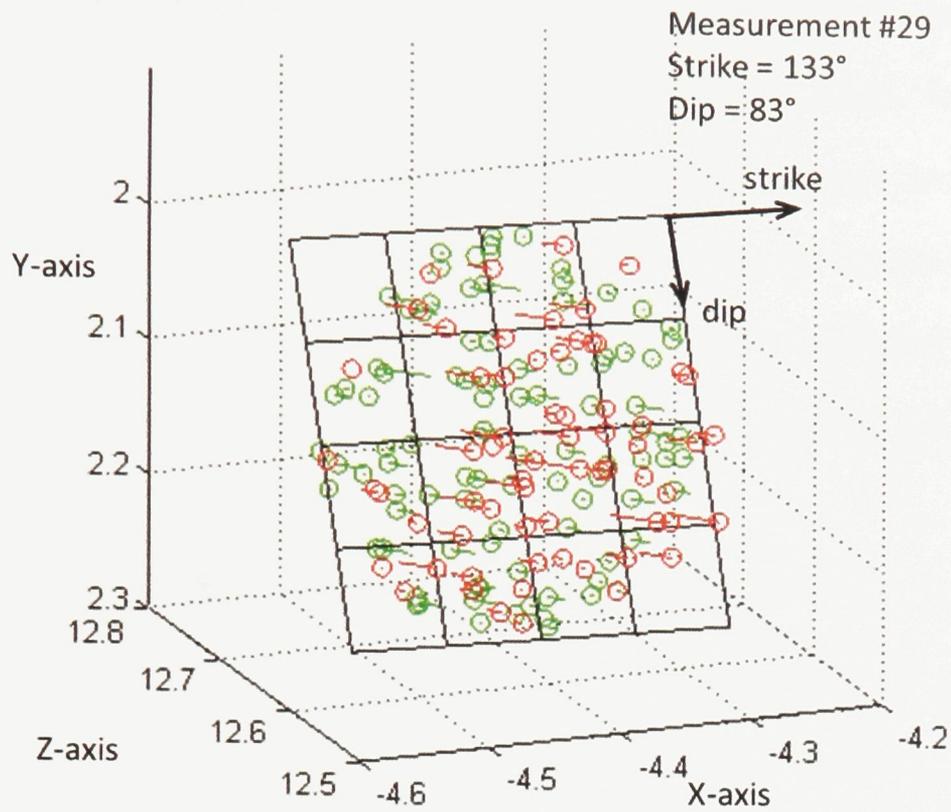


Figure 6.14: PCA results for measurement 29 on the west wall. Red and green points are below and above the best fit plane respectively.

from the west face worked as well as they did: the surfaces were relatively smooth and large and the point density was higher than any other point cloud. Additional error was incurred due to uncertainty with the strike and dip of the camera, particularly the former. Although the orientation of the 3D space is set with the first camera orientation used in ViDAR, it was not possible to determine this point exactly and an average of the orientation during the approximate period of the flight had to be used.

### **6.3.2 Visualization software**

Given the noise in the data, the triangles generated by the points (used as vertices) were not consistently oriented in a face. Since the points did not define a flat surface, but rather produced an approximately normal spread of points around where the surface would be, using these points as vertices in the mesh resulted in each small triangle demonstrating a different orientation. When the triangles were colour coded according to strike there were a wide range of colours in a single face that indicated substantially different orientations. With the dip overlay it was sometimes possible to determine broad structural features (steps and shelves), but no details could be picked out and no faces were obvious (Figure 6.15 (middle)). Using the strike overlay, there were no patterns at all (Figure 6.15 (bottom)). To address this challenge, the point clouds were additionally processed with a smoothing algorithm to see if this could offer any improvement in the visualization process.



Figure 6.15: The dip (middle) and strike (bottom) visualization for data that have not been smoothed. A photo (top) corresponding to the area is shown; the yellow circle indicates a common feature (the angle in the photo differs from that in the model).

With the smoothing applied, it became possible to identify some structural features although it could be challenging to determine the corresponding points between the mesh and the image (Figure 6.16). The noise in the original data obscured most of the detail of the structure, a fact that carried through after smoothing. Ten iterations were necessary to smooth most of the noise out of the point cloud and to display the broader features far more clearly; using fewer iterations did not make any features more visible and only retained some of the roughness caused by the noise. All of the features of the point clouds were made smoother: faces and edges as well as vegetation. As such, rock was indistinguishable from brush which made identifying features more difficult and it was necessary to know what features were supposed to be in the scene in order to correlate the corresponding shapes.

Seven point clouds from the mine site were processed with the visualization software with varying results. The dip overlay showed a distinction between shelves and steps (however, the angle was measured from the vertical rather than the horizontal). Mostly what was visible using the strike overlay were patches of colour. These rarely delineated actual faces, rather there could be multiple patches indicating very different strike in what was known from the dip overlay to be a single surface. In some of the point clouds the patches did roughly indicate the faces, but failed to clearly define them. Overall, the program could delineate broad features. Further, since the program does not yet georeference the point cloud, the values offered are relative values within the point cloud rather than absolute strike and dip values.

