Robotics and Autonomous Systems 70 (2015) 63–82

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

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

Incremental light bundle adjustment for structure from motion and robotics

Vadim Indelman [a,*], Richard Roberts [b], Frank Dellaert [b]

[a] Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel

[b] Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30332, USA

HIGHLIGHTS

- A novel approach for structureless incremental bundle adjustment is developed.
- Observed 3D points are algebraically eliminated.
- Optimization employs incremental smoothing for efficiently recovering MAP estimate.
- 3D points can be reconstructed at any time, but only if required.
- Probabilistic analysis and performance study are conducted.
- Application both to structure from motion and robotics problems.

ARTICLE INFO

Article history:
Received 15 May 2014
Received in revised form
15 December 2014
Accepted 17 March 2015
Available online 27 March 2015

Keywords:
Bundle adjustment
SLAM
Vision-aided navigation

ABSTRACT

Bundle adjustment (BA) is essential in many robotics and structure-from-motion applications. In robotics, often a bundle adjustment solution is desired to be available incrementally as new poses and 3D points are observed. Similarly in batch structure from motion, cameras are typically added incrementally to allow good initializations. Current incremental BA methods quickly become computationally expensive as more camera poses and 3D points are added into the optimization. In this paper we introduce incremental light bundle adjustment (iLBA), an efficient optimization framework that substantially reduces computational complexity compared to incremental bundle adjustment. First, the number of variables in the optimization is reduced by algebraic elimination of observed 3D points, leading to a structureless BA. The resulting cost function is formulated in terms of three-view constraints instead of re-projection errors and only the camera poses are optimized. Second, the optimization problem is represented using graphical models and incremental inference is applied, updating the solution using adaptive partial calculations each time a new camera is incorporated into the optimization. Typically, only a small fraction of the camera poses are recalculated in each optimization step. The 3D points, although not explicitly optimized, can be reconstructed based on the optimized camera poses at any time. We study probabilistic and computational aspects of iLBA and compare its accuracy against incremental BA and another recent structureless method using real–imagery and synthetic datasets. Results indicate iLBA is 2–10 times faster than incremental BA, depending on number of image observations per frame.

1. Introduction

Bundle adjustment (BA) has been at the focus of many research efforts for the past few decades. While its origin is in the photogrammetry community, it is an essential component in many applications in computer vision and robotics, where it is also known as structure from motion (SfM) and full simultaneous localization and mapping (SLAM). A thorough review of different aspects in BA can be found in [1].

Bundle adjustment can be considered as a large optimization problem, with the optimized variables being the camera poses and the observed structure (3D points/landmarks) and where one aims to minimize the difference between the actual and the predicted image observations. Computational complexity is scenario-dependent and is a function of several factors, including the number of images, observed 3D points, and the actual image

* Corresponding author.
E-mail address: vadim.indelman@technion.ac.il (V. Indelman).

http://dx.doi.org/10.1016/j.robot.2015.03.009

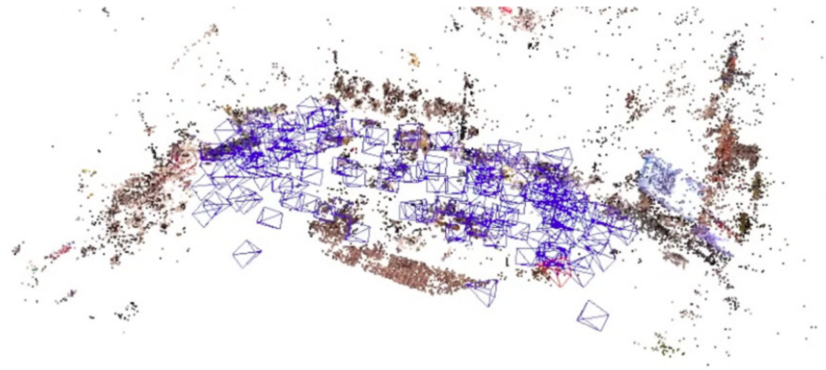

(a) *Cubicle* dataset: 148 cameras, 31,910 3D points and 164,358 image observations.

