Incremental Object Database: Building 3D Models from Multiple Partial Observations

Conference Paper**Author(s):**

Furrer, Fadri; Novkovic, Tonci; [Fehr, Marius](#) ; Gawel, Abel; Grinvald, Margarita; Sattler, Torsten; Siegwart, Roland; Nieto, Juan

Publication date:

2018

Permanent link:

<https://doi.org/10.3929/ethz-b-000318803>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/IROS.2018.8594391>

Incremental Object Database: Building 3D Models from Multiple Partial Observations

Fadri Furrer, Tonci Novkovic, Marius Fehr, Abel Gawel, Margarita Grinvald,
Torsten Sattler, Roland Siegwart, and Juan Nieto

Abstract—Collecting 3D object datasets involves a large amount of manual work and is time consuming. Getting complete models of objects either requires a 3D scanner that covers all the surfaces of an object or one needs to rotate it to completely observe it. We present a system that incrementally builds a database of objects as a mobile agent traverses a scene. Our approach requires no prior knowledge of the shapes present in the scene. Object-like segments are extracted from a global segmentation map, which is built online using the input of segmented RGB-D images. These segments are stored in a database, matched among each other, and merged with other previously observed instances. This allows us to create and improve object models on the fly and to use these merged models to reconstruct also unobserved parts of the scene. The database contains each (potentially merged) object model only once, together with a set of poses where it was observed. We evaluate our pipeline with one public dataset, and on a newly created Google Tango dataset containing four indoor scenes with some of the objects appearing multiple times, both within and across scenes.

I. INTRODUCTION

Humans have an excellent understanding of an environment’s structure. Once we observe the objects in a scene, we efficiently memorize them and can easily recall the complete shape of these objects regardless of occlusions or only partial observations. We can furthermore apply the knowledge of repeating objects to immediately recognize them in new scenes, and even hallucinate their full shape despite partial observations. However, modern robotic systems usually lack a notion of objects and merely work based on their sensor data, e.g., storing scenes as abstract point clouds. This can be highly inefficient for scene reconstructions, as any structure needs to be observed from multiple view-points and leads to large map sizes. During the reconstruction process, repetitive structures typically result in ambiguous data associations. Furthermore, such systems exhibit a lack of scene understanding, creating a need for post-processing steps that perform object recognition. Contemporary approaches for 3D object recognition require databases of pre-recorded objects to be able to match these to the scene data. Typically, the creation of these databases involves significant human labor in terms of either hand-modeled 3D meshes [1] or 3D models obtained via careful manual scanning [2].

F. Furrer, T. Novkovic, M. Fehr, A. Gawel, M. Grinvald, R. Siegwart, and J. Nieto are with the Autonomous Systems Lab, ETH, 8092 Zurich, Switzerland, e-mail: {fadri.furrer, tonci.novkovic, marius.fehr, abel.gawel, margarita.grinvald}@mavt.ethz.ch, {rsiegwart, nietoj}@ethz.ch.

T. Sattler is with the Computer Vision Group, Department of Computer Science, ETH, 8092 Zurich, Switzerland, e-mail: sattlert@inf.ethz.ch.

Find the accompanied video here: https://youtu.be/9_xg92qqw70.

Fig. 1: Input depth images, top left, get segmented using a geometrical segmentation (III-A), the individual depth segments, top right, are integrated into a Global Segmentation Map (GSM) (III-B), middle right, from which we can extract raw segments. We put these raw segments in a database of object models and describe their keypoints (III-C.2), middle left, and try to match and merge them with other models built from observations of instances of the same object kind (III-C.3, III-C.4), bottom left. With these merged object models, we can reconstruct the scene and complete it by combining the data of all observed object instances, bottom right. Different colors indicate different segments.

In this work, we present an automatic system to incrementally build a database of 3D object models using depth sensing. Rather than aiming to detect a fixed set of objects in the RGB images provided by an RGB-D sensor [3], [4], [5], our approach segments the depth image provided by the sensor based on an edge and convexity map similar to [6], [7]. In subsequent steps, the initial depth segments are integrated into a Truncated Signed Distance Field (TSDF) stored in a voxel grid resulting in a Global Segmentation Map (GSM). From this GSM, raw segments are extracted if the corresponding voxels are untouched for a certain time, described using features, and inserted in an object database as object models. In the database they are matched to existing object models, which get refined and completed over time. By re-projecting the merged object models to all locations

where the object instances were observed a scene can be completed. This process is illustrated in Fig. 1.

