

Benchmarking and Comparing Popular Visual SLAM Algorithms

Amey Kasar,
Dept. of Electronics & Telecom. Eng.,
Pune Institute of Computer Technology, Pune, India
kasaramey16@gmail.com

Abstract—This paper contains the performance analysis and benchmarking of two popular visual SLAM Algorithms: RGB-D SLAM and RTABMap. The dataset used for the analysis is the TUM RGBD Dataset from the Computer Vision Group at TUM. The dataset selected has a large set of image sequences from a Microsoft Kinect RGB-D sensor with highly accurate and time synchronized ground truth poses from a motion capture system. The test sequences selected depict a variety of problems and camera motions faced by SLAM algorithms for the purpose of testing the robustness of SLAM algorithms in different situations. The evaluation metrics used for the comparison are Absolute Trajectory Error (ATE) and Relative Pose Error (RPE). The analysis involves comparing the Root Mean Square Error (RMSE) of the two metrics and the processing time for each algorithm. This paper serves as an important aid in the selection of SLAM algorithm for different scenes and camera motions. The analysis helps to realize the limitations of both SLAM methods. This paper also points out some underlying flaws in the used evaluation metrics

Keywords—Simultaneous Localization And Mapping, Benchmark, RGBD SLAM, RTABMap

I.

INTRODUCTION

Robotic applications require modeling environment for various tasks, guidance, search and rescue, etc. A precondition for an autonomous robot is to obtain an accurate model of its environment. A major problem in obtaining this model is the problem of pose uncertainty. The mobile robot mapping problem under pose uncertainty is often referred to as the simultaneous localization and mapping (SLAM) or Concurrent Mapping and Localization (CML) problem [1]–[3].

SLAM is a complex problem because a robot needs a homogeneous map in order to localize itself. To obtain a map, the robot requires a good estimate of its location. The limited range and limited field of view of ranging sensor adds to the problems. The mutual dependency among the pose and the map estimates makes the SLAM problem hard. Even a small error in the map could prevent a robot from working in the environment. Hence it is important to tackle the SLAM problem.

Many different techniques to tackle the SLAM problem have been presented. There are different approaches to the problem like Extended Kalman Filter SLAM (EKF), Sparse

Extended Information Filter (SEIF), Extended Information Form (EIF), FastSLAM, GraphSLAM. There are also some proposed evaluation metrics [4]–[7] for comparing the results of different SLAM algorithms. People often use visual inspection to compare maps or overlays with blueprints of buildings for grid based estimation techniques. This kind of traditional evaluation becomes more and more difficult as new SLAM algorithms show increasing capabilities. Meaningful comparisons between different SLAM algorithms require some common performance metrics. The metrics should enable the user to compare the outcome of different mapping approaches when applying them on the same dataset. This research uses two such common evaluation metrics on the results obtained from popular SLAM algorithms.

This paper utilizes the metrics for evaluation proposed in [4] to evaluate the performance of two popular SLAM methods: RGB-D SLAM [8] and RTABMap [9]. The evaluation metrics used for the comparison are Absolute Trajectory Error (ATE) and Relative Pose Error (RPE).

The TUM RGB-D dataset [10] is applied to both the algorithms and the resulting trajectory estimate is compared to the ground truth by evaluating Absolute Trajectory Error (ATE), Relative Pose Error (RPE) and the time taken to process the sequence of images. The sequences are intentionally selected to depict the difficulties encountered by SLAM algorithms when operating in real world. The results of the evaluation are analyzed to determine the ideal algorithm for different situations.

II.

SYSTEM CONFIGURATION

The system used for testing has the following configuration

System Configuration	
CPU	Intel Core i7-7700HQ 2.80GHz x 8
Memory	8GB
Operating System	Ubuntu 14.04 LTS
OS Type	64-bit
GPU	GeForce GTX 1050
Cuda	Ti/PCIe/SSE2
SURF	8.0 Enabled


SIFT	Enabled
	

Fig. 1. ActiveMedia Pioneer 3 robot

III.

ALGORITHMS SELECTED FOR EVALUATION

A. RGB-D SLAM

RGB-D SLAM [8] works well with a hand-held Kinect-style depth sensor. It uses visual features such as SURF or SIFT to match pairs of acquired images, and uses RANSAC to robustly estimate the 3D transformation between them. To achieve online processing, the current image is matched only versus a subset of the previous images. Subsequently, it constructs a graph whose nodes correspond to camera views and whose edges correspond to the estimated 3D transformations. The graph is then optimized with HOG-Man [11] to reduce the accumulated pose errors.