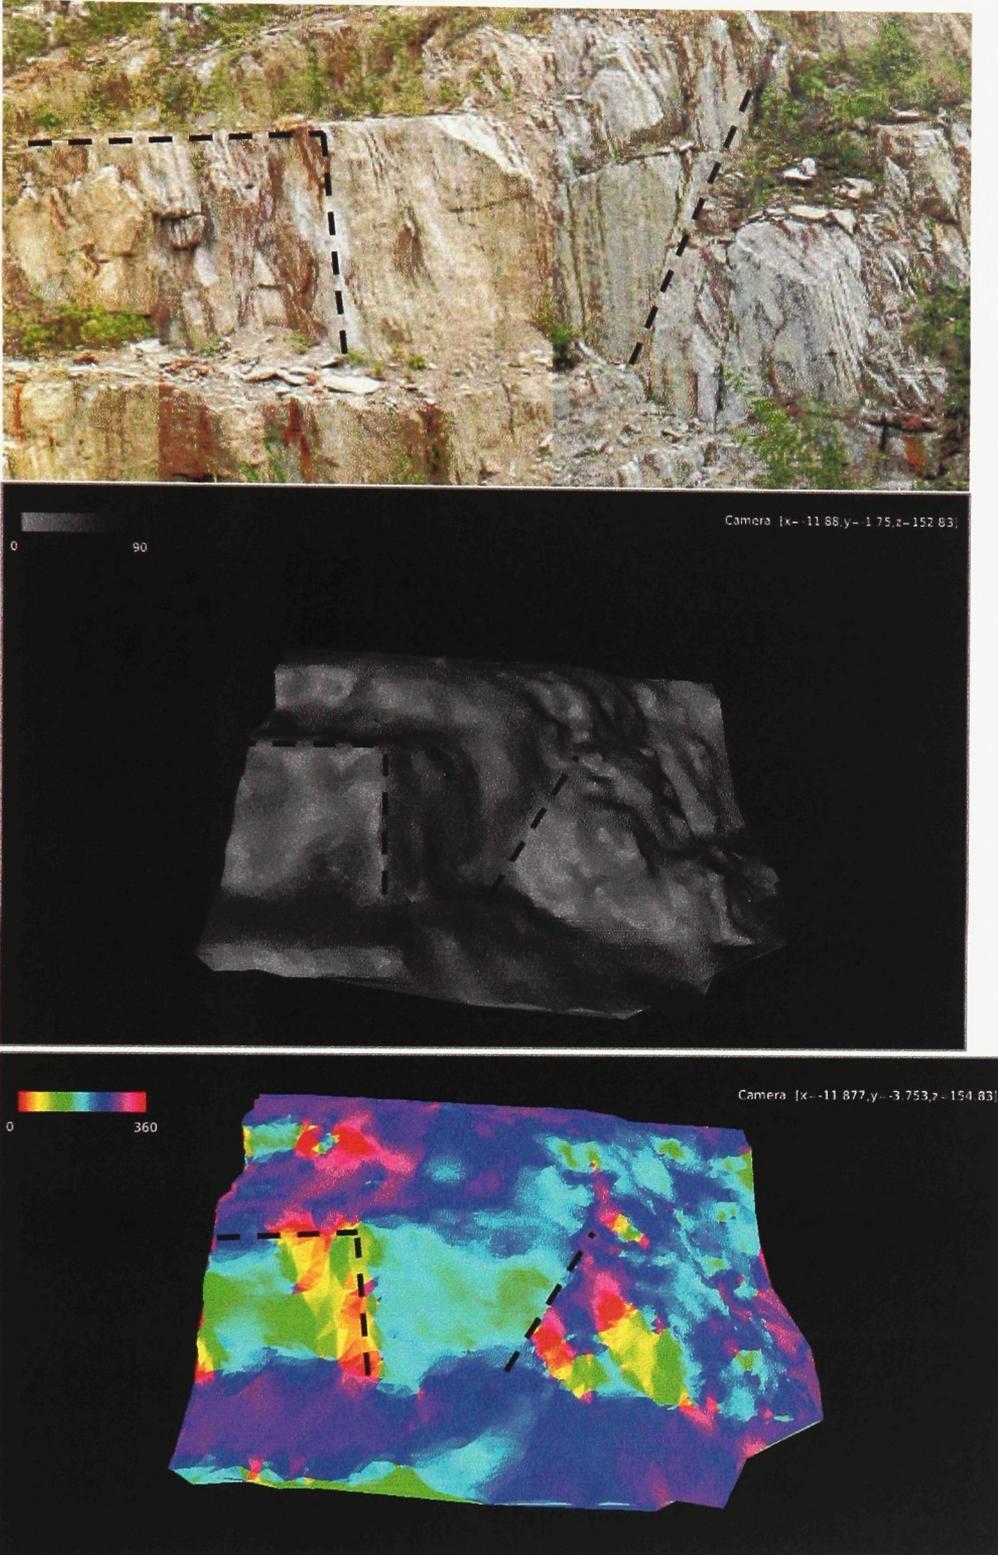


Figure 6.16: Dip (middle) and strike (bottom) visualization for the smoothed point cloud for the area shown in the photo (top). A few edges are delineated for reference. For strike, the colours do not correspond with the manual measurements, since the program does not yet georeference the point clouds.

## 7 Concluding remarks

The overarching goal of the project has been to examine the effectiveness of this particular combination of technologies. The two missions at York University and at the Canadian Wollastonite Mine had their own specific goals.

The data from the university allowed for an assessment of error and indicated the vast difference in point cloud concentration between differently textured surfaces. The low video resolution is responsible for the low overall point cloud concentration. However, there is a difference between scenes: scenes rich in intensity variations (e.g. leaves, rock faces) register more feature points in SfM resulting in higher point cloud concentrations, whereas scenes with minimal intensity variations (e.g. homogeneous brick wall) will register fewer features in SfM resulting in a lower point cloud concentration (§6.2.1): the point cloud concentration for the rock faces ranged from 0.01 to 0.1 points/cm<sup>2</sup>, while the west face of the Life Sciences building (a well represented man-made surface) had a point concentration of less than 0.002 points/cm<sup>2</sup>.

The points on a well represented flat surface provided an opportunity to quantify the range error (§6.2.2). In this instance the UAV was approximately 25 m from the building and the range error (one standard deviation from the mean) was on the order of 50 cm. Although this level of error might be acceptable in certain applications, it impeded the performance of the classification software. The error on a surface that should have been flat was comparable to that in the foliage of deciduous trees (§6.2.3)

which caused both such surfaces to register as planes, indicating that the classification software cannot provide useful information about the point clouds generated by the method used.

The data from the mine, despite being natural features, still produced low density point clouds with concentration ranging from 0.01 to 0.10 point/cm<sup>2</sup>. The rock mass characterization software (Mah et al., 2011) was originally designed for point cloud concentrations several orders of magnitude higher. Although the low point cloud concentration limited the areas where strike and dip measurements could be attempted (§6.3.1), on the few locations where the approach was feasible results were encouraging. The visualization software delineate variations in dip quite well, and when the faces are large enough give a reasonable indication of the locations of joint faces (§6.3.2).

The technologies were not as compatible as had been hoped. With so many different technologies involved in obtaining the results, it is difficult to pinpoint a single source of the problem; rather it would seem that the issue is this particular combination. The process did not return point clouds of sufficient density to work well with any of the processing software used in the project. It is likely that with higher resolution video, the results would improve. The technology is advancing rapidly and the current version of the Scout is now capable of acquiring video data at a resolution 1080p (1080 lines of vertical resolution). Different SfM programs would likely produce different results (whether better or worse is unknown). Further progress towards a robust system lie mainly 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 ViDAR (desired point cloud concentration, range error, etc.).

As it stands, this process could not fulfill the need for a less expensive, lightweight airborne 3D imaging system. However, technology continues to advance toward smaller, lighter-weight, faster, higher quality data collection devices – cameras and lidar systems among others. In the near future, small 2D lidar systems might provide 3D data directly, without the need for SfM. Future efforts will take advantage of these perpetual advances.

## References

- Aeryon Labs Inc. (2012). Aeryon Labs Inc, Waterloo, Ontario. [www.aeryon.com](http://www.aeryon.com).
- Anderson, D.E., and Pita, A.C. (2005). Geophysical Surveying with GeoRanger™ UAV. American Institute of Aeronautics and Astronautics.
- Barnard, J.A. (2008). Use of unmanned air vehicles in oil, gas and mineral E+P activities: 78<sup>th</sup> Annual International Meeting, SEG, Expanded Abstracts, 1132-1136.
- Bento, Maria de Fátima. (2008). Unmanned Aerial Vehicles: An Overview. *InsideGNSS*, 3(1): 54-61.
- Blais, F. 2004. Review of 20 Years of Range Sensor Development. *Journal of Electronic Imaging* 13(1):231-240.
- Bouguet, Jean-Yves. (2001). "Pyramidal Implementation of the Affine Lucas Kanade Feature Tracker – Description of the algorithm", Intel Corporation – Microprocessor Research Labs. Retrieved from [http://robots.stanford.edu/cs223b04/algo\\_affine\\_tracking.pdf](http://robots.stanford.edu/cs223b04/algo_affine_tracking.pdf).
- Bray, N. (2004). Notes on Mesh Smoothing. Retrieved from <http://mgarland.org/class/geom04/material/smoothing.pdf>.
- Campbell, J.B., and Wynne, R.H. (2011). *Introduction to Remote Sensing* (fifth edition). The Guilford Press. New York, New York.
- Canadian Wollastonite (2008). About Wollastonite. Canadian Wollastonite: The St. Lawrence Project. Retrieved from <http://www.canadianwollastonite.com/>
- Caron, R. 2011. Aeromagnetic surveying using a simulated unmanned aircraft system.

M.Sc. thesis, Department of Earth Sciences, Carleton University.

Caron, R., Samson, C., Straznicky, P., Ferguson, S., Archer, R., and Sander, L. 2011.

Magnetic and magneto-gradiometric surveying using a simulated unmanned aircraft system. 81<sup>st</sup> Annual Meeting of the Society of Exploration Geophysicists, San Antonio, TX, 18-23 September. SEG Expanded Abstracts **30**, 861-865.

Crutchley, S. and Crow, P. (2009). *The Light Fantastic: Using airborne laser scanning in archaeological survey*. Swindon: English Heritage.

Derpanis, K.G. (2010). Overview of the RANSAC Algorithm . Retrieved from

[http://www.cse.yorku.ca/~kosta/CompVis\\_Notes/ransac.pdf](http://www.cse.yorku.ca/~kosta/CompVis_Notes/ransac.pdf).

English, C., Zhu, S., Smith, C., Ruel, S., Christie, Iain. 2005. TriDAR: A hybrid sensor for exploiting the complementary nature of triangulation and LiDAR technologies. Proc. of 'The 8<sup>th</sup> International Symposium on Artificial Intelligence, Robotics and Automation in Space.

Frueh, C. and Zakhor, A. (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.

Harrison, J. (2009). Efficient Initial Segmentation of Point-cloud data. Unpublished manuscript.

Hefford, S.W., Samson, C., Iles, P.J., Harrison, J.W., Ferrie, F.P., Kusevic, K., and Mrstik, P. 2009. Augmenting the Iterative Closest Point (ICP) alignment algorithm with intensity. *Geomatica*, **64**(4): 407-418.

Hogervorst, M.A. (1966). Limitations of the recovery of structure from motion.  
Netherlands.

Johnson, P. (2006). Unmanned aerial vehicle as the platform for lightweight laser sensing to produce sub-meter accuracy terrain maps for less than \$5/km<sup>2</sup>. Retrieved from <http://s122642433.websitehome.co.uk/uavlidar/report.pdf>.

