Application and Accuracy of Structure from Motion Computer Vision Models with Full-Scale Geotechnical Field Tests

Logan M. Palmer[1], S.M. ASCE, Kevin W. Franke[2], M. ASCE, R. Abraham Martin[1], Brendan E. Sines[1], S.M. ASCE, Kyle M. Rollins[3], M. ASCE, and John D. Hedengren[4]

[1]Undergraduate student researcher, Brigham Young Univ., Provo-UT 84602
[2]Assistant Professor, Dept. of Civil and Environ. Eng., Brigham Young Univ., Provo-UT 84602
[3]Professor, Dept. of Civil and Environ. Eng., Brigham Young Univ., Provo-UT 84602
[4]Assistant Professor, Dept. of Chemical Eng., Brigham Young Univ., Provo-UT 84602

ABSTRACT: Structure from motion (SfM) computer vision is a relatively new technology that allows engineers to reconstruct a three-dimensional (3D) model of a given scene using two-dimensional digital photographs captured from a single, moving camera. SfM computer vision provides an economic and user-friendly alternative to other 3D scene-capture and modeling tools such as light distance and ranging (LiDAR). Although the resolution and accuracy of laser-based modeling methods are generally superior to vision-based modeling methods, the economic advantages associated with the latter may make it a useful and practical alternative for many geotechnical engineering applications. Although other engineering disciplines have investigated the potential usefulness of SfM computer vision for years, its application to geotechnical engineering generally remains unexplored. Researchers are currently investigating the application of this technology to select full-scale geotechnical field experiments and assessing its potential usefulness as a high-resolution instrumentation/monitoring tool. This paper presents preliminary computer vision results and findings from these studies. The field experiments, as well as the hardware and software details used to develop 3D SfM computer models of the experiments are summarized. The developed 3D models are presented, and displacements measured in the models are compared against ground truth to evaluate accuracy. Observed advantages and limitations of SfM computer vision are discussed, and several potentially useful applications of the technology in geotechnical engineering are listed.

INTRODUCTION

This paper explores and documents the accuracy and speed of Structure from Motion (SfM) computer vision modeling in geotechnical engineering applications. SfM modeling is a cost-effective way to quickly generate accurate computer models of objects and terrain and offers significant contributions to the geotechnical engineering industry. This newer and lesser known technique for data collection utilizes optical sensors (e.g., regular commercial cameras) to capture images and generate point clouds with sub-centimeter resolution in a relatively short amount of time. Such an approach allows engineers with the appropriate tools to develop accurate models and to obtain precise measurements from those models without the high cost of other remote sensing methods such as LiDAR. Furthermore, while LiDAR models are generally recognized as having superior resolution and accuracy over SfM computer vision models, the resolution and accuracy that can be obtained from the SfM approach may be suitable for most geotechnical engineering applications. The purpose of this paper is to evaluate some of these potential geotechnical applications of SfM by investigating its effectiveness in modeling a few full-scale geotechnical research sites.

SfM modeling is a technique in which a set of images of a scene, taken from different viewpoints, are used to recover the 3D structure of the scene along with the camera parameters (Koenderink, and Van Doom 1991). Essentially, SfM involves the act of taking two-dimensional images and projecting them to three-dimensional models. SfM software takes groups of pictures of an object of interest and generates 3D computer models for analysis. Based on recent research, the use of high resolution cameras can produce 3D models with accuracies and resolutions that achieve sub-centimeter levels.

Currently, one of the most well-known and frequently used methods of data collection and modeling is LiDAR, or 3D laser scanning. This method uses a three dimensional mapping technology that employs a laser and a rotating mirror or housing to rapidly scan and image volumes and surficial areas (Kemeny and Turner 2008). This method of modeling can create models with fine detail but the costs are significantly higher. A LiDAR system is approximately $70,000 to $150,000 (based on 2008 prices) with relatively high maintenance costs (Kemeny and Turner 2008). SfM is an alternative modeling technique that offers competitive accuracy at a fraction of the cost.