B. RTABMap

RTAB-Map (Real-Time Appearance-Based Mapping) [9] is a RGB-D, Stereo and Lidar Graph-Based SLAM approach based on an incremental appearance-based loop closure detector. The loop closure detector uses a bag-of-words approach to determine how likely a new image comes from a previous location or a new location. When a loop closure hypothesis is accepted, a new constraint is added to the map's graph, then a graph optimizer minimizes the errors in the map. A memory management approach described in [9] is used to limit the number of locations used for loop closure detection and graph optimization, so that real-time constraints on large-scale environments are always respected. RTAB-Map can be used alone with a handheld Kinect, a stereo camera or a 3D LIDAR for 6DoF mapping, or on a robot equipped with a laser rangefinder for 3DoF mapping.

IV.

DATASET

The TUM RGBD dataset [10] is a large set of data with sequences containing both RGB-D data and ground truth pose estimates from a motion capture system. The following seven sequences used in this analysis depict different situations and intended to test robustness of algorithms in these conditions.

C. freiburg2 desk with person

This sequence is a typical office scene with a person sitting at a desk. The person moves and interacts with some of the objects on the table. This sequence is intended for checking the robustness of a SLAM system against dynamic objects and

persons, but it can also be used to differentiate maps and find changes in the scene.

A. freiburg2 pioneer 360 robot slam

This sequence was recorded from a Kinect mounted on top of an ActiveMedia Pioneer 3 robot (See Fig. 1.). With these sequences, it becomes possible to demonstrate the applicability of SLAM systems to wheeled robots. Due to the large dimensions of the hall, the Kinect could not observe the depth of the distant walls for parts of the sequence. Several objects like office containers, boxes, and other feature-poor objects are placed through the scene. As a consequence, this sequence has depth, but is highly challenging for methods that rely on distinctive keypoints.

B. freiburg2 360 kidnap

In this sequence the Kinect sensor is covered several times while it is pointed to a different location ("kidnap") for testing algorithms that can recover from tracking problems.

C. freiburg3 long office household

The RGB-D sensor is moved along a large round through a household and office scene with much texture and structure. The end of the trajectory overlaps with the beginning so that there is a large loop closure.

D. freiburg3 no structure no texture near with loop

The RGB-D sensor is moved in approximately one meter height along a planar, wooden surface of approximately 3m x 3m. This sequence intentionally has little to no visible structure and texture. The beginning and the end of the sequence overlaps, i.e., there is a loop closure.

E. freiburg3 no structure texture near with loop

The RGB-D sensor has been moved in one meter height in a circle a textured, planar surface. The texture is highly discriminative as it consists of several conference posters. The beginning and the end of the trajectory overlap, so that there is a loop closure.

F. freiburg3 sitting static

In this sequence, two persons sit at a desk, talk, and gesticulate. The sensor has been kept in place manually. This sequence is intended to evaluate the robustness of visual SLAM and odometry algorithms to slowly moving dynamic objects.

V.

EVALUATION METRICS

The output of a SLAM algorithm is an estimated camera trajectory along with an estimate of the resulting map. While it is in principle possible to evaluate the quality of the resulting map, accurate ground truth maps are impossible to obtain as there are various uncontrollable factors involved. So we base our analysis on the quality of the estimated trajectory obtained from a sequence of RGB images. Even though an accurate trajectory does not imply a good map and error free operation

of robot, it is the most common parameter used to measure the accuracy of SLAM algorithms. For the evaluation, we assume that the output of the algorithm is a sequence of poses from the estimated trajectory $P_1, \dots, P_n \in SE$ and from the ground truth trajectory $Q_1, \dots, Q_n \in SE$.

For simplicity of notation, it is assumed that the two sequences are time-synchronized, equally sampled, and both have length n . However, in reality, these two sequences vary in sampling rates, lengths and have missing data, so an additional data association and interpolation step is required. Both sequences consist of homogeneous transformation matrices that give the pose of the RGB optical frame of the RGBD sensor from some other arbitrary reference frame. This reference frame does not have to be the same for both sequences, i.e., the estimated sequence might start at the origin, while the ground truth sequence is an absolute coordinate frame which was defined during calibration. While, in principle, the choice of the reference frame on the RGBD sensor is also arbitrary, the RGB optical frame is used as the reference because the depth images in the dataset have already been registered to this frame. In the remainder of this section, two common evaluation metrics for visual odometry and visual SLAM evaluation given in [4][4] are defined.