One of the advantages of our approach is that we do not need to pre-define which objects we want to detect, rather, we automatically detect objects based on a segmented scene reconstruction. This is a prerequisite for building a system that makes scanning objects as simple as walking through an everyday scene, without the need to place the objects in front of a scanning setup, e.g. a turn-table. Our method can be used to construct 3D maps consisting of planar segments, which are geometrically not distinctive, and a set of (potentially incomplete) object models. Since these maps are constructed incrementally, our approach is directly applicable to scenarios where a robot (or another type of autonomous agent) needs to interact with the world on an object level in order to solve its tasks.

This paper makes the following contributions:

- A novel approach to incrementally build a database of objects from one or multiple sessions.
- Object completion capabilities, facilitating TSDF merging of multiple occurrences of the same object.
- Integration of the proposed automatic object database building with a full scene reconstruction framework.
- The release of the Tango RGB-D dataset to serve as benchmarks for comparisons.

II. RELATED WORK

In order to create a database of 3D objects, our approach first segments the input depth data, constructs a GSM and then inserts all the segments into the database. Newly detected segments are then compared with the object models in the database to determine whether a new object model is required or an existing model needs to be updated. In the following, we review literature relevant to the individual stages of our approach.

Scene segmentation is a very active research topic in the computer vision community and often a necessary prerequisite for certain tasks in robotics such as manipulation [8], object detection [9], scene understanding [10], etc. Numerous approaches have been developed in order to obtain meaningful segments from a scene. These can be extracted from RGB images [11], RGB-D data [6], [9], [7], image sequences [12], or 3D data such as point clouds [13], meshes [14], [10], voxel grids [15], or using deep neural networks [11], [9], [12]. However, due to the requirement for large computational resources, deep neural networks are not applicable on most hardware constrained robotic systems. Uckermann et al. [6] demonstrated a real-time capable segmentation method on depth images, based on a surface normal edge map which captures discontinuities of depth measurements, sudden changes of surface normals at object edges, and performs region growing of surface patches into segments. The same idea was explored by Tateno et al. [7] to provide a fast segmentation method for depth images. Karpathy et al. [10] demonstrated an approach that partitions a scene mesh.

Our segmentation approach is similar to that of [16]. We modified the normal estimation such that instead of just using three points in the vicinity of the midpoint, we use a kernel, resulting in a smoother normal map. Additionally, instead of using the original edge map, we consider edges from a depth discontinuity map calculated based on [17].

Often, it is beneficial to fuse multiple segmented frames into a consistent map. Finman et al. [18] implemented a method that incrementally stores segmented depth data. They use an incremental variation of the algorithm from [14] to segment new data and a voting algorithm for recomputing parts of the **GSM** based on this new data, which ensures global consistency. In their approach, the GSM relies on a 3D TSDF representation [19]. Tateno et al. [20] use a TSDF volumetric surface representation for the GSM and a SLAM system to keep track of the camera poses. Our approach is based on [20]. In contrast to [20], we keep track of the entire history of segment labels associated to a TSDF voxel and assign that voxel to the segment with the highest count. This approach increases robustness to noisy per-frame segmentation and allows us to deal more efficiently with merges of multiple segments into one, or splits of one segment into multiple ones.

From the GSM, raw segments can be extracted and added to the **database of objects**. In contrast to some of the object datasets that are generated offline in controlled conditions [17], [2], the GSM allows us to add segments to the database at any point during a mapping session. Furthermore, if new information is observed, object models can easily be updated and completed online, without any post-processing required. Dai et al. [21] developed a **shape completion** method that predicts and fills in the missing data from the input. First, a fixed-size voxel volume is predicted using a 3D CNN and then a higher resolution model is synthesized based on an offline shape database. A similar approach was developed by Han et al. [22]. However, both approaches can only complete shapes that they have been trained for. Our approach does not require any a-priori knowledge about the objects and, furthermore, does not approximate objects but rather uses previously observed data to complete unobserved parts. This also means that objects in the database will not be completed unless the missing data is observed in one of the sessions on at least one of the object instances.