BACKGROUND AND RELATED WORK

SfM computer vision techniques continue to be refined, and applications in various engineering geology fields have already begun to emerge. Golparvar-Fard (2010) and Golparvar-Fard et al. (2009) presented methods of efficient reconstruction of 3D models from unordered image sets, and apply the models created using this method to track the construction progress of buildings. Fathi and Brilakis (2011) presented a method for creating sparse 3D point clouds of infrastructure using stereo videogrammetry. Castillo et al. (2012) included SfM modeling in a study evaluating the accuracy of several methods for measuring gully erosion. James and Robson (2012) demonstrated the use of SfM in coastal cliff erosion surveys, and compared their results to terrestrial laser scans. They reported relative model precisions of about 1:1000 using consumer grade SLR cameras. Niethammer et al. (2012) developed computer vision models of the Super-Sauze landslide in France, and measured displacements over a period of 18 months.

DESCRIPTION OF GEOTECHNICAL FIELD TESTS

Two geotechnical field tests sites were used for the SfM modeling analysis presented in this paper. The first test site is located adjacent to the Salt Lake City International Airport in Utah. The purpose of this test was to better understand the passive soil resistance behind skewed pile caps. For this test, a 3.35-meter wide pile cap was supported by six 0.3-meter diameter steel pipe piles. Two hydraulic actuators (600 kip extension, 450 kip contraction) were installed in parallel on the back side of the pile cap to produce horizontal loading. The rams were then used to push the pile cap 2.54 cm, 3.81 cm, and 7.62 cm into the compacted structural fill in front of the pile cap. Passive resistance during the test was measured, and displacements of the pile cap were measured and recorded. After each push of the pile cap, ground surveys were performed to map observed cracks in the soil and to measure any possible heaving. SfM models were used with the pile cap test site to evaluate the ability of SfM to provide accurate object dimensions of commonly used geotechnical engineering materials (e.g., steel, concrete, and soil).

The second test site is located in Christchurch, New Zealand. The purpose of this test site was to better understand liquefaction-induced drag loads on piles and subsequent pile settlements. Three instrumented 0.6-meter diameter test piles were driven into liquefiable soils and were loaded with weights to simulate static loading. Following pile installation, liquefaction was induced in the soil from depths between 3 to 10 meters below the ground surface using controlled blasting while the build-up and dissipation of excess pore pressure in the ground was monitored. The test piles were monitored with strain gauges to measure the change in skin friction along the pile length as the soil liquefied and then reconsolidated. Liquefaction-induced settlement in the surrounding soil was measured to investigate its effect on observed pile settlements and drag loads. SfM models were developed both before and after liquefaction induction to investigate the accuracy and usefulness of SfM change and settlement detection capabilities in geotechnical engineering field applications.

DEVELOPMENT OF SfM COMPUTER VISION MODELS

Three dimensional computer reconstruction of each test site was undertaken using a combination of structure-from-motion and multi-view stereopsis (MVS) techniques. The SfM-MVS process results in a 3D point cloud similar to that produced using sensors such as LiDAR. The density of the point cloud is a function of image resolution and camera object positioning (Kemeny and Turner 2008). One hundred and twenty seven images of the Salt Lake City test site were collected using a Sony NEX-5R camera, which has a resolution of 16 megapixels. The NEX-5R camera has an adjustable focus lens, which was set at 25mm for the Salt Lake City test site. Images of the New Zealand test site were collected using a Nikon Digital SLR D7100 camera. Three hundred and twenty two images were collected before blasting, and five hundred and ten images were collected after blasting. The New Zealand images were captured at 24 megapixels, using a fixed focal length of 35 millimeters. All image processing and 3D reconstruction was performed on an MSI GT70 laptop with 24 GB of RAM, and an NVIDIA GeForce GTX 780M graphics card.