D. Absolute Trajectory Error (ATE)

For a visual SLAM system, the global consistency of the estimated trajectory is an important quantity. The Absolute Trajectory Error (ATE) is evaluated by comparing the absolute distances between the estimated and the ground truth trajectory. As both trajectories can be specified in arbitrary coordinate frames, they first need to be aligned. This is done in closed form by using the method of Horn [12], which finds the rigid-body transformation S corresponding to the least-squares solution that maps the estimated trajectory $P_{1:n}$ onto the ground truth trajectory $Q_{1:n}$. The absolute trajectory error at time step i can be calculated as

$$F_i := Q_i^{-1} S P_i \quad (1)$$

Evaluating the root mean squared error (RMSE) over all time indices of the translation components, we get,

$$RMSE(F_{1:n}, \Delta) := \left(\frac{1}{n} \sum_{i=1}^n \left\| \text{trans}(F_i) \right\|^2 \right)^{\frac{1}{2}} \quad (2)$$

A visualization of the absolute trajectory error is given in Fig. 2.. Here, RGB-D SLAM [8] was used to estimate the camera trajectory from the “fr2 360 pioneer slam” sequence.

E. Relative Pose Error (RPE)

The Relative Pose Error measures the local accuracy of the trajectory over a fixed time interval Δ . Therefore, the Relative Pose Error corresponds to the drift of the robot’s trajectory which is useful for the evaluation of visual odometry systems. We define the relative pose error at time step i as

$$E_i := (Q_i^{-1} Q_{i+\Delta})^{-1} (P_i^{-1} P_{i+\Delta})^{-1} \quad (3)$$

From a sequence of n camera poses, we obtain in this way $m = n - \Delta$ individual relative pose errors along the sequence.




Fig. 2. Visualization of Absolute Trajectory Error (ATE) of the estimated trajectory wrt ground truth trajectory

From these errors, we propose to compute the Root Mean Squared Error (RMSE) over all time indices of the translation component as

$$RMSE(E_{1:n}, \Delta) := \left(\frac{1}{m} \sum_{i=1}^m \left\| \text{trans}(E_i) \right\|^2 \right)^{\frac{1}{2}} \quad (4)$$

Where $\text{trans}(E_i)$ refers to the translation components of the Relative Pose Error E_i . Some prefer to evaluate the mean error instead of the root mean squared error as it affords less influence to outliers. Some use the median instead of the mean, which attributes even less influence to outliers. For visual odometry systems that match consecutive frames, the time parameter Δ is $\Delta=1$ which is an intuitive choice; $RMSE(E_{1:n})$ then gives the drift per frame. For systems that use more than one previous frame, larger values of Δ can also be appropriate, for example, $\Delta=30$ gives the drift per second on a sequence recorded at 30 Hz. A common choice is to set $\Delta=n$ which means that the start point is directly compared to the end point. This metric can be misleading as it penalizes rotational errors in the beginning of a trajectory more than towards the end [5], [13] and must not be used. It therefore makes sense to average over all possible time intervals Δ , i.e., to calculate

$$RMSE(E_{1:n}) := \frac{1}{n} \sum_{\Delta=1}^n \text{trans}(E_{1:n}, \Delta) \quad (5)$$

Note that the computational complexity of this expression is quadratic in the trajectory length. Therefore, it is approximated by calculating it from a fixed number of relative pose samples.

The RPE can be used to evaluate the global error of a trajectory by averaging over all possible time intervals. The RPE assesses both translational and rotational errors, while the ATE only assesses the translational errors. Therefore, the RPE is always greater than the ATE (or equal if there is no rotational error). Thus, the RPE metric gives us a way to combine rotational and translational errors into a single measure. However, rotational errors are also indirectly captured by the ATE as it manifest itself in wrong translations. From a practical perspective, the ATE has an intuitive visualization which facilitates visual inspection. Nevertheless, the two metrics are strongly correlated.

VI.

RESULTS

The evaluation of both the SLAM methods was carried out on the selected sequences from the dataset. The values of Absolute Trajectory Error, Relative Pose Error are root mean squared. The time required to process the entire sequence was recorded. All the values were plotted for ease of comparison (See Fig. 3.). The estimated trajectories and detailed results including min, max, mean, mode, and median values of errors are available at [14].

