International Journal of Rock Mechanics & Mining Sciences 90 (2016) 1–14

Contents lists available at ScienceDirect

International Journal of Rock Mechanics Mining Sciences

journal homepage: www.elsevier.com/locate/ijrmms

Automated rapid mapping of joint orientations with mobile LiDAR

Marc J. Gallant*, Joshua A. Marshall

Mining Systems Laboratory, Queen's University, 25 Union Street, Kingston, Ontario, Canada K7L 3N6

ARTICLE INFO

Keywords:
Joint orientation
Mobile
LiDAR
Laser
Discontinuity
Stereonet

ABSTRACT

This paper introduces three-dimensional axis mapping (3DAM), a new method for joint orientation estimation that rapidly collects data from a mobile platform containing a scannerless LiDAR and an inertial measurement unit (IMU). The mobile platform is moved through the environment either as a handheld device or by mounting it to a remotely operated or robotic vehicle. 3DAM is formulated as a global state estimation problem that estimates the orientation of the mobile platform and the joint set orientations by minimizing the uncertainty introduced by the inherently noisy sensors. This requires a dual-parameterization of both the orientation of joint sets and the mobile platform to permit the use of state estimation techniques. 3DAM was field tested in three separate locations and is directly compared to hand measurements and stationary LiDAR. In all experiments, it is shown that 3DAM produces stereonets comparable to other methods, yet does so with lower-cost hardware and significantly reduced effort.

1. Introduction

LiDAR has recently been recognized as a valuable tool for rock mass characterization.[1] In particular, there has been considerable interest in using LiDAR for joint orientation estimation.[2–12] Generally, these methods use a stationary tripod-mounted LiDAR to capture one or more point clouds of the exposed rock face. The orientation of the sensor is measured beforehand such that the collected data can be transformed to the global coordinate frame. Using the geometry of the captured point cloud, the orientations of planar surfaces are extracted and plotted on a stereonet.

Joint orientation estimation using LiDAR has several advantages compared to manual techniques (i.e., using a compass and inclinometer). Safety is improved by limiting human exposure to potentially unstable rock faces, it allows greater access to hard-to-reach locations, it may be faster and less labor intensive, and the opportunity for introducing erroneous measurements due to procedural difficulties, human bias, and human error is significantly reduced. Despite these advantages, there are barriers preventing widespread adoption of stationary LiDAR. Its high cost can be prohibitive, it is often physically too large and heavy for some remote deployments, processing the resulting point clouds often requires manual intervention to remove outliers, and several measurements of the same rock face from different points of view are usually required to avoid occlusions. This last drawback can make the data collection process very time-consuming (i.e., moving, reorienting, and taking new measurements with the sensor can take tens of minutes) and is particularly challenging in enclosed spaces (e.g., underground mines).

This paper introduces a new algorithm for joint orientation estimation called *three-dimensional axis mapping* (3DAM) that retains the advantages of using LiDAR measurements for joint orientation estimation, while addressing many of its disadvantages. 3DAM achieves this by using a lightweight, inexpensive *mobile* platform containing a low-cost, compact three-dimensional LiDAR and an inertial measurement unit (IMU). Joint orientations are derived from the collected data by probabilistically determining the orientation of the mobile platform along its trajectory. This is calculated by combining high frequency measurements of the angular velocity of the mobile platform, the direction of gravity, the direction of the Earth's magnetic field, and by repeatedly measuring the orientations of the joint sets themselves.

Many of the disadvantages of using stationary LiDAR are avoided with 3DAM. By capturing a larger number of lower resolution point clouds, the use of a low cost and mobile LiDAR is permitted. Occlusions are easily and quickly eliminated by simply maneuvering around the rock face. The techniques used to process the resulting 3DAM point clouds are fast, fully automated, and require no manual pruning of outliers. The time and effort required to collect data is significantly reduced compared to both manual techniques and stationary LiDAR. Finally, the mobile platform can take any form capable of carrying the sensors, from a handheld configuration to remote operation with a mobile robot.

The focus of this paper is on automatic estimation of joint orientations, with visualization as stereonets. The problems of estimat-

* Corresponding author.
E-mail addresses: m.gallant@queensu.ca (M.J. Gallant), joshua.marshall@queensu.ca (J.A. Marshall).

http://dx.doi.org/10.1016/j.ijrmms.2016.09.014
Received 4 December 2015; Received in revised form 12 May 2016; Accepted 20 September 2016

ing other characteristics such as, for example, joint spacing, roughness, area/volume estimates, and change detection lie beyond the scope of this work. For these problems, the specific approach developed here may not provide similar advantages.

1.1. Related work

Research about joint orientation estimation using LiDAR has focused on different strategies for extracting the orientation of planar surfaces from point clouds. In general, these strategies can be categorized as either *top-down* or *bottom-up* approaches. The top-down approach (sometimes called *surface reconstruction*) uses a mesh generated from the point cloud (usually by interpolating between points) to approximate the surface of the measured rock face, and calculates the normal vector at each segment in the mesh. Many researchers depend on commercial software (e.g., PolyWorks[13]) to perform surface reconstruction. Slob et al.[14] evaluated this approach and compared it to manual techniques from a cost-benefit perspective. Other researchers have studied optimizing the point cloud processing procedure,[5] correcting for biases in data collection,[7] and have analyzed the sensitivity to point cloud resolution.[6,10] The top-down approach often requires manual intervention to prevent unwanted points from being included in the surface reconstruction [10] (e.g., vegetation, outliers).