The creation of dense 3D point clouds from imagery is a multistep process, and variety of workflows and software packages exist for the purpose (Harwin and Lucieer 2012). These include commercial packages such as PhotoScan (AgiSoft 2014), and PhotoModeler (SStems Inc., 2014), as well as free platforms such as Photosynth (Microsoft Corporation 2014), and open source alternatives such as Bundler (Snavely 2010). The software used and the processing steps taken to produce the point clouds analyzed in this study are presented below.

Feature Extraction and Structure from Motion for Point Cloud Development

The first step in the 3D reconstruction process is the extraction and matching of features in the images. This is done using an automatic algorithm called SIFT (Scale Invariant Feature Transform) (Harwin and Lucieer 2012). An image feature is an area of image texture containing patterns likely to be recognizable in other images. Matching these features across an image set produces a sparse point cloud, typically comprised of only a few tens of thousands of points, as well as the back-calculated position and orientation of the camera. Feature extraction and sparse point cloud construction were implemented using the software package VisualSfM (Wu 2011).

Point Cloud Enhancement and Model Meshing

The sparse point cloud and camera positions supplied by VisualSfM were subsequently processed using MVS techniques to filter, enhance, and densify the point cloud. The MVS technique utilizes an algorithm that outputs groupings of triangular patches that cover the surfaces generated from the gathered images and add further detail to the point cloud. Following point cloud enhancement, a 3D mesh was developed using the free software package CMPMVS (Jancosek and Pajdla 2011). Developing a 3D mesh is typically the most computationally intensive step of the process, and requires a high end graphics card for parallel GPU processing. Meshed models are generally better for presentation purposes and qualitative assessments. Fortunately, most change detection and measurement activities that might be useful to geotechnical engineers do not require a meshed model and can be performed solely with an enhanced point cloud.

Processing Time

Processing time is generally a function of the computing and video hardware that is available. Using the hardware described above, the Salt Lake City test site photographs required approximately 12 hours to process and to produce the textured mesh. Processing time for the Christchurch pre-blast photographs was approximately 24 hours, while the processing time for post-blast photographs was approximately 27 hours.

Scaling

3D reconstructed computer models require scaling to obtain useful measurements in actual units. This was done by measuring point-to-point distances on the model, comparing them with actual measurements taken at the site, and computing a scaling factor for the model. The point cloud was then scaled and analyzed using the open software package CloudCompare (Girardeau-Montaut 2012).

Figure 1 presents an original photograph of the Christchurch test site, along with three of the different SfM model formats produced by the reconstruction process.

COMPARISONS AND RESULTS

Salt Lake City Model

At the Salt Lake City test site, the primary goal of SfM modeling was to evaluate the potential for SfM models to make accurate measurements of objects at geotechnical field sites. The SfM model was scaled using the 1.22-meter width cap dimension of the reinforced concrete drilled shaft used as a reaction from the hydraulic actuators (Figure 2). After model scaling, the actual dimensions of the concrete pile cap, reinforced concrete shafts and the I-beam from the test site were compared with dimensions obtained from the SfM model. A screenshot of the SfM model and the locations of the obtained measurements are presented in Figure 2. Comparisons between the actual dimensions and the model dimensions for the Salt Lake City test site are presented in Table 1.

FIG. 1. Photograph of Christchurch test site (top left); screenshot of sparse point cloud (top right); screenshot of enhanced SfM point cloud with wireframe (bottom left); screenshot of fully meshed SfM model (bottom right)