VII.

CONCLUSIONS

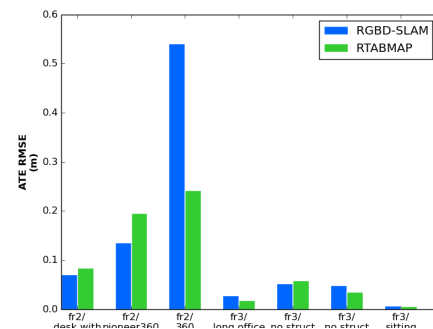
From the obtained results it can be concluded that RTABMap takes longer to process a sequence than RGB-D SLAM. The exception being fr2 360 kidnap; the reasons for which are discussed later on in this section. A general trend can not be observed from the Absolute Trajectory Error (ATE) RMSE comparison to clearly suggest which SLAM algorithm performs better. We must observe the estimated trajectory from both SLAM algorithms to analyze the problems faced by each.

The estimated trajectory comparison for fr2 360 kidnap sequence when juxtaposed with the ground truth (See Fig. 4.) reveals that RTABMap failed to recover after the sensor was covered. No estimated trajectory exists for the motion after the covering of the sensor for RTABMap. This is clearly a flaw with the used evaluation metrics as the evaluation ignores the motion for which no poses were generated. RGB-D SLAM recovers from the covering of the sensor and generates poses for the motion afterwards. Because of how the evaluation metrics are defined, RGB-D SLAM could be misidentified as the worse of the two algorithms in this situation despite its recovery from the tracking error.

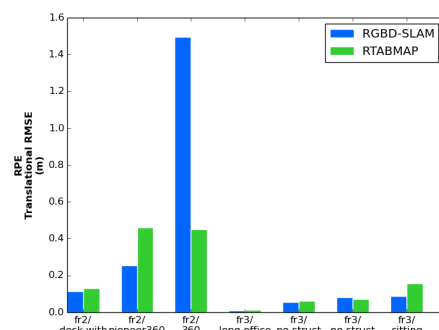
The time taken by RTABMap to process the sequence is less than the time taken by RGB-D SLAM due to the same reason. RTABMap fails to recover from the tracking problem and does not process the remaining sequence thus resulting in a shorter processing time.

In case of fr3 no structure no texture sequence, the same flaw is observed (See Fig. 5.). RTABMap fails to generate valid poses for the latter half of the motion. The algorithm lost tracking due to lack of features in the frame. The algorithm never recovered and hence we have an incomplete estimation of the trajectory.

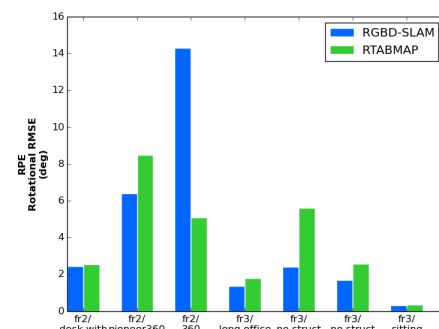
It must be emphasized that RTABMap generates significantly lesser number of pose pairs for evaluation than RGB-D SLAM. Hence an accurate comparison is not possible with the raw estimated trajectories generated by both methods.



(a) Absolute Trajectory Error (RMSE)



(b) Relative Pose Error Translational (RMSE)



(c) Relative Pose Error Rotational (RMSE)

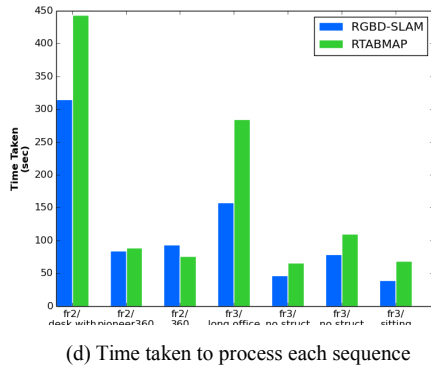
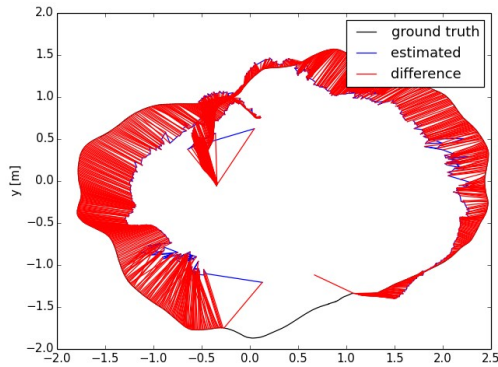