Karni, Z., and Gotsman, C. (2000). Spectral Compression of Mesh Geometry. Computer Graphics (Proc. SIGGRAPH), pp. 279-286.

Mah, J., Samson, C., and McKinnon, S. 2011. 3D Laser imaging for joint orientation analysis. Int. J. of Rock Mech. and Mining Sci. 48: 932-941.

Mah, J., Samson, C., McKinnon, S., and Thibodeau, D. 2012. 3D laser imaging for surface roughness analysis. Int. J. of Rock Mechanics and Mining Sciences, in press.

Morris, W., Dryanovski, I., and Xiao, J. 2010. 3D indoor mapping for micro UAVs using hybrid range finders and multi-volume occupancy grids. Retrieved from [http://www.seattle.intel-research.net/RGBD/RGBD-RSS2010/papers/morris-  
RGBD10-3D-indoor.pdf](http://www.seattle.intel-research.net/RGBD/RGBD-RSS2010/papers/morris-RGBD10-3D-indoor.pdf)

Murphy, D., and Cycon, J. (1998). Applications for mini VTOL UAV for law enforcement. Retrieved 10 September 2012 from [http://www.dtic.mil/cgi-  
bin/GetTRDoc?AD=ADA422459](http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA422459).

Robertson, D.P., Cipolla, R. (2009). Structure from Motion. In Varga, M. (Eds.), Practical Image Processing and Computer Vision, John Wiley. From [http://mi.eng.cam.ac.uk/~cipolla/publications/contributionToEditedBook/2008-  
SFM-chapters.pdf](http://mi.eng.cam.ac.uk/~cipolla/publications/contributionToEditedBook/2008-SFM-chapters.pdf)

- Runge, H., Rack, W., Ruiz-Leon, A., Hepperle, M. (2007). A Solar Powered HALE-UAV for Arctic Research. Preprint from the 1<sup>st</sup> CEAS European Air and Space Conference.
- Samson, C., Straznicky, P., Laliberté, J., Caron, R., Ferguson, S., and Archer, R. (2010). Designing and building an unmanned aircraft system for aeromagnetic surveying. 80<sup>th</sup> Annual Meeting of the Society of Exploration Geophysicists, Denver, CO, 17-22 October. SEG Expanded abstracts **29**, 1167-1171.
- Sofge, E. (2009). Britain's Police Drone: Could it stop next terror plot? Retrieved from Popular Mechanics website:  
<http://www.popularmechanics.com/technology/military/4218776>.
- Smith, L.I. (2002). A tutorial on Principal Component Analysis. Retrieved from [http://www.cs.otago.ac.nz/cosc453/student\\_tutorials/principal\\_components.pdf](http://www.cs.otago.ac.nz/cosc453/student_tutorials/principal_components.pdf)
- Suhr, J.K. (2009). Kanade-Lucas-Tomasi (KLT) Feature Tracker. Computer Vision (EEE6503). Retrieved from <http://web.yonsei.ac.kr/jksuhr/articles/Kanade-Lucas-Tomasi%20Tracker.pdf>.
- 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.
- Trucco, E. and Verri, A. (1998). *Introductory Techniques for 3-D Computer Vision*. Upper Saddle River (New Jersey): Prentice-Hall Inc.
- Wegener, S.S., Schoenung, S.M., Totah, J., Sullivan, D., Frank, J., Enomoto, F., Frost, C., and Theodore, C. (2004). UAV Autonomous Operations for Airborne Science Missions. American Institute of Aeronautics and Astronautics.

Whitehead, K. (2010). Unmanned Aerial Vehicles for Glaciological Studies. *GIM International*. 24(10).

## **Appendix A: Basics of structure from motion**

### ***Feature points***

Trucco and Verri (1998) define image features as “local, meaningful, detectable parts of the image”. As such, image features are interesting scene elements such as uniform regions (surfaces) or sharp intensity variations resulting from contours (edges or corners). Detection algorithms produce feature descriptors that specify properties of the feature such as position. Although there are algorithms that use surfaces and edges, ‘corners’ are easier to characterize in mathematical terms and are the type of feature used by ViDAR. ‘Corners’, in this case, means not only intersection points of object contours and geometric corners, but also local intensity patterns.

### ***Correspondence problem***

To determine corresponding points, an algorithm computes a similarity measure between the specific feature point in the first image and the feature points in the second image. It then selects the feature point from the second image that maximizes the similarity measure.

### ***Triangulation in a simple stereo system***

The diagram (Figure A1) shows a stereo system involving two pinhole cameras with coplanar image planes:  $I_l$  and  $I_r$ .  $O_l$  and  $O_r$  are the centers of projection. To determine the real world positions of P and Q, a ray is projected from the respective center of projection through the image points. The intersection of the rays from the two images sets the 3D location of the feature. Triangulation depends on the correspondence issue: if  $p_l$  and  $q_r$  were determined to be corresponding points, the location of the 3D point would be quite different.

### ***Point clouds***

A point cloud is simply the collection of the (reconstructed) 3D points that delineate features of the scene.

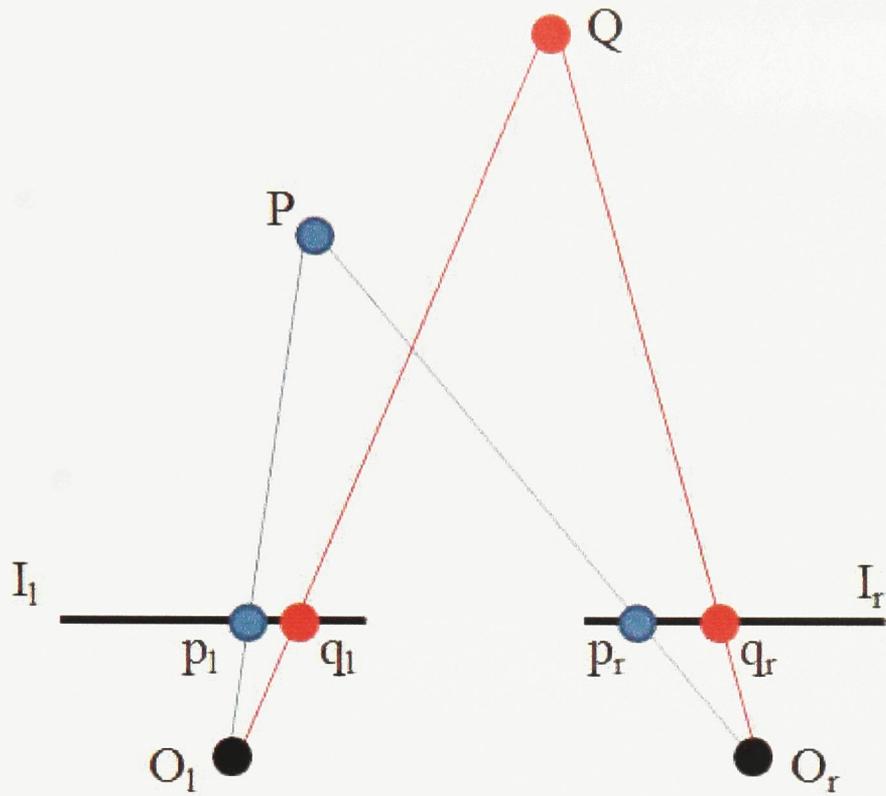


Figure A1: 3D reconstruction from a simple stereo system.

