Robotics and Computer-Integrated Manufacturing 42 (2016) 1–16

Contents lists available at ScienceDirect

Robotics and Computer-Integrated Manufacturing

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

High-accuracy vehicle localization for autonomous warehousing

Goran Vasiljević [a], Damjan Miklić [a,*], Ivica Draganjac [a], Zdenko Kovačić [a], Paolo Lista [b]

[a] LARICS Laboratory, Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia
[b] Euroimpianti S.r.l., Schio (VI), Italy

ARTICLE INFO

Article history:
Received 15 September 2015
Received in revised form
23 March 2016
Accepted 3 May 2016
Available online 14 May 2016

Keywords:
High-accuracy localization
Autonomous warehousing
Autonomous ground vehicle

ABSTRACT

The research presented in this paper aims to bridge the gap between the latest scientific advances in autonomous vehicle localization and the industrial state of the art in autonomous warehousing. Notwithstanding great scientific progress in the past decades, industrial autonomous warehousing systems still rely on external infrastructure for obtaining their precise location. This approach increases warehouse installation costs and decreases system reliability, as it is sensitive to measurement outliers and the external localization infrastructure can get dirty or damaged. Several approaches, well studied in scientific literature, are capable of determining vehicle position based only on information provided by on board sensors, most commonly wheel encoders and laser scanners. However, scientific results published to date either do not provide sufficient accuracy for industrial applications, or have not been extensively tested in realistic, industrial-like operating conditions. In this paper, we combine several well established algorithms into a high-precision localization pipeline, capable of computing the pose of an autonomous forklift to sub-centimeter precision. The algorithms use only odometry information from wheel encoders and range readings from an on board laser scanner. The effectiveness of the proposed solution is evaluated by an extensive experiment that lasted for several days, and was performed in a realistic industrial-like environment.

1. Introduction

Notwithstanding the recent very successful examples of large scale use of autonomous mobile delivery vehicles such as Kiva Systems [1], today's manufacturing and logistics facilities are still largely dependent on manually operated vehicles [2]. Although numerous commercial solutions using autonomous ground vehicles (AGVs) do exist, e.g., automated material handling vehicles by Swisslog[1] and Euroimpianti[2] or Komatsu's autonomous haul systems for the mining industry[3], significant improvements can be made in terms of their level of autonomy and deployment cost. In particular, the industrial state of the art for vehicle localization requires additional infrastructure, usually in the form of reflective markers or electromagnetic guides, for accurately determining vehicle pose [3]. This approach suffers from numerous disadvantages, as the markers have a high installation cost, requiring many man-hours of work by qualified personnel, they are sensitive to false positive readings and to changes in the environment which may obstruct the field of view of the vehicle, and they can get damaged or dirty in harsh industrial environments. Relevant contemporary literature, such as [4–6] unanimously identifies the ability of AGVs to self-localize without additional infrastructure as one of the key technologies that will increase their performance and flexibility, thus enabling their widespread use.

Because self-localization is a basic prerequisite for autonomous vehicle operation, this subject has been receiving significant scientific attention from the earliest beginnings of mobile robotics. Different sensors have been used for this purpose, the most common ones being wheel encoders, contact switches, sonar arrays, 2D and 3D laser scanners, mono, stereo and RGBD cameras. A nice overview of relevant references for the different approaches can be found in [7]. For localization in the plane, i.e. on flat terrain, 2D laser scanners have been shown to provide the best accuracy, robustness and speed [8]. Additionally, the vast majority of AGVs deployed today is already equipped with laser scanners for safety and localization (using artificial landmarks) purposes, so we are focusing on this type of sensor throughout the rest of the paper. Advances in computational power at the turn of the century have

---

* Corresponding author.
E-mail address: damjan.miklic@fer.hr (D. Miklić).

[1] http://www.swisslog.com/en/Solutions/HCS/Material-Handling-Automation

[2] http://www.skilledrobots.com/products/agv

[3] http://www.komatsu.com.au/AboutKomatsu/Technology/Pages/AHS.aspx

http://dx.doi.org/10.1016/j.rcim.2016.05.001
0736-5845/

made new classes of probabilistic methods computationally feasible. Since then, the AMCL[4] method [9] has established itself as the de facto standard in academic research. The method is based on a particle filter which fuses odometry information provided by wheel encoders with laser scanner range readings to provide robust localization with reported accuracy between 0.05 m and 0.1 m [10,11]. Because it is a multi-hypothesis method, it is able to deal with situations when the robot is temporarily “lost”. These features make the method suitable for most academic research work in mobile robotics. Its effectiveness has led researchers to consider robot localization in the plane to be a “solved problem” and move on to other research topics, such as localization in 3D space. However, the accuracy required by material handling applications is typically within 0.01 m and 0.5°, so the industry has continued relying on additional infrastructure to ensure the required accuracy.