Fig. 3. Result of the benchmark evaluation over seven different sequences



(a) RGB-D SLAM recovers and generates estimated trajectory for the entire motion except when it sensor was covered



(b) RTABMap fails to generate estimated trajectory after the sensor was covered


Fig. 4. ATE visualization of fr2 360 kidnap sequence with RGB-D SLAM and RTABMap

After the analysis of the results it can be concluded that RGB-D SLAM has a better performance overall in the tested circumstances while ignoring the underlying flaws in the evaluation metrics. It has a shorter processing time. It is better at recovering from tracking errors than RTABMap. It also sustains fast camera motion much better than RTABMap.


A need for a better evaluation metric was realized due to the aforementioned flaws in the evaluation metrics as observed from the results.

Acknowledgment

The author would like to thank Jurgen Sturm from the Computer Vision Group at TUM for without his dataset and benchmark evaluation method this paper would not have been possible. The author also appreciates the support and resources provided by Research Center at Pune Institute of Computer Technology (PICT), India. The author also thanks Dr. Gaurav Bansod and Dr. Geetanjali Kale for their guidance and support.



(a) RGB-D SLAM generates estimated trajectory for the



(b) RTABMap fails to generate estimated trajectory after losing track due to lack of features

Fig. 5. ATE visualization of fr3no structure no texture sequence with RGB-D SLAM and RTABMap

References

- [1] R. C. Smith and P. Cheeseman, "On the representation and estimation of spatial uncertainty," *The international journal of Robotics Research*, vol. 5, no. 4, pp. 56–68, 1986.
- [2] G. Dissanayake, H. Durrant-Whyte, and T. Bailey, "A computationally efficient solution to the simultaneous localisation and map building (slam) problem," in *Robotics and Automation, 2000. Proceedings.*

- ICRA'00. IEEE International Conference on, vol. 2, pp. 1009–1014, IEEE, 2000.
- [3] M. Montemerlo and S. Thrun, “Simultaneous localization and mapping with unknown data association using fastslam,” in ICRA, pp. 1985–1991, 2003. K. Elissa, “Title of paper if known,” unpublished.
- [4] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, “A benchmark for the evaluation of rgb-d slam systems,” in Proc. of the International Conference on Intelligent Robot Systems (IROS), Oct. 2012.
- [5] R. Kümmerle, B. Steder, C. Dornhege, M. Ruhnke, G. Grisetti, C. Stachniss, and A. Kleiner, “On measuring the accuracy of slam algorithms,” *Autonomous Robots*, vol. 27, no. 4, p. 387, 2009
- [6] F. Steinbrücker, J. Sturm, and D. Cremers, “Real-time visual odometry from dense rgb-d images,” in Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on, pp. 719–722, IEEE, 2011
- [7] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, pp. 3354–3361, IEEE, 2012.
- [8]] F. Endres, J. Hess, N. Engelhard, J. Sturm, D. Cremers, and W. Burgard, “An evaluation of the rgb-d slam system,” in Robotics and Automation (ICRA), 2012 IEEE International Conference on, pp. 1691–1696, IEEE, 2012
- [9] M. Labbé and F. Michaud, “Memory management for real-time appearance-based loop closure detection,” in Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on, pp. 1271–1276, IEEE, 2011
- [10] TUM, “Rgb-d slam dataset and benchmark.” <https://vision.in.tum.de/data/datasets/rgbd-dataset>, 2012. Accessed: 2018-09-30
- [11] G. Grisetti, R. Kümmerle, C. Stachniss, U. Frese, and C. Hertzberg, “Hierarchical optimization on manifolds for online 2d and 3d mapping,” in Robotics and Automation (ICRA), 2010 IEEE International Conference on, pp. 273–278, IEEE, 2010
- [12] B. K. Horn, “Closed-form solution of absolute orientation using unit quaternions,” *JOSA A*, vol. 4, no. 4, pp. 629–642, 1987
- [13] A. Kelly, “Linearized error propagation in odometry,” *The International Journal of Robotics Research*, vol. 23, no. 2, pp. 179–218, 2004
- [14] A. Kasar, “Git repository for the results and scripts.” <https://github.com/ameykasar/benchmarking-slam>, 2018. Accessed: 2018-11-22