## Appendix A: Manual measurements from mine site

*East wall strike and dip measurements from north to south*



Figure B1



Figure B2

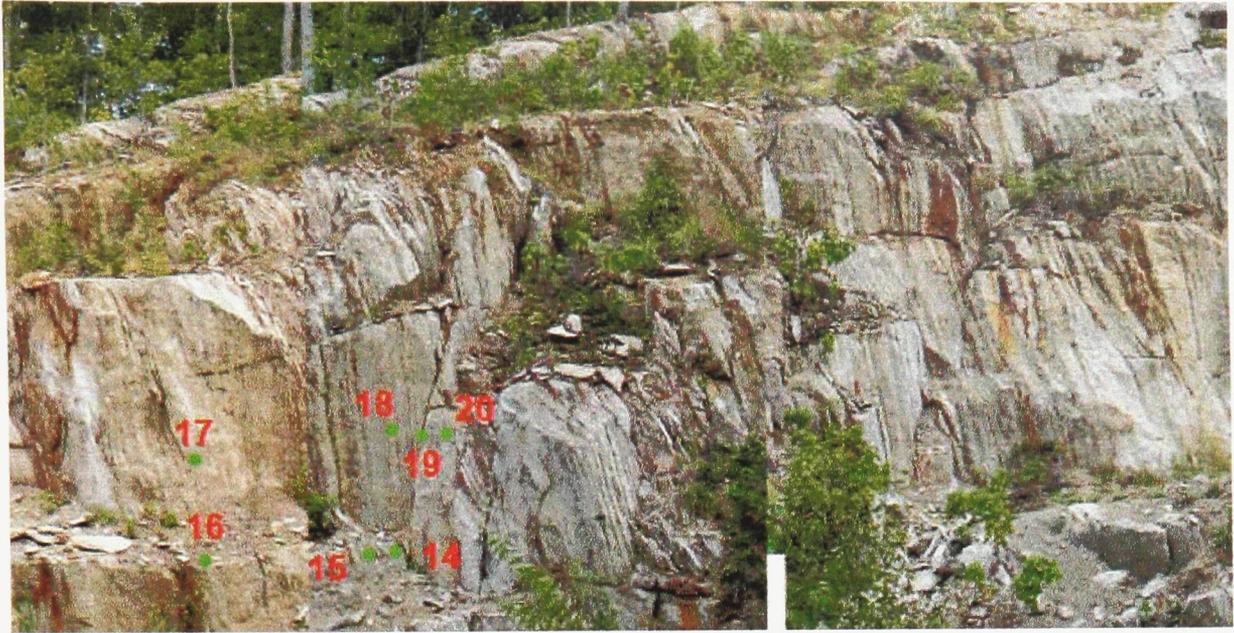


Figure B3



Figure B4

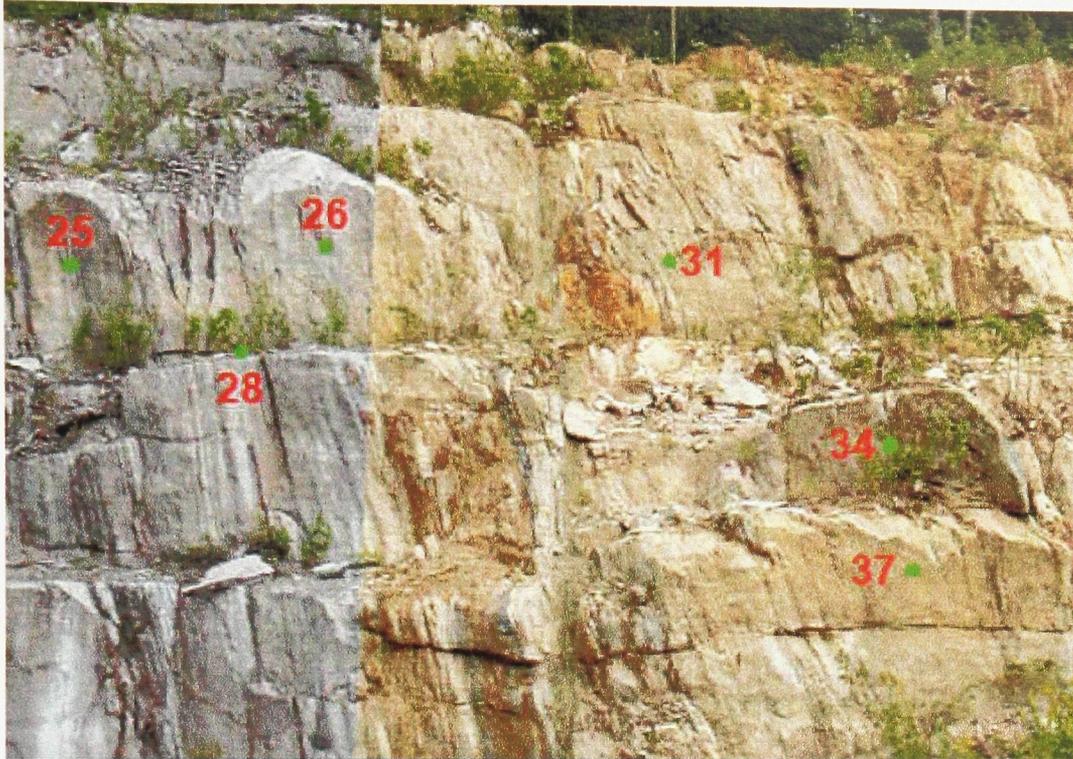


Figure B5



Figure B6



Figure B7

Table: Manual Measurements for East wall

Measurement Number	Strike	Dip	Additional notes
1	106	70	
2	108	10	subhorizontal
3	206	68	
4	114	12	subhorizontal
5	238	74	
6	39	74	
7	39	72	
8	4	78	
9	283	80	
10	186	88	
11	24	72	
12	150	12	subhorizontal
13	132	11	subhorizontal
14	40	80	
15	347	69	
16	207	26	sloping surface
17	354	70	
18	358	88	
19	278	71	
20	178	89	
21	140	90	
22	44	75	
23	94	15	sloping surface
24	196	89	
25	333	71	
26	130	82	
27	94	14	subhorizontal
28	205	33	subhorizontal
29	105	11	
30	172	87	
31	132	88	
32	121	65	
33	180	85	
34	331	76	
35	150	75	Lighter colour
36	145	84	Green colour
37	348	85	
38	332	78	
39	322	72	
40	350	74	
41	95	3	subhorizontal
42	346	83	

