

Removing Moving Objects from Point Cloud Scenes

Krystof Litomisky and Bir Bhanu

University of California, Riverside
krystof@litomisky.com, bhanu@ee.ucr.edu

Abstract. Three-dimensional simultaneous localization and mapping is a topic of significant interest in the research community, particularly so since the introduction of cheap consumer RGB-D sensors such as the Microsoft Kinect. Current algorithms are able to create rich, visually appealing maps of indoor environments using such sensors. However, state-of-the-art systems are designed for use in static environments, which severely limits the application space for such systems. We present an algorithm to explicitly identify and remove moving objects from multiple views of a scene. We do this by finding corresponding objects in two views of a scene. If the position of an object with respect to the other objects changes between the two views, we conclude that the object is moving and should therefore be removed. After the algorithm is run, the two views can be merged using any existing registration algorithm. We present results on scenes collected around a university building.

Keywords: SLAM, 3D Mapping, RGB-D sensors, Kinect

1 Introduction

Although point clouds and sensors that provide point cloud data have been around for several decades, the introduction in 2010 of the Microsoft Kinect RGB-D (RGB color + per-pixel depth) sensor reinvigorated the field. One popular area of research has been using RGB-D sensors for Simultaneous Localization and Mapping (SLAM) in primarily indoor environments. State of the art systems have achieved impressively accurate results, producing visually-appealing maps of moderately-sized indoor environments [1]. Such algorithms typically rely on the Iterative Closest Point (ICP) for point cloud registration, and also incorporate loop-closing techniques to detect when an agent has returned to a previously visited area [2].

In published work on SLAM using RGB-D sensors in indoor environments, such as [1], [3], and [4], the maps are created in static environments: there are no moving objects, and, in particular, no people walking around. In more dynamic environments, not addressing moving objects can lead to maps that contain moving objects as permanent features, inconsistent maps, or even registration failure. Our work addresses this issue by introducing a novel preprocessing step to explicitly identify and remove moving objects from point cloud frames prior to registration. A sample result of our algorithm is in Fig. 1, and an overview of our approach is in Fig. 2.

To identify moving objects, we compare two frames with significant overlap. The viewpoint between the two frames can change, and some amount of time must elapse before the second frame is captured so that the position of moving objects changes between the frames. For each frame, we segment out individual clusters, and find which cluster from the first frame corresponds to each cluster from the second. We then analyze the spatial relationship of each cluster in a frame to all other clusters in the frame. If this relationship changes from one frame to the next, we conclude that the cluster in question must be moving and remove it. Having done this, we can apply any existing registration algorithm to register, align, and merge the two clouds.




Fig. 1. A sample result of our algorithm. The top row shows two input point clouds. Our algorithm identified the person as a moving object and removed him from each point cloud. We then aligned and merged the clouds to produce the cloud in the bottom row.

2 Related Work

3D SLAM has been an area of immense research interest. Over the years, a number of approaches using different technologies have been developed, including range scans [3, 4], stereo cameras [7], monocular cameras [8], and recently also consumer RGB-D cameras [1][3][4]. However, explicitly identifying and removing moving objects from point cloud data has not been a topic of great interest in the research community.

In their RGB-D mapping work, Henry et al. [1] do not address moving objects, and do not present any results from environments with moving objects. However, the use of a surfel representation [9] allows them to deal with some instances of moving objects. Because surfels have an area, they permit reasoning about occlusion. This ultimately enables this representation to remove moving objects by eliminating surfels which occlude other surfels. However, this approach will fail to eliminate moving objects when these objects do not occlude other known objects. Furthermore, converting all data to surfels is computationally inefficient.

Similarly, the Kinect Fusion work [10], which implements a real-time version of ICP that runs on CUDA-enabled GPUs, deals with moving objects implicitly by using a volumetric representation. Kinect Fusion is not primarily a SLAM algorithm; as a result, the approach does not scale to larger environments. In particular, with current

GPUs, the approach cannot be used for environments larger than a moderately small room (approximately a 5m x 5m x 5m volume).

Other volumetric approaches, such as the Octree-based OctoMap [6], also identify moving objects only implicitly. The OctoMap’s probabilistic nature means that it may require a substantial amount of measurements until a moving object is discarded.

Unlike previous work, which deals with moving objects in only an implicit and incomplete way, this paper presents a focused and systematic approach for removing moving objects from point cloud scenes captured from a moving platform. An advantage of our approach is that it deals directly with the point cloud data. This means that after eliminating moving objects with our approach, any existing algorithm for point cloud registration – or any other application – can be applied straightforwardly.