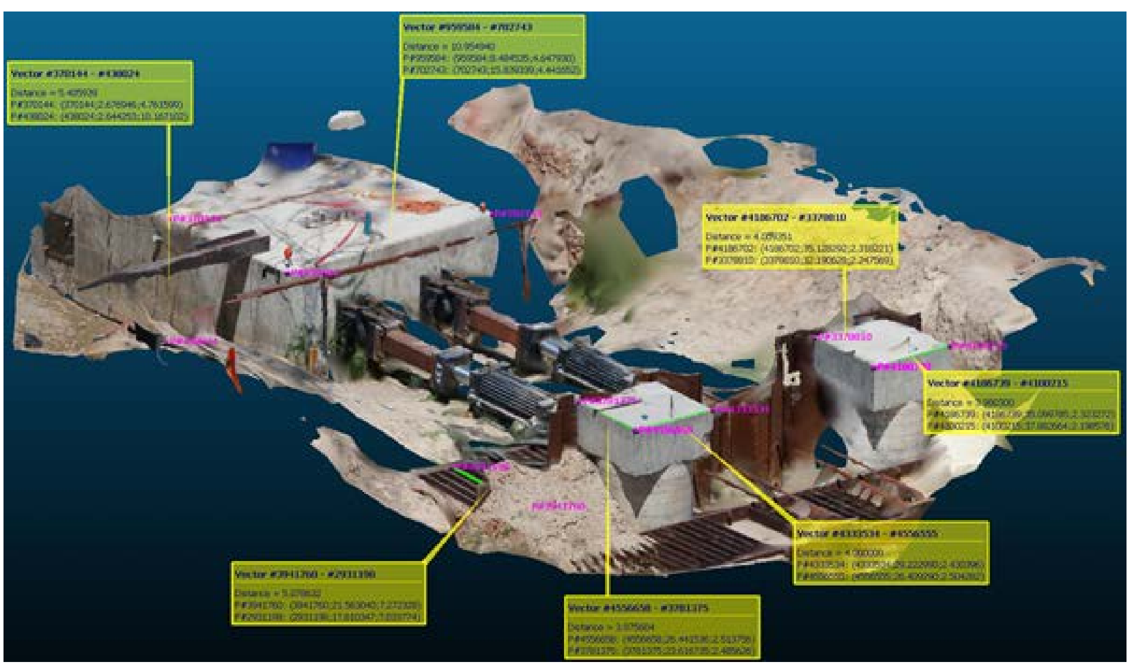

FIG. 2. Screenshot of the Salt Lake City pile cap SfM model. Captions indicate locations were measurements were obtained and compared against actual dimensions

Table 1. Dimension Comparisons for Salt Lake City test site

Object	Actual Dimension (m)	Model Dimension (m)	Error (%)
I-Beam length	1.52	1.55	1.6
Pile cap width	3.35	3.34	-0.5
Pile cap depth	1.68	1.65	-1.6
Left drilled shaft length	1.22	1.21	-0.5
Right drilled shaft length	1.22	1.22	0.3
Right drilled shaft width	1.22	1.21	-0.5
		Average:	-0.2

As shown in Table 1, the average observed dimensional measurement error for the six evaluated test site objects was -0.2%. Additionally, all measured dimensions from the SfM model were within 3-mm of the actual dimensions at the test site. These small values are quite remarkable when one considers that all of the data necessary to produce the SfM model (i.e., 127 photographs) was collected in less than 3 minutes. These results demonstrate the ability of the SfM modeling method to generate accurate models that may be useful in many different engineering applications. It is recognized that some engineering applications may require greater accuracy than 0.2%, which would certainly require the use of traditional instrumentation and/or more accurate remote sensing like LiDAR. However, the observed measurement accuracy of the SfM model observed in this study would likely be sufficient for most geotechnical engineering analyses.

<u>New Zealand Model</u>

For the New Zealand test site, pre- and post-blast SfM models were developed and scaled. A change detection analysis was performed to measure the amount of relative settlement between the tops of the piles and the ground surface due to liquefaction. Figure 3 presents a screenshot of the post-blast SfM model of the test site. The software package CloudCompare (Girardeau-Montaut 2012) was used to compare the two point clouds and measure vertical changes (i.e., displacements or settlements) as referenced from the top of the piles. These measured changes can be represented as heat maps, such as the one presented in Figure 4. Evaluation of the change detection results indicated that the largest relative settlements occurred at the corners of the weights, with one corner settling significantly more than the others. This observation is consistent with observations made and documented by field personnel who performed the testing.