The bottom-up approach attempts to find subsets of points on planar surfaces and uses least-squares techniques (e.g., principal component analysis) to fit planes to the data. This requires finding subsets of points in the point cloud measuring planar surfaces. A plane is then fit to the subset of points and its normal vector is calculated. Various methods have been used to find the appropriate planar subsets. For example, random sample consensus (RANSAC) has been shown to be an effective way to find subsets belonging to planar surfaces by iteratively comparing subsets to the model of a plane.[2,4,11] Gigli et al.[9] divide the point cloud into a cubic grid of various resolutions to evaluate the planarity of points in each grid cell. Olariu et al.[3] apply a k-means clustering algorithm at different resolutions and measure the planarity of the clusters.

A thorough comparison by Slob[15] of the two techniques concluded that the bottom-up approach is preferable because it is easier to automate, retains the original point cloud, and is better at dealing with outliers. 3DAM uses a modified form of the bottom-up approach for orientation extraction. Instead of fitting planes to discrete planar surfaces, it estimates the planarity at each point in the point cloud and then clusters similar points across the entire point cloud.

There has been limited research on the use of a mobile platform for joint orientation estimation. De Agostino et al.[16] mounted a GPS, IMU, LiDAR, and camera to a truck, but only captured point clouds while the LiDAR was stationary. As a result, this system was very similar to the established techniques, with the exception that the sensors were more easily moved from one scanning location to the next. A mobile LiDAR system constructed by Terrapoint Canada[17] (now Ambercore) and employed by Lato et al.[18] captured point clouds while in motion. This system combines multiple LiDAR sensors, a high-end IMU, and a differential GPS system to track the movement of a truck mounted on railway tracks. The system constructs a massive point cloud in the global coordinate frame, after which a top-down approach was applied to manually selected portions of the point cloud.

Rather than constructing a large, high-resolution point cloud and applying standard techniques to extract the joint orientations, 3DAM is built from the ground up with joint orientation estimation in mind. It is designed to be relatively inexpensive, easily carried and operated by a single person, and have no dependence on GPS. By meeting these constraints, it is deployable in a large variety of environments (i.e., anywhere where a person or remotely operated vehicle can reach, including underground) and is both economically and physically viable as a drop-in replacement for a compass and inclinometer. Although the expensive stationary LiDAR devices (i.e., tens of thousands of dollars[1]) used in the literature discussed above produce large, high-resolution point clouds—which may be useful for solving problems beyond only joint orientation estimation—they must remain stationary during the data acquisition process that can take tens of minutes per scan. On the other hand, 3DAM uses a low-cost (i.e., hundreds of dollars), lower resolution and mobile LiDAR. The need for high resolution is mitigated by using a LiDAR with a much higher frequency, which captures multiple complete point clouds *per second*. However, one tradeoff is that this type of low-cost and mobile LiDAR has a shorter range (usually less than 5–10 m) compared with stationary LiDAR (can be tens of meters), making it applicable only in low-range scenarios.

1.2. About this paper

This paper begins with mathematical preliminaries used in 3DAM (Section 2). In particular, background information about the parameterization of joint orientations and state estimation is presented. Next, the 3DAM algorithm is described by decomposing it into smaller sub-algorithms (Section 3), which begins with an overview of the algorithm as a whole (Section 3.1). Information about the field tests is presented (Section 4), followed by a discussion of the experimental results obtained from tests at three different sites (Section 5).

2. Mathematical preliminaries

The 3DAM algorithm represents the orientations of planar surfaces in an environment as *axes* (Section 2.1) and maximizes the likelihood of all measurements using a *state estimation* algorithm (Section 2.2). Background material about these two topics is presented here.

2.1. Axes

An *axis* is an unordered pair of opposing directions. A phenomenon that is well-represented by an axis is the orientation of a plane in $\mathbb{R}^3$. In this example, the unordered pair of opposing directions is the normal vectors on the two sides of the plane. When the planar surfaces are joints, the two degrees of freedom of an axis are often parameterized as dip and dip direction. Although this parameterization is intuitive and useful for visualization purposes, it is not a suitable choice for use in state estimation. This is because dip and dip direction do not form a vector space near the origin. For example, a small perturbation of a near-zero dip can significantly change the dip direction. Furthermore, the dip direction is undefined when dip is zero. Instead, two other parameterizations of axes are used to avoid these issues; namely, unit axes and axis vectors.

A *unit axis* is an over-parameterization (three parameters, one constraint) of an axis that is free of singularities (i.e., a small perturbation of an axis will always result in a small perturbation of its unit axis parameterization). A unit axis $\mathbf{n}$ is the 3×1 column

$$\mathbf{n} := \begin{bmatrix} \lambda \\ \boldsymbol{\kappa} \end{bmatrix}, \tag{1}$$

where $\lambda \in \mathbb{R}$ and $\boldsymbol{\kappa} \in \mathbb{R}^2$ are the scalar and vector parts of the unit axis, respectively. Unit axes are constrained to the unit sphere S^2; that is, $\lambda^2 + \boldsymbol{\kappa}^T\boldsymbol{\kappa} = 1$. Additionally, $\mathbf{n} \equiv -\mathbf{n}$ because they both represent the same axis. An illustration of the scalar and vector components of a unit axis is shown in Fig. 1a. Because they are free of singularities, unit axes are a good choice for a *global* parameterization of axes (i.e., every axis is well-represented by a unit axis).

An *axis vector* is a minimal parameterization (two parameters) of an axis that is well-suited for state estimation problems (e.g., near the origin, axis vectors behave like real vectors). An axis vector $\boldsymbol{\phi}$ is the 2×1 column