3 Technical Approach

A high-level overview of our approach is given in Fig. 2. For notational clarity, we will refer to the first point cloud we process as the “source” cloud, and to the second cloud we process as the “target” cloud. Our approach is symmetric, however, and which cloud is designated as “source” does not affect our results. For some of the point cloud manipulation tasks, we rely on the Point Cloud Library (PCL) [11].

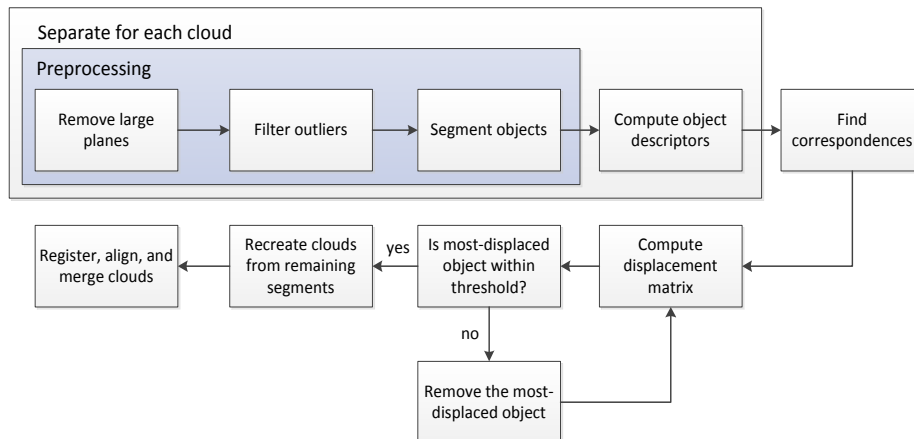


Fig. 2. An overview of our system.

3.1 Preprocessing

Before identifying moving clusters, we temporarily remove large planes, filter outlier points, and run segmentation to identify individual clusters.

We identify planes using a Random Sample Consensus (RANSAC) algorithm. The reason for removing large planes is twofold: first, we can safely assume that large planes are not parts of moving objects, so we do not need to consider them in subsequent steps of the algorithm. Second, removing large planes – in particular, the floor –

improves the performance of our segmentation algorithm. Having completed this step, we filter out outlier points to remove artifacts due to sensor noise.

We now use a Euclidean Cluster Extraction algorithm to get individual clusters. This algorithm places points that are less than the cluster tolerance apart in the same cluster. Since large planar objects as well as statistical outliers have already been removed from the clouds at this point, this is an acceptable approach. Due to quantization of the depth data, which is inevitable with consumer RGB-D sensors, we need to use a somewhat large cluster tolerance value. We settled on 15 cm, meaning two points will be in the same cluster if they are less than 15 cm apart. We use a k-d tree to speed up the cluster extraction process.

Each point cloud can now be represented as a set of clusters. Let $C_s = \{s_1, s_2, \dots, s_{n_s}\}$ be the source point cloud, and $C_t = \{t_1, t_2, \dots, t_{n_t}\}$ be the target point cloud. Here, s_i is the set of points representing the i^{th} cluster in the source cloud, and t_i is the set of points representing the i^{th} cluster of the target cloud.

3.2 Cluster Descriptors and Identifying Correspondences

We use the Viewpoint Feature Histogram (VFH) proposed by Rusu et al. [12] as our feature descriptor. The VFH has good object classification performance, outperforming spin images for these tasks [12]. The VFH stores the relationships between the pan, tilt, and yaw angles between pairs of normals in a cluster, as well as the relationships between the viewpoint direction and the surface normals.

We now find cluster correspondences, or which cluster from the target cloud corresponds to each cluster from the source cloud. We calculate the distance between a pair of clusters in feature space in two ways and compare the performance in Section 4.

The first distance measure is the sum of absolute differences between two clusters' VFH descriptors, while the second distance measure comes from an insight on the ordering of the VFH: the bins for each angle "category" are consecutive. For example, if the yaw angle between two points changed slightly, the pair would move from one bin into a bin immediately next to it. As a result, small object deformations (such as a person's pose changing as he walks) as well as changes in sensor position cause non-linear local deformations to the object's histogram. Fig. 3 illustrates the local deformations of a part of the VFH due to sensor motion.