Learning-based **object detection** methods usually try to classify patches in images [23], [24], however, the number of classes such algorithms can detect is restricted by the training data provided to the algorithm. Alternatively, database-assisted approaches such as [1] rely on a limited number of pre-defined objects in the database, which are then detected in the scene using keypoints and descriptors. In this work, we do not make assumptions on the objects in the scene. Therefore, once we obtain the raw segments from the GSM, we use keypoints [25], [26] and descriptors [27] to represent those, and then match them to the previously obtained object models in the database. Our approach is similar to [7], however, instead of global descriptors we use local descriptors and keypoints, allowing us to match

Fig. 2: The output of our depth segmentation for a single frame of the Tango dataset. Top left: rectified depth image, bottom left: filtered depth image, top center: depth discontinuity edges, bottom center: convexity map, top right: normal map, bottom right: labeled image.

individual parts of the models. After we obtain the good correspondences from matching, we detect the pose using Random Sample Consensus (RANSAC) [28] and Iterative Closest Point (ICP), as in [29]. The main difference of our approach to [20] is that we are able to, by merging individual segments, complete models of objects that we have not seen before, and therefore do not need to know them beforehand.

III. METHOD

In this section, we describe how we get from RGB-D images to object models in the database. First, we use a geometric approach to *segment* input depth images. In a second step, we employ a TSDF-based *Global Segmentation Map (GSM)*. The GSM not only fuses depth measurements into a 3D reconstruction but obtains improved and temporally consistent labels, which ideally means one label per object instance. Segments which remain unchanged in the GSM for a certain time are extracted (raw segments) and inserted into an initially empty *database*. Once these raw segments are in the database, we call them object models. Afterwards, we match them to other models in an attempt to combine them into more complete and accurate object models. Thus, we are trying to match them to either models of other instances of the same object or partial models of the same instance, created due to the oversegmentation of the GSM. A third option is to match to object models from separate recordings.

A. Depth Segmentation

In order to segment RGB-D images, we use the fact that the objects we want to segment always have a closed shape. Therefore, we apply similar filters as described in [6], [7] to the depth images. In a first step, we inpaint the depth image to obtain more continuous areas with valid depth values. This step is optional and its necessity depends on the specifications and configuration of the RGB-D sensor. In a second step we detect edges that exhibit strong depth discontinuities. Next, we compute surface normals based on a local pixel neighborhood and use them to determine the local convexity of each pixel. In a final step, we combine the convexity map and the depth discontinuity filter to form

closed regions and extract contours which we fill with a label. However, these labels are not consistent across different frames and a nearly identical region might be assigned a different label in the next frame. The depth measurements of each region can then be used to obtain a labeled 3D segment. The results of the different steps are shown on a sample depth image of a Lenovo Tango Phab 2 Pro in Fig. 2

B. Global Segmentation Map

The goal of the GSM is to merge the frame-wise segmentation into a more accurate and globally consistent object instance segmentation. To that end, the segments extracted from each depth image are incrementally integrated into a TSDF-based voxel grid capable of storing and fusing labels for every voxel. The proposed GSM serves the same purpose and shares some of the concepts with the surfel-based segmentation map proposed in [7]. Since our proposed system does not include a camera pose estimation stage, the camera poses for each depth map must be provided by an external estimation pipeline or ground truth data.

The GSM builds on top of Voxblox [30], a real-time reconstruction framework based on a volumetric TSDF surface representation. The Voxblox framework has been extended with a second volume, the label volume, storing the segment label associated with each voxel in the TSDF grid. At each frame, the set of point clouds representing the 3D shapes of all the identified segments are fused into the voxel-based representation, with the GSM ensuring consistency in the segment labels across different frames. The computational complexity of the method does not depend on the size of the map or the number of merged frames, and the resulting segmentation and reconstruction of the scene are obtained at interactive rates.