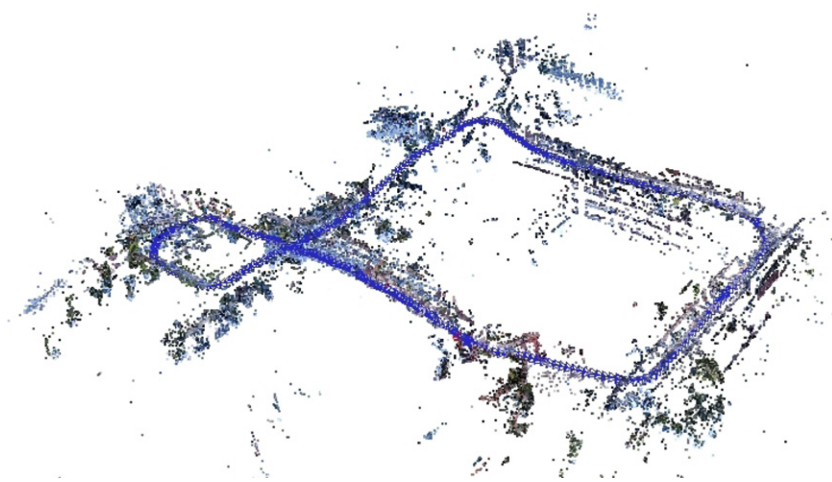

(b) *Outdoor* dataset: 308 cameras, 74,070 3D points and 316,696 image observations.

Fig. 1. Optimized camera poses using iLBA followed by (sparse) structure reconstruction. Computational time of iLBA is 8–10 times faster compared to incremental BA.

observations. Research in recent years has been focused on developing efficient and fast approaches for BA optimization, which is required to process a large amount of information in a reasonable time. Proposed approaches include exploiting the typical sparseness of the problem, performing parallelization and distribution of the involved calculations, and optimizing a reduced system that approximates the full system.

This paper introduces incremental light bundle adjustment (iLBA), an efficient optimization method that is applicable both to SfM and SLAM. The method incorporates two key components to reduce computational complexity: structureless BA and incremental smoothing.

The idea in the recently-developed structureless BA methods [2–4] is to improve computational complexity by reducing the number of variables involved in the optimization. Specifically, in structureless BA only the camera poses are optimized, while the 3D points are algebraically eliminated using multi-view constraints. The optimization minimizes the errors in satisfying these constraints, instead of optimizing the re-projection errors as in standard BA. iLBA utilizes three-view constraints [5], which in contrast to using only epipolar constraints in structureless BA [2], allow consistent motion estimates even when the camera centers are co-linear, a situation common in mobile robotics. Alternatively, one can apply trifocal constraints [6,7] within the same framework. If required, all or some of the 3D points can be calculated using standard structure reconstruction techniques using the optimized camera poses. Fig. 1 shows an example of iLBA-optimized camera poses, which are then used for sparse 3D reconstruction in two real–imagery datasets considered in this paper.

The second component in iLBA is incremental smoothing, an efficient incremental optimization that uses a recently-developed technique [8] to update the solution using adaptive partial calculations each time a new camera (image) is incorporated into the optimization. Being able to incorporate new information and efficiently perform the optimization is important in many applications, such as robot vision, where the images are not given ahead of time (e.g. gradually captured by a single or group of mobile robots) and the maximum a posteriori (MAP) estimate is required each time a new image is received. Furthermore, since a good initialization is essential not to be trapped in a local minima, incremental optimization is also used in batch scenarios where all the imagery data is available. Typically, incremental smoothing involves recalculating only a small number of camera poses as opposed to always optimizing all camera poses in previous incremental SfM [9] and structureless BA methods [10,2,3]. To the best of our knowledge, incremental smoothing has not been suggested to structureless BA thus far.

We demonstrate, using different real–imagery and synthetic datasets, that these two components result in a significant reduction in computational complexity while exhibiting high accuracy for reasonable image noise levels, compared to standard BA. Our implementation of iLBA is publicly available.[1]

In this paper we also present a probabilistic analysis of iLBA, analyzing how well the probability distribution corresponding to

[1] Code is available on the website of the first author.