FIG. 3. Screenshot of the Christchurch meshed SfM model (after liquefaction blasting)

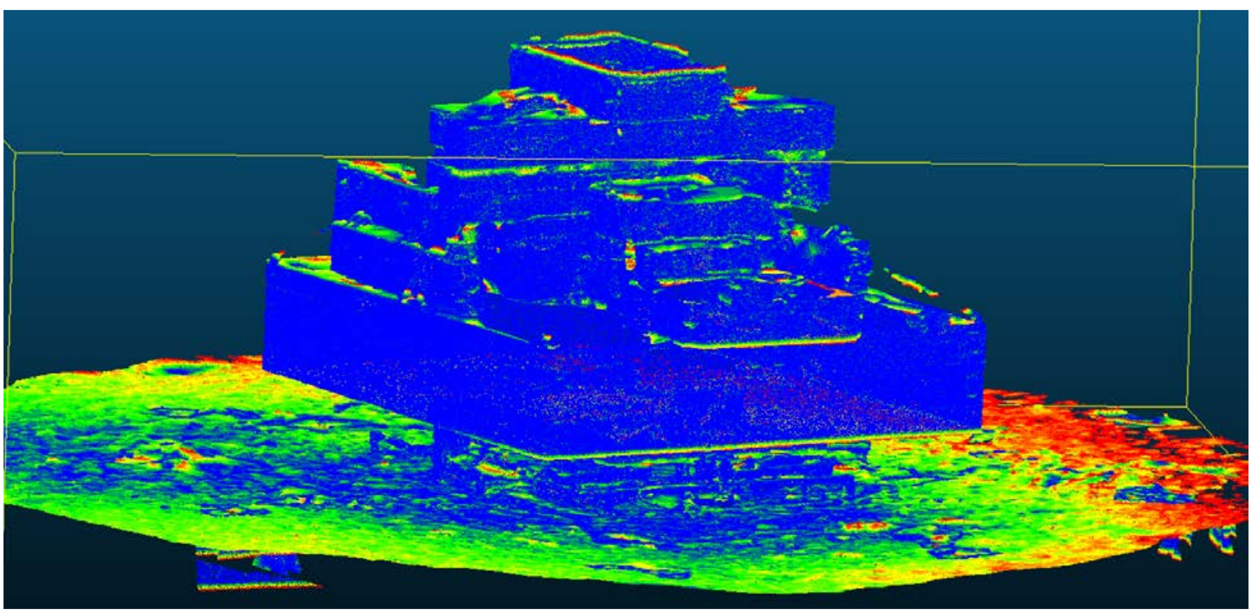

FIG. 4. Screenshot of the Christchurch change detection model showing vertical changes between pre- and post-liquefaction blasting

Table 2 presents the comparisons between the relative pile settlements obtained from the change detection analysis and the actual relative pile settlements measured from string potentiometers in the field. Three string potentiometers were used on each pile to generate more accurate results. The average value from the three potentiometers is what is reported in Table 2.

Table 2. Comparison between actual and measured relative pile settlements at Christchurch test site

Pile	Actual Relative Settlement (cm)	Measured Relative Settlement (cm)	Error (%)
#1 (14-m length)	5.18	5.44	4.9
#2 (8.5-m length)	8.1	8.92	10
#3 (12-m length)	4.09	3.99	-2.5
		Average:	4.1

As shown in Table 2, the average error in measuring relative pile settlements using the SfM change detection analysis was only 4.1%. Furthermore, the SfM change detection analysis was consistently within about 8-mm of the actual relative pile settlements. Much of this error can be attributed to the manual alignment of the pre- and post-blast models necessary to perform change detection. Had more consistent anchor points between the pre- and post-blast SfM models been available for the change detection analysis, it is possible that the observed error between actual and measured pile settlements could have been significantly reduced.

