Robotics and Autonomous Systems 82 (2016) 35–45

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

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

The LiDAR compass: Extremely lightweight heading estimation with axis maps

Marc J. Gallant *, Joshua A. Marshall

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

HIGHLIGHTS

- A two-dimensional heading estimation algorithm called the LiDAR compass is developed.
- LiDAR measurements are associated with the orientations of flat surfaces.
- In many environments, the growth of heading errors is bounded.
- The LiDAR compass is lightweight because it uses mostly linear operations on scalars.
- When paired with odometry, a good initial estimate for graph-based SLAM is generated.

ARTICLE INFO

Article history:
Received 30 March 2015
Received in revised form
5 February 2016
Accepted 12 April 2016
Available online 4 May 2016

Keywords:
Heading estimation
LiDAR
Compass
Localization

ABSTRACT

This paper introduces the *LiDAR compass*, a bounded and extremely lightweight heading estimation technique that combines a two-dimensional laser scanner and *axis maps*, which represent the orientations of flat surfaces in the environment. Although suitable for a variety of indoor and outdoor environments, the LiDAR compass is especially useful for embedded and real-time applications requiring low computational overhead. For example, when combined with a sensor that can measure translation (e.g., wheel encoders) the LiDAR compass can be used to yield accurate, lightweight, and very easily implementable localization that requires no prior mapping phase. The utility of using the LiDAR compass as part of a localization algorithm was tested on a widely-available open-source data set, an indoor environment, and a larger-scale outdoor environment. In all cases, it was shown that the growth in heading error was bounded, which significantly reduced the position error to less than 1% of the distance travelled.

1. Introduction

Localization of mobile robots or other ground vehicles is an active area of research that has important applications in mapping, planning, and control. In the absence of an absolute positioning system (e.g., GPS), localization is traditionally performed by measuring the internal state of the vehicle with interoceptive sensors (e.g., accelerometers, gyroscopes, wheel encoders), and/or measuring the local environment surrounding the vehicle with exteroceptive sensors (e.g., cameras, LiDAR). Accurately estimating the robot's heading is particularly important when the localization algorithm involves dead reckoning. In this common scenario, the positional components of the motion model are often tightly coupled with the heading component, causing inaccuracies in heading estimation to be quickly propagated into substantial position errors. As a result, accurate heading estimation can be essential for localization. Although sensors exist that can directly or indirectly be used for heading estimation (e.g., a compass or gyroscope), environmental limitations and/or dead reckoning errors can render these estimates unreliable for many applications. To this end, this paper introduces the *LiDAR compass* (LC), which transforms data from a horizontally-oriented 2D scanning LiDAR into an absolute heading estimate that can be used in a variety of indoor and outdoor environments.

The LC uses *axis maps* (AMs), which are a minimal representation of the orientations of surfaces in the environment. To emphasize the fact that the orientation of a flat two-dimensional surface (i.e., a line) is invariant to rotations of $\pm\pi$, it is called the *axis* of the surface. This naming convention is consistent with other types of axial data in directional statistics [1]. Lines are extracted from 2D LiDAR scans and compared to both an *a priori* and local AMs, which provide information about the absolute or relative heading of the robot. As a result, it is assumed that the operating environment

* 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.robot.2016.04.005
0921-8890/

contains some approximately straight surfaces. This assumption is met in a great many common scenarios, including both indoor (e.g., walls, furniture) and outdoor (e.g., buildings, cars) environments. In these environments, the LC is essentially a virtual heading sensor that – when combined with a means to measure translation – can be used to aid localization. Using the LC in this way is analogous to augmenting wheel odometry with a compass and gyroscope. Much like how a compass provides an absolute heading reference by measuring the direction of the Earth's magnetic field, the LC uses the dominant axes of the surfaces in the environment as the absolute reference. And similar to how a gyroscope provides a dead reckoning heading estimate by integrating angular velocity, the LC also provides relative heading estimates by tracking the axes of local surfaces not in the *a priori* AM. The result is an easy-to-implement, very lightweight, virtual heading sensor that (unlike a gyroscope) provides an absolute heading reference, and (unlike a compass) can be used in any environment from which an AM can be derived.

There are several applications where the advantages of using an LC would be especially useful. For example, the state of the art in autonomous mapping is graph-based simultaneous localization and mapping (SLAM), which is formulated as a nonlinear least-squares problem that maximizes the likelihood of sensor measurements. This formulation requires an initial estimate of the robot trajectory, and the quality of this initial estimate can greatly affect the accuracy of the resulting map [2]. The LC can be combined with odometry to provide an easily-implemented solution, while not requiring any additional sensors. For instance, Fig. 1 illustrates the trajectories estimated by raw odometry, an LC with odometry, and SLAM [3] for the MIT Killian Court data set [4]. The detailed results using this data set are provided in Section 5.1.