In the past five years, the prospect of widespread industrial adoption of autonomous vehicles, with enhanced performance and flexibility and reduced deployment costs, has spurred vigorous research activity related to precise localization without artificial landmarks. A popular approach for obtaining high precision while maintaining robustness is to combine AMCL with scan matching. Estimates obtained by AMCL are refined by matching the laser readings to map features. This approach is used very successfully in [7]. The authors performed extensive experiments with a holonomic vehicle in a laboratory environment, using a vision-based tracking system for evaluating the results with millimetre accuracy. In a static environment, they were able to achieve localization and positioning errors at taught-in locations well below 0.005 m and 0.2°. Slightly higher, but still sub-centimetre errors were reported in the presence of dynamic obstacles (people walking by). In a similar setup, tracking accuracy of taught-in trajectories was examined [12]. By manually teaching-in the trajectories, the authors avoid using a global map. Point-to-line scan matching [13] is used for trajectory tracking in the robot's local coordinate system. Accuracy was evaluated by measuring the minimal distance to the reference trajectory and milimeter accuracy is reported. It should be noted that the experimental environment was small and completely static. The work presented in [14] focuses on seamless transition between marker based, map based and pure SLAM (localization without an a priori map) scenarios. The reported preliminary evaluation in a large scale warehouse, on a single trajectory, had errors within 0.05 m and 0.5° most of the time. In their later work, the same authors examine the accuracy of 2D localization as a prerequisite for 3D mapping [15], in extensive experiments in a large scale warehouse. They do not use odometry information from wheel encoders. Instead they use Point-to-Line Iterative Closest Point (P-L-ICP) [13] to simulate odometry. They do not use the localization estimate as feedback for the position controller (they rely on the artificial landmarks for control) and report a mean positioning error of 0.052 m and 0.012°. Another very promising approach is based on the Normal distributions transform (NDT) introduced by [16]. The NDT is a piecewise continuous representation, which represents the space as a set of normal distributions. It enables a more compact and more accurate representation of the environment, compared to grid-based representations. Authors in [17] formulate Monte Carlo localization (MCL) using the Normal distributions transform (NDT) as underlying representation for map and data. They evaluate their approach using offline data sets, in closed loop with a smaller AGV in a laboratory environment and in a real warehouse with an industrial AGV in open loop (localization estimate is not used as control feedback for the vehicle). They achieve localization accuracy of 0.014 m and 0.074° in laboratory conditions and below 0.03 m in the industrial setting. The NDT-MCL algorithm is extended in [18] to dual-timescales, i.e., a dynamically updated short-term map is used for localization in addition to the a priori provided static map. This approach improves localization performance in highly dynamic environments, reducing localization errors below 0.02 m in a laboratory setting. In our previous work [19], we evaluated an approach which fused readings from several laser scanners mounted at different heights. The average localization error was below 0.06 m, while the maximum error was kept within 0.1 m at all times during an experiment in an industrial setting, which does not quite satisfy industrial requirements.

The approach pursued in this paper combines AMCL, scan matching and Discrete Fourier Transform (DFT)-based pose estimate refinement into one algorithm stack for high-precision localization in industrial indoor environments. We describe all components of the localization algorithm and provide extensive experimental results, which confirm the accuracy, robustness and reliability of our approach. The experiments have been performed in an industrial warehouse setting, with a full sized autonomous forklift. The localization module is working in closed loop with the vehicle control module, so any significant localization error would cause a failure in path execution. During three days and 19 hours of total travel time, the vehicle has logged over eight kilometers, relying only on map information and its sensor readings, without a single failure or operator intervention. The results are evaluated using a methodology similar to [7]. Because vehicle positioning is most critical at docking stations, when picking up and delivering pallets, we defined 6 docking stations in the warehouse layout. The vehicle was repeatedly visiting the docking stations and these positions were used for evaluating localization accuracy. Making a fair comparison with results from other authors is difficult, mainly because of a lack of standardization. Some attempts at standardization are beeing made [20,21], however, these are currently not applicable to industrial systems. In state of the art research mentioned above, the vehicles, experimental environment and evaluation methods vary vastly. Experimental platforms range from holonomic laboratory robots to non holonomic industrial forklifts weighting several tons. Some experiments are performed in laboratory conditions, while others are performed in actual in-production warehouses. Different evaluation methods include absolute errors with respect to a reference localization system, relative errors computed with respect to taught-in poses, static errors, dynamic errors or Absolute Trajectory Errors (ATE) [20]. Notwithstanding the lack of standardized methods for comparison, to the best of our knowledge, no localization algorithm to date has been evaluated in comparably realistic industrial conditions which provides better localization accuracy at taught-in docking positions. Therefore, the main contribution of this paper is a rigorously validated AGV localization algorithm, which satisfies demanding industrial requirements and provides an improvement over the state of the art. Furthermore, in addition to the localization algorithm, we also describe the path planning and tracking algorithms that have been implemented on the vehicle and used for the experiment. Together with the localization module, they constitute a complete, experimentally verified positioning solution for autonomous warehousing.

The paper is organized as follows. The focal point of the paper is Section 2, where we provide a detailed description of our localization algorithm stack. In Section 3 we describe the industrial platform and the environment in which we performed our experiments. The positioning algorithm used for experimental validation is described in Section 4. The results of extensive experimental validation are described in Section 5. Concluding remarks and an outline of our future work plans are provided in Section 6.

---

[4] Adaptive Monte Carlo localization