DISCUSSION

Based on the accuracy results from this study, there appears to be significant potential and many possible applications for SfM modeling in the field of geotechnical engineering. Geotechnical and geomatic professionals are already using SfM computer reconstructions of project sites to develop digital elevation and/or terrain models. The rapid advancement and increased use of unmanned aerial vehicles (UAVs) in modern society introduces numerous possible engineering and site characterization applications for computer vision models. For example, computer vision models developed from UAVs could be used two develop 2D or 3D ground surface profiles for various engineering analyses such as watershed studies or slope stability analysis. Investigators could use SfM computer vision to rapidly and accurately perform detailed geotechnical site reconnaissance following extreme events such as earthquakes or floods. Computer vision models could be used to develop a "virtual site reconnaissance" in which many engineers can simultaneously evaluate a given site and look for evidence of significant site details such as marked utilities and potential drilling obstructions. Engineers could potentially combine computer vision models of the ground surface with subsurface investigation data to develop 3D renderings of both the ground surface and the subsurface for a given site. Such renderings could aid engineers in explaining analysis results to non-technical clients and/or to aid in marketing presentations.

SfM computer vision models offer superior speed, satisfactory accuracy for most engineering applications, and significantly reduced costs when evaluated against other popular forms of remote sensing such as LiDAR. The largest costs associated with SfM modeling typically involve a quality camera, the hardware necessary to process and manipulate/edit the computer vision models within a reasonable amount of time, and the desired software. Additionally, data collection for SfM modeling is much more rapid than most other forms of remote sensing. Several hundreds of images of a site can be collected in a matter of minutes rather than hours. The image capture also does not require extensive training, and even basic field personnel can learn how to effectively gather images. The development of the models themselves is exceptionally fast and models can typically be generated without about 24 hours. Also, as seen in Figure 2, SfM modeling can capture many of the rich details of a site that can be valuable in characterizing a site and helping to make engineering judgment.

While there are numerous potential applications for SfM modeling in geotechnical engineering, there are also numerous disadvantages and limitations that must be discussed. Based on the experience of the authors, application of SfM in the field can depend heavily on the available lighting, weather conditions, camera hardware and settings, and the manner in which the photographs are taken. Changes in any of these factors can cause significant inconsistency and unpredictability in the quality of the computer vision model. The presence of shadows or excessively bright objects at a site can result in poor SfM models. Simple point-and-shoot cameras or cameras with fisheye lenses can results in poorer SfM models. Generally, digital SLR cameras with a fixed focus (or "prime") lens tend to generate better models. Significant overlap and redundancy must be present in the image dataset that is being processed. Based on the authors' experience, the best computer vision models are developed when the photographs are taken from the "outside looking in," in which the camera operator moves around the exterior of the site taking photographs while constantly pointing the camera towards the center of the site. If it is desired to capture richer details of specific objects at the site, then additional detailed photographs of those objects can be captured. It is generally ineffective to stand in a stationary point and capture images by rotating. Rather, the camera must be translating for SfM to be effective. Not all objects are well-suited to be modeled with SfM. For example, objects that are very uniform and/or repetitive such as stadium bleachers may result in poor SfM models without additional telemetry metadata being included the digital images. Finally, modeling of very large or long objects (>600 meters) can result in distorted or "curved" SfM models that must be corrected for distortion (Wu 2014).

CONCLUSIONS

Structure from motion computer vision is a virtual modeling technique in which a set of images of a scene, taken from different viewpoints, are used to recover the 3D structure of the scene along with the camera parameters. This study evaluated the accuracy of SfM computer vision models on two full-scale geotechnical field testing projects. Dimensions measured from the SfM models generally were within 3-mm of the actual dimensions. SfM change detection analysis produced relative pile settlements within about 8-mm of true pile settlements.