Such issues have been addressed in the literature with Dynamic Time Warping [13]. Dynamic time warping finds a nonlinear, monotonic mapping from one sequence to the other, allowing for small shifts and deformations between the series. To limit the amount of warping, plausible values, we use the Sakoe-Chiba band [14] to limit warping to 2% of the total histogram length, corresponding to the shift in histogram resulting from a camera pan of about 45 degrees.

To identify correspondences, we iteratively take the closest pair of clusters (s_i, t_j) in feature space as corresponding to each other, until there are no clusters left in at least one cloud. Due to sensor and/or object motion, there may be leftover objects, which we remove from their cloud. This may remove objects that could provide useful data, but this is an acceptable tradeoff to ensure the removal of any moving objects that appear in only one frame. This step leaves $n = \min(n_s, n_t)$ clusters for each

cloud. Before proceeding further, we reorder the clusters in C_t such that cluster s_i corresponds to cluster $t_i \forall i$.




Fig. 3. Local deformations of the yaw component of the Viewpoint Feature Histogram of the mannequin extracted from the two views above. Histogram for the mannequin on the left is solid blue; right is dashed red. Dynamic Time Warping addresses such local deformations.

3.3 Identifying and Removing Moving Objects

We now calculate the Euclidean distance in world-coordinate space between each pair of clusters for each cloud. Let $d_{i,j}^s$ be the world-coordinate-space Euclidean distance between cluster i and cluster j in the source point cloud, and $d_{i,j}^t$ be the world-coordinate-space distance between the corresponding clusters in the target cloud.

To get a measure of how the position of each cluster has changed from one cloud to the next, we calculate the displacement vector $\Delta = [\delta_1, \delta_2, \dots, \delta_n]^T$, where δ_i is the displacement of cluster i ,

$$\delta_i = \sum_{k=1}^n |d_{i,k}^s - d_{i,k}^t| \quad (1)$$

In essence, δ_i is the sum of how much the distance of cluster i to each other cluster has changed from one cloud to the other cloud.




Fig. 4. Object removal threshold ROC curve. A true positive is a removed moving object, and a false positive is a removed static object.

We now iteratively remove the cluster which has the greatest displacement value as long as this value is above a threshold ε . After removing each cluster, we recalculate Δ . In order to find the optimal value of ε , we generated a ROC curve (Fig. 4). For the ROC curve, a true positive is a moving object that was removed from the scene, and a false positive is a static object that was removed from the scene. See Section 4 for details regarding the data used.

We now iteratively remove the cluster which has the greatest displacement value as long as this value is above a threshold ε . After removing each cluster, we recalculate Δ . In order to find the optimal value of ε , we generated a ROC curve (Fig. 4). For the ROC curve, a true positive is a moving object that was removed from the scene, and a false positive is a static object that was removed from the scene. See Section 4 for details regarding the data used.

We ran our algorithm for different thresholds for both the histogram difference and dynamic time warping distance measures, achieving better results with dynamic time warping. In particular, we were able to correctly remove all moving objects with $\varepsilon = 0.7$ meters when using dynamic time warping. We therefore use this value in our experiments (Section 4).

Having removed all of the moving objects from C_s and C_t , we reconstruct each cloud from its remaining clusters as well as the planes that had been removed from it. After this, the clouds are in a standard point cloud format, and any existing registration, alignment, and merging algorithms can be used to concatenate the clouds. The result of this operation is a point cloud model of the environment with no moving objects in the point cloud, even though such objects might have been present in one or both of the original clouds.

4 Experiments

We tested our algorithm on 14 scenes collected around a university building. Each scene consists of two views, and each view has up to 2 moving people. Fig. 5 shows some results, demonstrating several different scenarios our algorithm can deal with.

Fig. 5(a) shows a scene with multiple moving people, as well as significant camera motion between the two frames. As a result of this, we cannot find good correspondences for some objects in the scene, such as the desk to the right of the second view. Note that there is a third person in the scene, sitting at a desk in the background. This person is put into the same cluster as the desk he is sitting at by our algorithm, and since he is not moving he is included in the recreated cloud.

Fig. 5(b) shows scene where a person who is not present in one point cloud moves into view by the time the second point cloud is captured. Our algorithm correctly identifies that there is no equivalent object for the person in the other frame, and removes the person from the point cloud.

Fig. 5(c) shows a scene where a walking person completely changes direction from one frame to the next. Nevertheless, our algorithm matches the person in one frame to the person in the next frame and correctly removes him from both clouds.