***West wall strike and dip measurements from north to south***



Figure B8

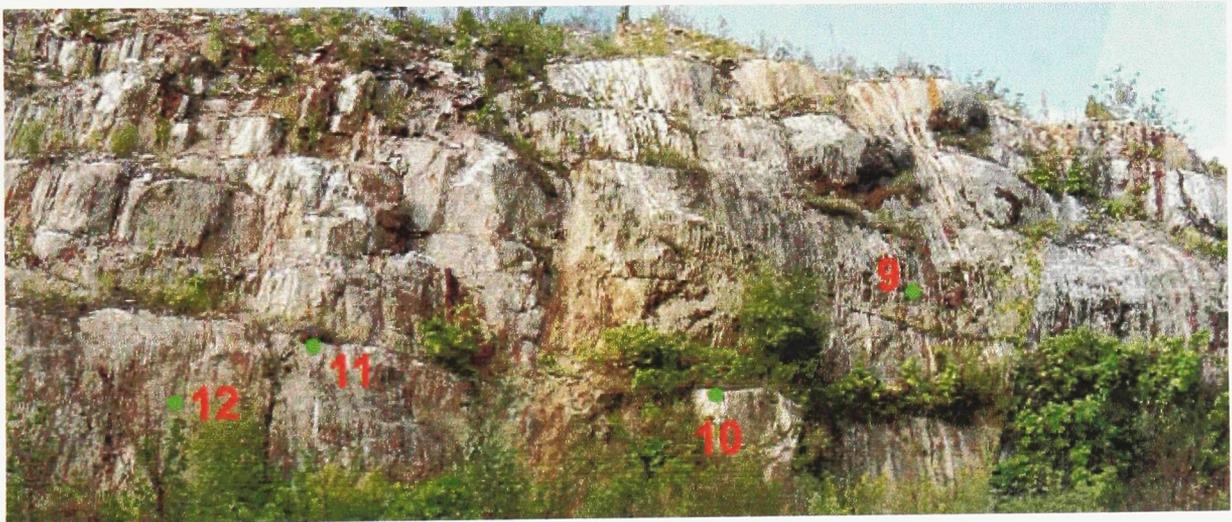


Figure B9



Figure B10

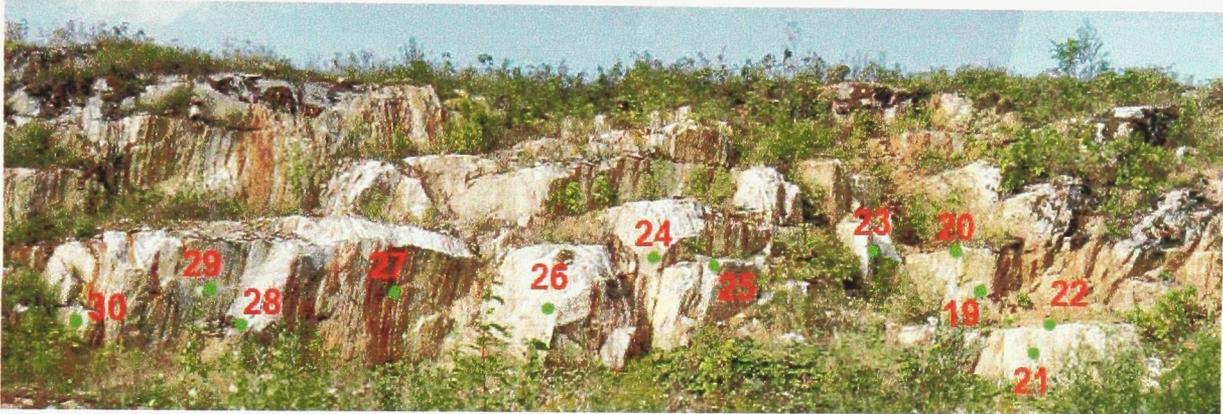


Figure B11



Figure B12

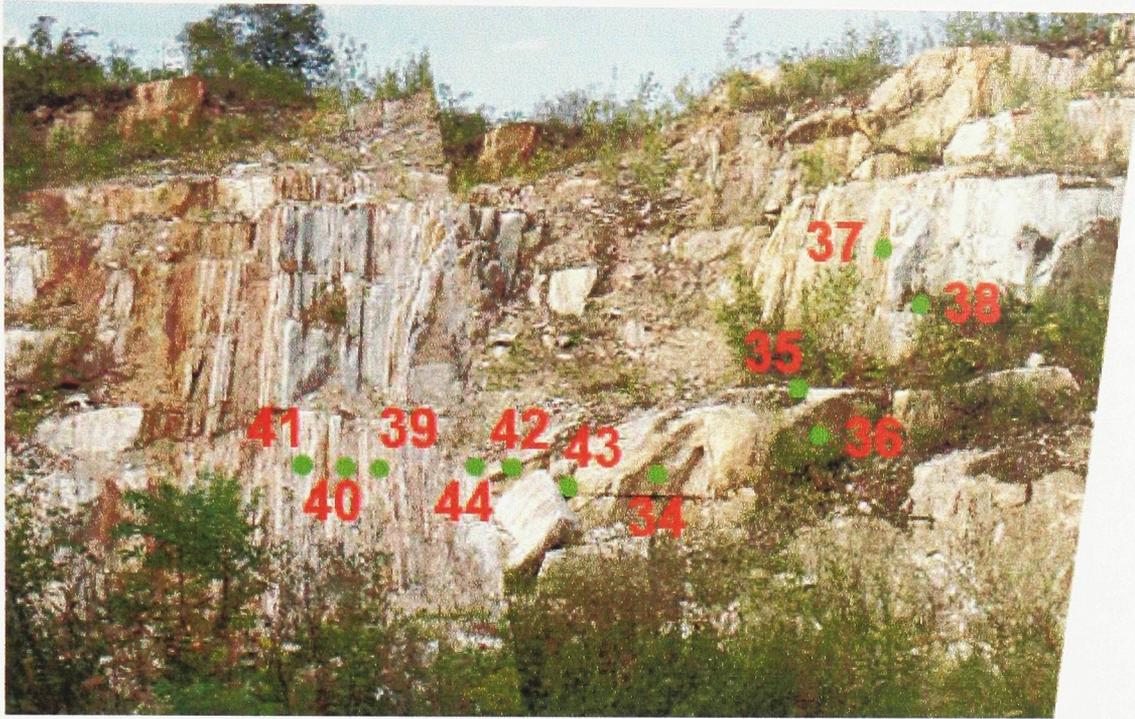


Figure B13

Table: Manual Measurements for West Wall

Measurement Number	Strike	Dip	Additional notes
1	49	79	
2	311	79	
3	220	88	
4	131	79	
5	60	78	
6	317	86	
7	47	79	
8	141	80	
9	327	72	
10	209	16	subhorizontal
11	274	14	subhorizontal
12	340	78	
13	274	18	subhorizontal
14	339	71	
15	6	70	
16	334	79	
17	242	80	
18	218	6	subhorizontal
19	3	82	
20	178	5	subhorizontal
21	23	80	
22	204	7	subhorizontal
23	101	81	
24	19	82	
25	186	10	subhorizontal
26	43	72	
27	136	85	
28	29	84	
29	134	81	
30	18	74	
31	329	64	
32	335	73	
33	349	60	
34	140	84	
35	206	4	subhorizontal
36	136	77	
37	37	78	
38	224	16	subhorizontal
39	336	83	
40	70	81	
41	18	65	
42	335	88	