Another application for the LC is in small, consumer robots, where the embedded computing environment would benefit from a lightweight, real-time localization algorithm (e.g., the Neato Robotics robot vacuums, which are equipped with a 2D LiDAR). These products tend to operate in semi-structured environments from which an AM could be easily derived. Finally, vehicles equipped with an integrated GPS and inertial navigation system (GPS/INS) in urban environments rely on dead-reckoning from accelerometers and gyroscopes when GPS signals are unavailable (e.g., in tunnels or when surrounded by buildings). Augmenting such a system with an LC could easily and drastically improve localization by using the obstructions themselves to populate the AM.

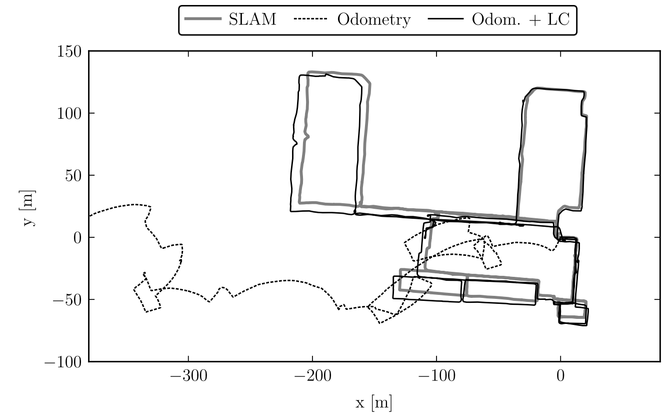

Fig. 1. The trajectory of the MIT Killian Court data set as estimated by odometry, odometry with a LiDAR compass (LC), and the full SLAM solution [3]. Much of the odometry-estimated path is outside of the scale of this figure. The LC estimate provides a good initial guess for graph-based SLAM. A major source of error of the LC trajectory is from translational error in the odometry estimate, which may be partially mitigated with proper calibration.

1.1. Related work

The LC requires re-occurring, line-extractable surfaces in the environment. Although its map consists only of the axes of lines extracted from these environments, line-based map representations for localization and mapping have considerable heritage in mobile robotics research. Before relatively low-cost LiDAR became available, sonar was used to observe *a priori* line-based maps of the environment [5]. Here, the idea of improving localization by incorporating geometric constraints in the environment was exploited. One of the first implementations using LiDAR provided online localization given an *a priori* line-based map [6], where it was noted that many indoor environments are suitable for this type of map representation. An early implementation that actually constructed line-based maps [7] predates modern SLAM and actually decouples mapping and localization. However, even with a rudimentary 2D LiDAR, line-based maps of less structured environments such as underground mines were shown to be effective representations.

More recent efforts have used line-based maps in SLAM implementations by employing the position and axis of line segments to update the pose estimate of the robot. One EKF SLAM implementation [8] demonstrated the accuracy of using line segments in the SLAM state, whose compactness also reduces the burden of the computational complexity of SLAM. However, as line segments have four degrees of freedom, effective data association among line segments requires several correspondence tests. This issue is non-existent in AMs because AM entries have only a single degree of freedom.

The popular line-based Orthogonal SLAM [9] takes advantage of an orthogonality assumption of the surfaces in the environment (e.g., the perpendicular walls common in most indoor areas). This approach was later extended to create a lightweight Rao–Blackwellized particle filter Orthogonal SLAM [10]. Here, only orthogonal lines extracted from the environment are used to update the SLAM state. By fixing the possible axes of lines to an absolute reference (and discarding lines that do not meet this criterion), remarkable mapping and localization accuracy is achieved by limiting the growth of heading errors. Recent work [11] has explored automatically identifying additional types of structure (e.g., point-to-point distances, circles) and incorporating these constraints into the graph-based optimization of the map. The LC presented in this paper shares some of the benefits of these structure-sensitive SLAM algorithms with its *a priori* axis map. Unlike the aforementioned approaches, the LC is intended to use the structure of the environment for real-time, lightweight, and easily implemented heading estimation. It is effective in environments whose surfaces are not orthogonal or not even in the *a priori* map.

A different approach of heading estimation called a *visual compass* (VC) is currently an active research topic in mobile robot localization. Although many implementations of VCs exist, most implementations use a camera to track changes or features in the image frame to infer rotational information. VCs differ from visual odometry because they usually discard all positional information in the sensor data, not unlike the LC. A common approach is to unwrap sequential omnidirectional images and observe how simple extracted features appear to be displaced as the robot rotates [12,13]. An overview of this method concluded that its dead reckoning heading estimation performed similarly or better than inertial sensors in appropriate environments [12]. Other forms of VCs do not necessarily require omnidirectional cameras and instead track the motion of specific vanishing points at far distances [14,15], which work well in large, open environments. Finally, a recent VC implementation addresses the restrictions of many other VCs (e.g., computation, prior environmental