Most engineers are generally concerned about time and cost, and SfM appears to provide a relatively rapid and cost-effective alternative to other remote sensing methods such as LiDAR. Several potentially useful applications of SfM in geotechnical engineering were discussed including the use of unmanned aerial vehicles. Limitations of the SfM approach were discussed, and general recommendations for producing quality SfM models based on the authors' experience were provided.

ACKNOWLEDGMENTS

Funding for this study was provided in part by the NSF/IUCRC Center for Unmanned Aircraft Systems (C-UAS) Project BYU13-03, and NSF RAPID Grant # 1408892. These sources of funding are gratefully acknowledged. However, the conclusions and opinions presented in this paper do not necessarily reflect those of the C-UAS and the NSF.

REFERENCES

AgiSoft. (2014). Agisoft PhotoScan, Version 1.04. http://www.agisoft.ru/products/photoscan/.

Castillo, C., Pérez, R., James, M.R., Quinton, J.N., Taguas, E.V., and Gómez, J.A. (2012). Comparing the accuracy of several field methods for measuring gully erosion. *Soil Sci. Soc. Am. J.*, 76(4), 1319-1332.

Fathi, H. and Brilakis, I. (2011). Automated sparse 3D point cloud generation of infrastructure using its distinctive visual features. *Adv. Eng. Informatics*, 25(4), 760–770.

Girardeau-Montaut, D. (2012). Cloud Compare, Version 2.4. http://www.danielgm.net/cc/.

Golparvar-Fard, M. (2010). Model-based detection of progress using D 4 AR models generated by daily site photologs and building information models. *Proc.*, ASCE International Conference on Computing in Civil and Building Engineering, No. 1, 2–7.

Golparvar-Fard, M., Pena-Mora, F., and Savarese, S., (2009). D4AR - A 4-Dimensional augmented reality model for automating construction progress monitoring data collection, processing, and communication, *J. Inf. Technol. Constr.*, (14), 129–153.

Harwin, S. and Lucieer, A. (2012). Assessing the accuracy of georeferenced point clouds produced via multi-view stereopsis from unmanned aerial vehicle (UAV) imagery. *Remote Sensing*, 4(12), 1573-1599.

James, M.R. and Robson, S. (2012). Straightforward reconstruction of 3D surfaces and topography with a camera: Accuracy and geoscience application. *J. Geophys. Res.*, 117(F3), F03017.

Jancosek, M. and Pajdla, T. (2011). Multi-view reconstruction preserving weakly-supported surfaces. *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference*, IEEE, 3121-3128.

Koenderink, J.J. and Van Doom, A.J. (1991). Affine Structure from Motion. *Jour. of the Optical Society of America*, A(8.2), 377-385.

Kemeny, J., and Turner, K. (2008). Ground-based LiDAR – Rock slope mapping and assessment. Accessed online Oct. 30, 2013 at <http://www.cflhd.gov>, 3-7.

Microsoft Corporation (2014). Photosynth, Version 2.0110.0317.1042. http://photosynth.net/.

Niethammer, U., James, M.R., Rothmund, S., Travelletti, J., and Joswig, M. (2012). UAV-based remote sensing of the Super-Sauze landslide: Evaluation and results. *Eng. Geol.*, Vol. 128, 2–11.Snavely, N. (2010). Bundler: Structure from Motion (SfM) for Unordered Image Collections, Version 0.4. http://www.cs.cornell.edu/~snavely/bundler/.

SStems Inc. (2014). PhotoModeler, Version 2014.0.2. http://www.photomodeler.com/products/photomodeler.htm.

Wu, C. (2011). VisualSFM - A visual structure from motion system, Version 0.5.24. http://ccwu.me/vsfm/.

Wu, C. (2014). Critical configurations for radial distortion self-calibration. *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference*, IEEE.