43	242	85	
44	204	11	subhorizontal

***East wall length measurements north to south***

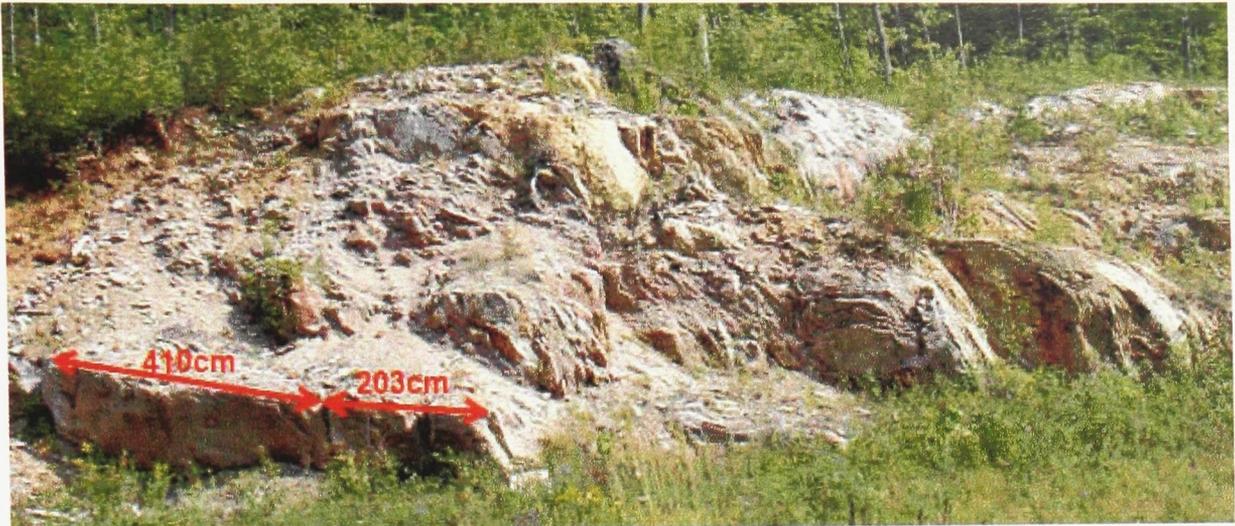


Figure B14

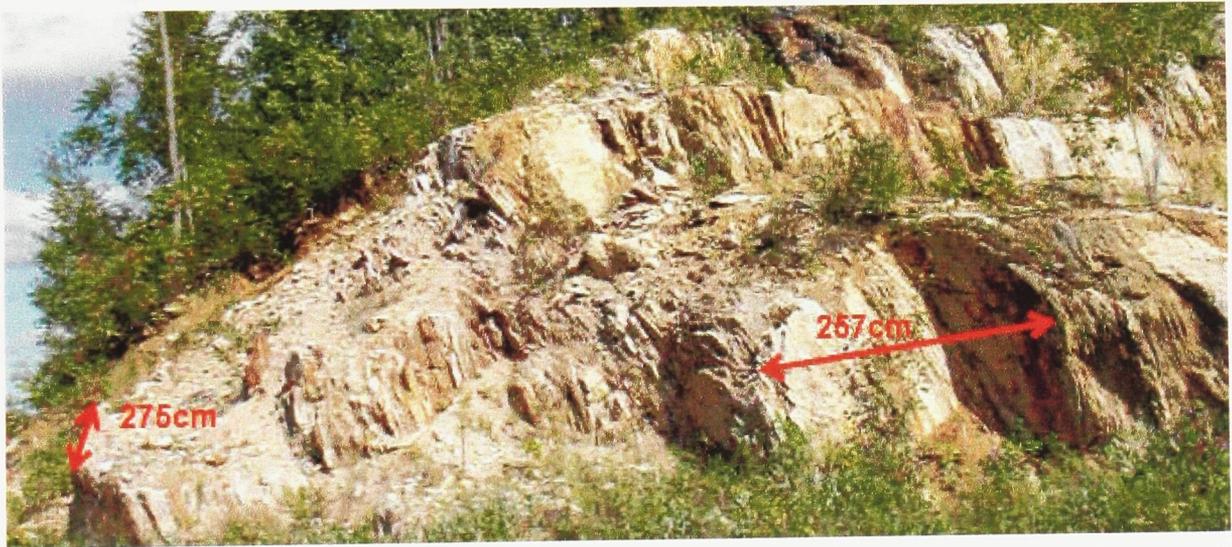


Figure B15

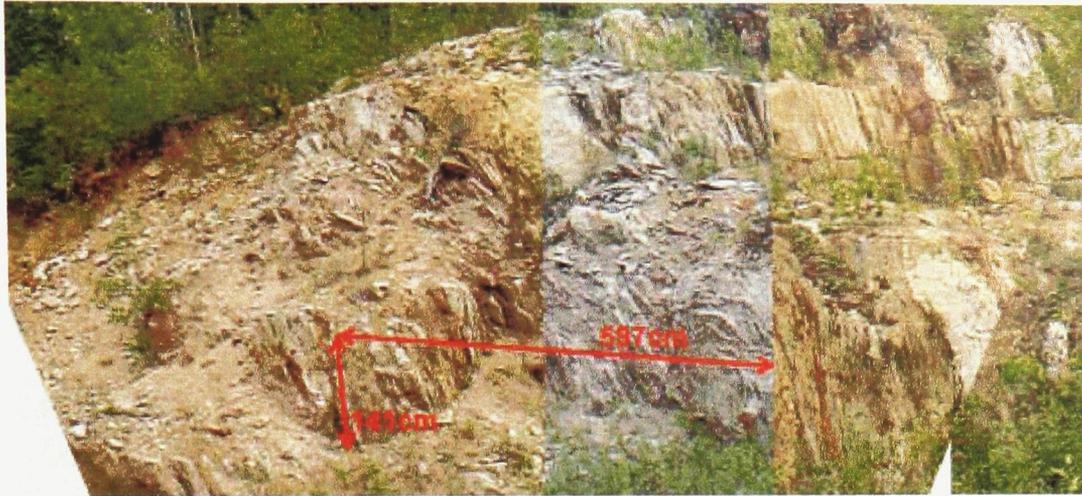


Figure B16



Figure B17



Figure B18