Fig. 5. Example results. (a) office scene with three people (two walking, one sitting). The top two frames from columns (a) and (b) were merged together to produce the bottom point cloud. (b) corridor scene where the person just entered the corridor through a door on the right. (c) a walking person changing direction.

4.1 Quantitative Analysis

Fig. 4 shows the ROC curve obtained by running our algorithm on the dataset, where a true positive is defined as a moving object that is removed from the scene and a false positive is defined as a static object that was removed from the scene. We get better results with the dynamic time warping distance measure than with the simple histogram difference. In particular, for the object removal threshold $\epsilon = 0.7$, we correctly remove all moving objects while also removing 47% of the static objects. In a

SLAM scenario, we would keep a greater fraction of static objects, since objects that are removed due to occlusion or sensor motion at first would likely be seen again at a later time and thus ultimately kept.

We also evaluate what fraction of the total points that belong to stationary objects we keep in the recreated clouds. For each scene, we calculate this number separately for each of the two frames, and then report their mean and standard deviation. Fig. 6 shows the results. On average, we keep 85% of the static points in a scene.




Fig. 6. Fraction of static points kept for each of the 14 scenes on which we evaluated our algorithm. Numbers along the x-axis indicate which figure corresponds to the particular scene, if applicable. The mean is 85% of static points retained.

5 Conclusions and future work

We introduce early work on an algorithm for identifying moving objects in point cloud scenes. We can eliminate moving objects from scenes while retaining most of the static objects. Thus, when used as a preprocessing step, our approach can complement existing point cloud-based SLAM algorithms. This will allow existing SLAM approaches to create consistent maps even in the presence of moving objects, making the system applicable in more scenarios.

Future work will include incorporating the approach presented here into an autonomous SLAM system, and making any necessary adjustments necessary to make the system work properly in real-time. Considerations will include tuning plane removal as well as segmentation, and studying the temporal stability of the approach.

References

- [1] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox, “RGB-D Mapping: Using depth cameras for dense 3D modeling of indoor environments,” in *the 12th International Symposium on Experimental Robotics (ISER)*, 2010.
- [2] P. Newman, “SLAM-Loop Closing with Visually Salient Features,” in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pp. 635-642.
- [3] H. Andreasson and A. J. Lilienthal, “6D scan registration using depth-interpolated local image features,” *Robotics and Autonomous Systems*, vol. 58, no. 2, pp. 157-165, Feb. 2010.

- [4] H. Du et al., "Interactive 3D modeling of indoor environments with a consumer depth camera," in *Proceedings of the 13th international conference on Ubiquitous computing - UbiComp '11*, 2011, p. 75.
- [5] S. May et al., "Three-dimensional mapping with time-of-flight cameras," *Journal of Field Robotics*, vol. 26, no. 11–12, pp. 934-965, Nov. 2009.
- [6] K. M. Wurm, A. Hornung, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: A probabilistic, flexible, and compact 3D map representation for robotic systems," in *Proc. of the ICRA 2010 Workshop on Best Practice in 3D Perception and Modeling for Mobile Manipulation*, 2010.
- [7] T. Lemaire, C. Berger, I.-K. Jung, and S. Lacroix, "Vision-Based SLAM: Stereo and Monocular Approaches," *International Journal of Computer Vision*, vol. 74, no. 3, pp. 343-364, Feb. 2007.
- [8] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse, "MonoSLAM: real-time single camera SLAM.," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 6, pp. 1052-67, Jun. 2007.
- [9] H. Pfister, M. Zwicker, J. Van Baar, and M. Gross, "Surfels: Surface elements as rendering primitives," in *Proceedings of the 27th annual conference on Computer Graphics and Interactive Techniques*, 2000, pp. 335–342.
- [10] R. Newcombe et al., "KinectFusion: Real-time dense surface mapping and tracking," in *2011 10th IEEE International Symposium on Mixed and Augmented Reality*, 2011, pp. 127-136.
- [11] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in *Proceedings of the 2011 IEEE International Conference on Robotics and Automation, Shanghai, China*, 2011.
- [12] R. B. Rusu, G. Bradski, R. Thibaux, and J. Hsu, "Fast 3D Recognition and Pose Using the Viewpoint Feature Histogram," in *International Conference on Intelligent Robots and Systems*, 2010, pp. 2155-2162.
- [13] D. Berndt and J. Clifford, "Using dynamic time warping to find patterns in time series," in *Conference on Knowledge Discovery and Data Mining - KDD*, 1994, vol. 398, pp. 359-370.
- [14] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 26, no. 1, pp. 43-49, Feb. 1978.