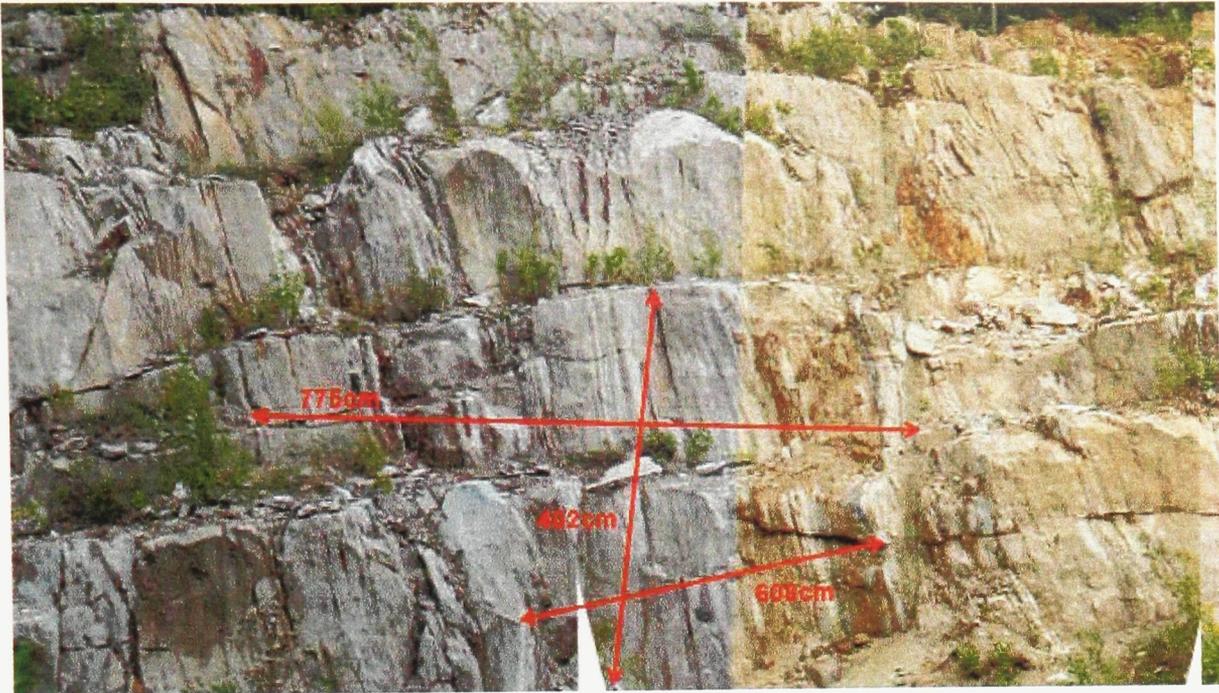


Figure B19

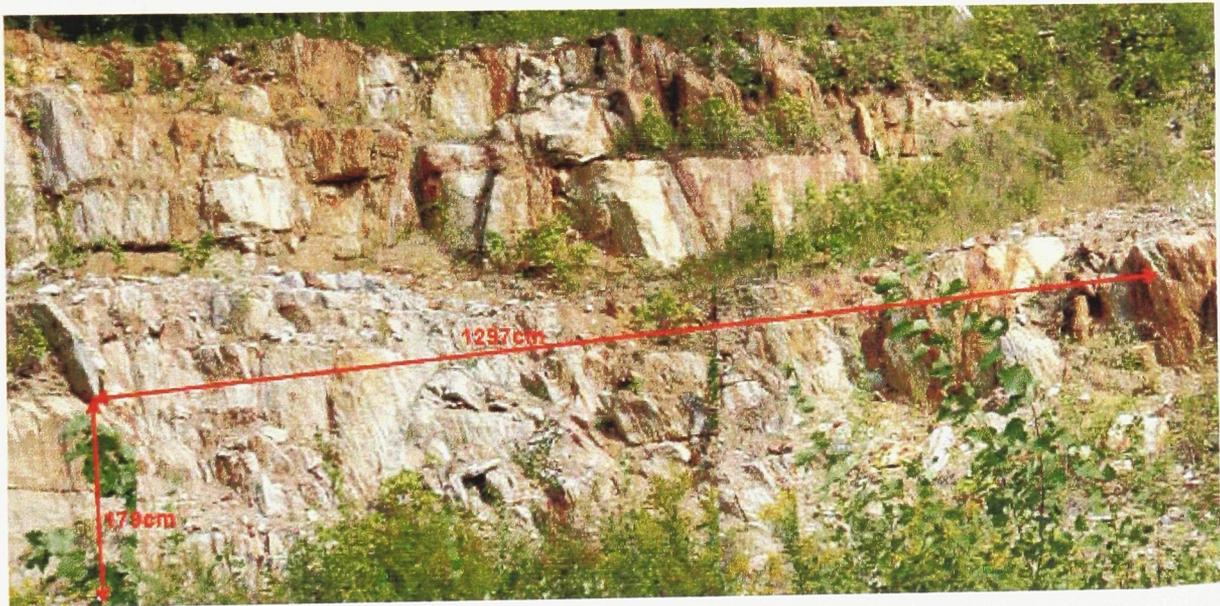


Figure B20



Figure B21

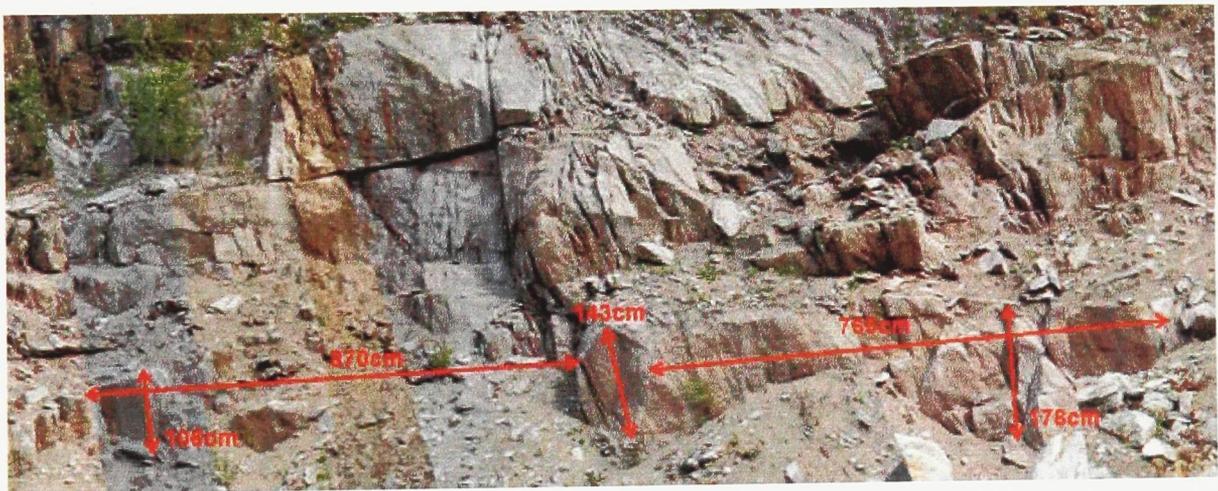


Figure B22



Figure B23

***West wall length measurements south to north***

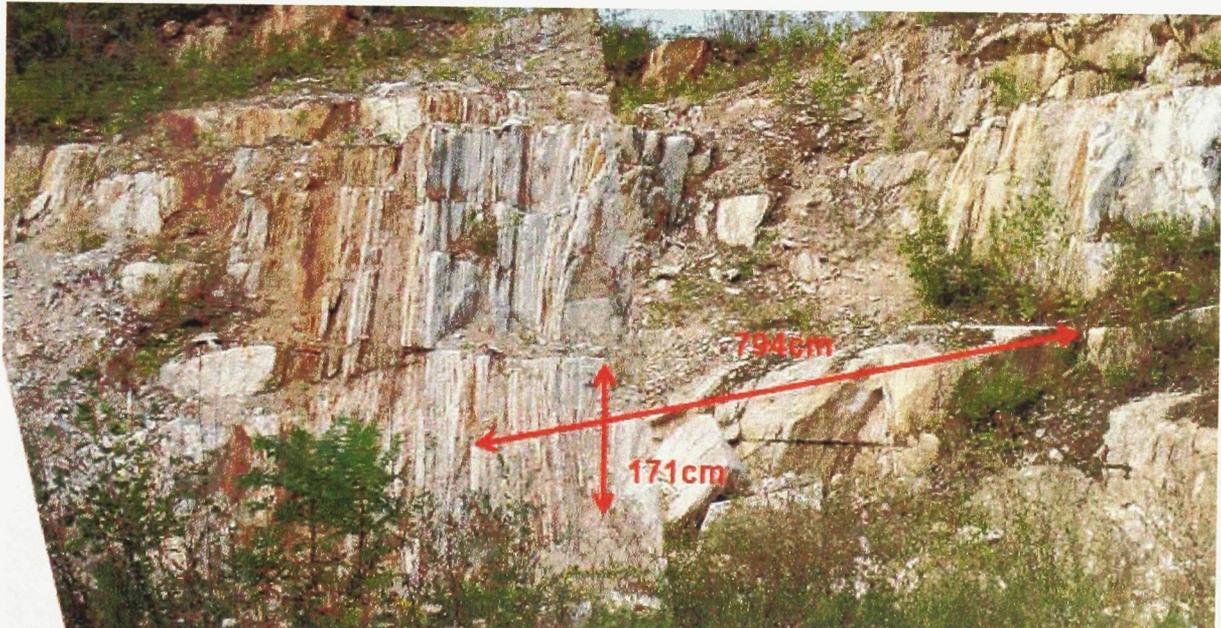


Figure B24

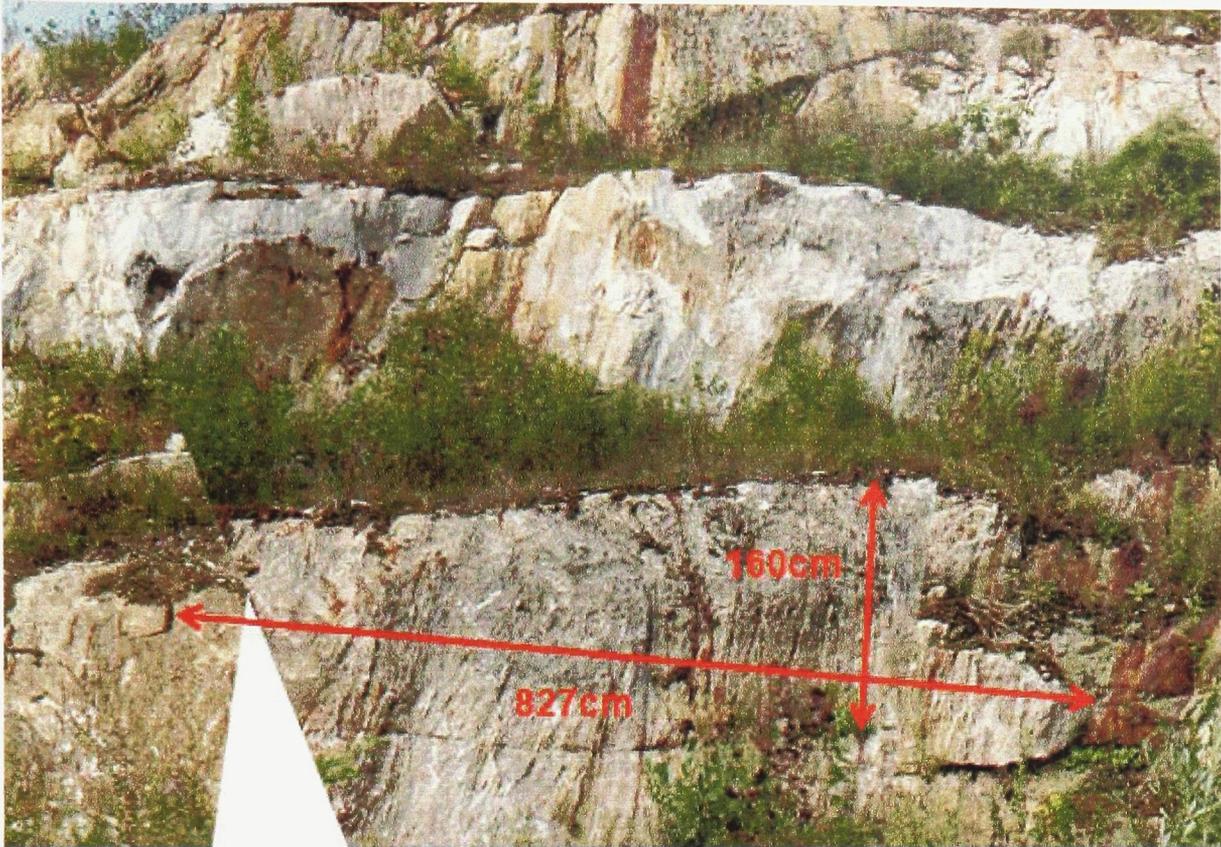


Figure B25

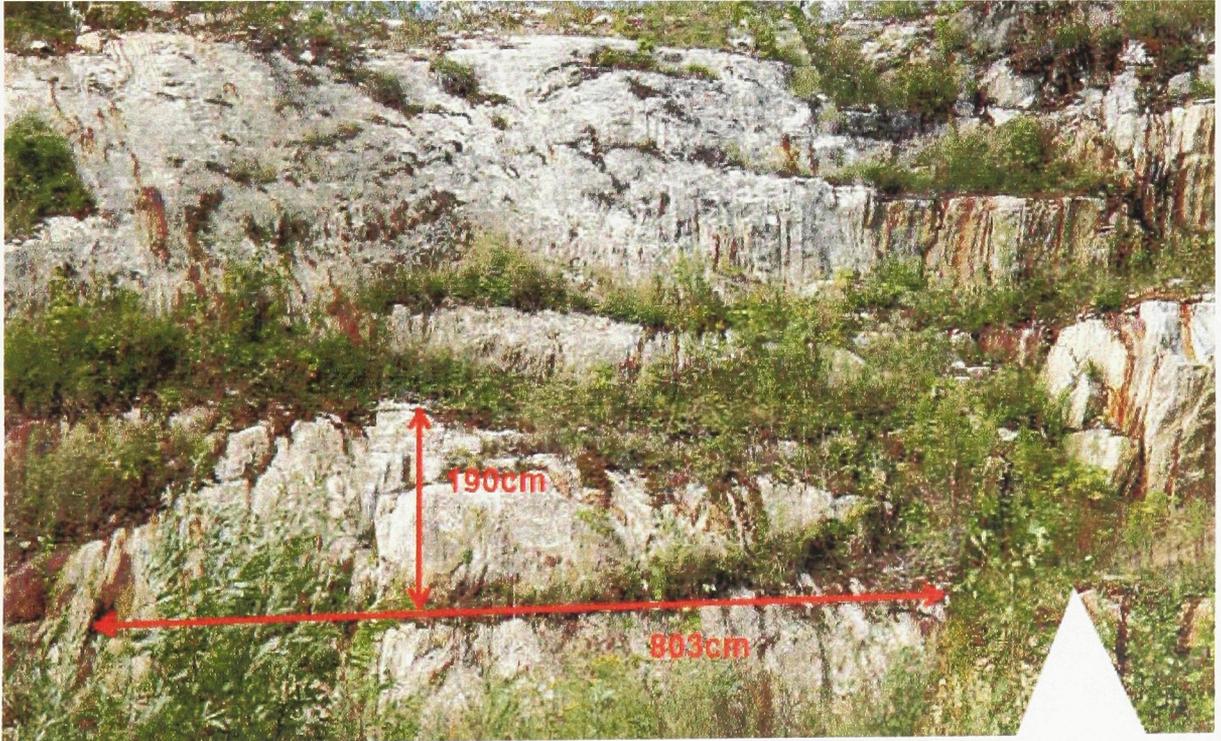


Figure B26



Figure B27

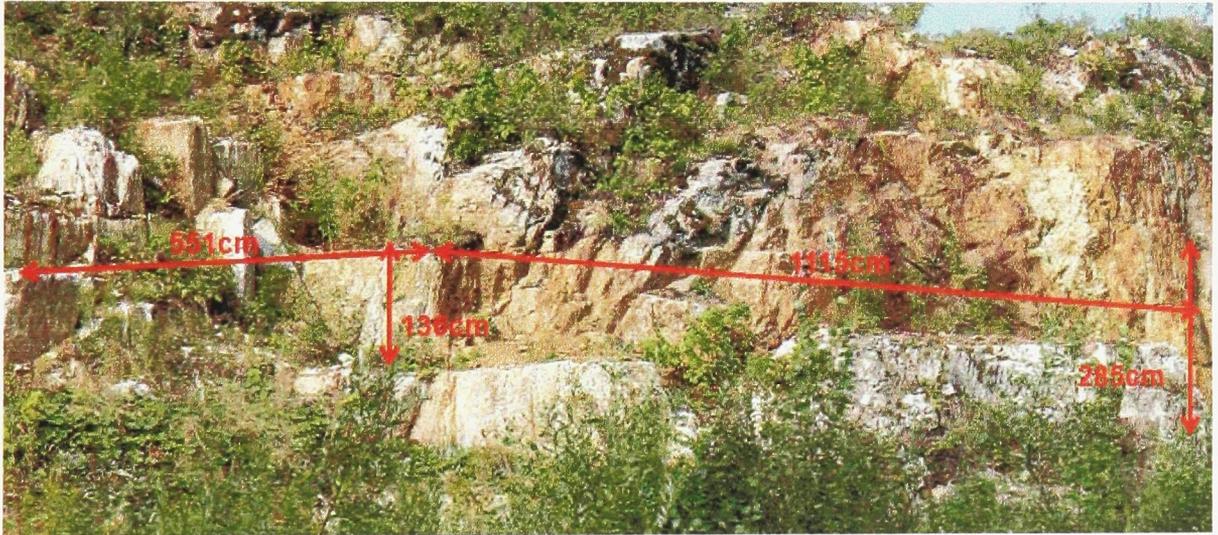


Figure B28

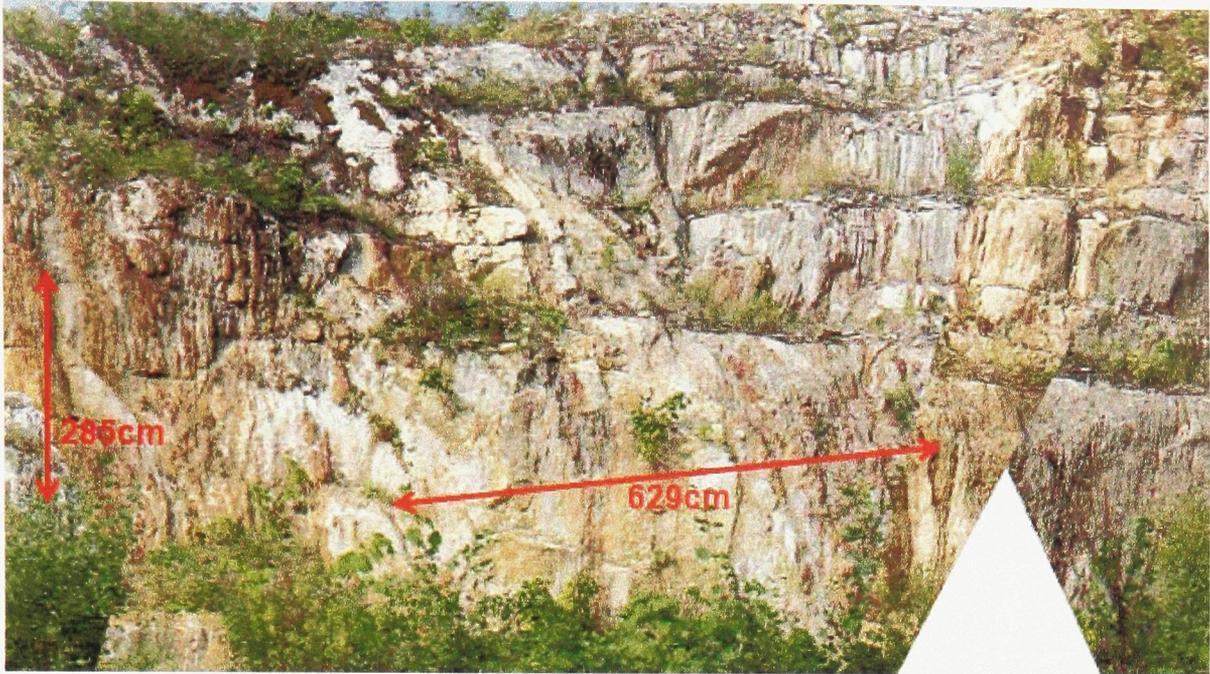


Figure B29



Figure B30

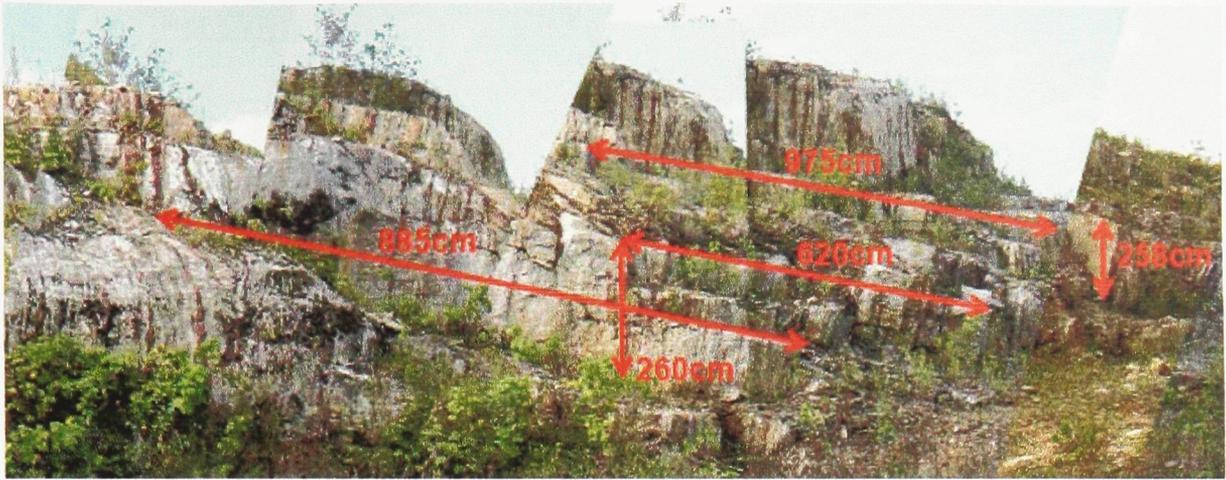


Figure B31