The main difference between our method and [7] lies in the way the label volume is updated at each new segmented input depth map. Instead of storing just one segment label and one confidence value at each voxel, we store the complete history of all the segment labels that have ever been merged into this voxel, together with the respective counts. The label with the highest count is then set to be the main segment label associated with that voxel. At the cost of additional memory usage, this approach is more robust towards noise in the per-frame segmentation. Correct frame-wise segmentation outputs can contribute to the highest label observation count in a voxel, independently of whether they have been followed by a number of different improper ones. This, in contrast, is not possible when approximating all the votes with a single value as in [7]. Furthermore, our method enables lossless merging of two or more segments in the map which have been detected to be part of the same object. In the frames leading up to the segments getting merged, some of them are updated with the labels of the others, until it is detected that those multiple different labels are actually one. When only one label count is stored in a voxel as in [7], these updates lead to the lowering of the previously gained confidence for the merged segments, while our approach allows to recover

Fig. 3: The keypoints are matched among two point clouds (green and red), the green cloud is registered to the red one using RANSAC (blue) and this registration is refined using ICP (magenta). Green lines indicate the good matches. Both images result in a match with a good ICP Root-mean-square error (RMSE). (a) will be rejected by the geometric verification step described in Section III-C.4, whereas (b) is a correct match.

the total number of correct votes for each voxel by simply summing up the label counts for the merged segments.

For each 3D segment representing a distinct object in the scene and uniquely identified by a label, the GSM also implicitly stores its pose in a world frame. A segment is extracted from the GSM when none of the voxels corresponding to the segment have been updated for a certain period of time. The TSDf grid of this extracted raw segment is sent to the database together with the corresponding segment label and its world pose.

C. Incremental Object Database

The object models s_i in the database Ω consist of the following components: a set $\mathbb{T}_i$ of poses T_k ¹ for all model observations constituting s_i , where the first element of the set $\mathbb{T}_i$ defines the base frame of the model s_i , a TSDf-grid t_i retaining the 3D shape of the model, the corresponding surface as a point cloud with normals c_i , as well as 3D keypoints k_i , and their 3D feature descriptors d_i .

Upon insertion into the database, the labeled segments extracted from GSM are not full object models yet. In this section we first explain how we complete the object model with keypoints and descriptors followed by the matching and merging of object models.

1) *Point Cloud Extraction*: In order to merge raw segments extracted from the GSM into complete object models, we need to match and register these segments. First we use the marching-cubes surface reconstruction algorithm [31] on the TSDf grid to obtain a point cloud with surface normals, which is required to extract keypoints and descriptors for matching. We use a RANSAC-based planarity check to exclude planar segments from merging as they do not provide enough constraints to allow for meaningful matches, based on their geometry.

¹All the transformations with one subscript denote transformations from the world origin to the object model base frame.

Fig. 4: Geometric verification of object models. The evaluated models are represented by the *cyan* and *magenta* point clouds. The TSDf errors between the two voxel grids of the models are visualized as one point per overlapping voxel, colored *green* to *red*, where *red* represents a high error.

2) *Keypoints and Descriptors*: From the remaining segments, we extract keypoints and describe their neighborhood in the point cloud using descriptors. We use a combination of Intrinsic Shape Signature (ISS) [26] and Harris3D [25] keypoints. The two types of keypoint detectors complement each other well for our application. The ISS detector is very efficient at the cost of reduced repeatability [32], therefore, we used it to extract a larger number of keypoints even in smoother areas, whereas Harris3D was tuned to provide fewer but more repeatable keypoints.

To describe the resulting keypoints, we use Fast Point Feature Histogram (FPFH) descriptors [27], a very efficient but low-dimensional descriptor that is based on the surface normals in a spherical neighborhood of radius r . While the efficiency of FPFH allows fast matching to a large number of object model candidates, it requires strong geometric consistency checks to compensate for the limited expressiveness.

3) *Matching and Registration*: After obtaining an object model as defined at the beginning of this section, we want to find database entries that match this newly created object model. In contrast to other state-of-the-art approaches [33], [2], we do not assume any prior information about the environment, local relation between several segments, their semantics, or rely on any prior object models. However, it is possible to incorporate prior knowledge (e.g. semantic

segmentation) to support the matching process and improve speed and accuracy.

We apply a three-step registration process that allows for global registration of two point clouds based on their geometry. First, we conduct an efficient nearest neighbor search in descriptor space using a kd-tree to obtain matching 3D descriptors. For every keypoint we select the best k_{NN} matches with a descriptor similarity score above t_{sim} .

Secondly, we use RANSAC to find a geometrically consistent set of inlier matches as well as the initial transformation $T_{i,j}^{coarse}$ between models s_i and s_j . If no consistent set of feature matches is found, the candidate is rejected.

In the refinement step, Point-to-Plane ICP is applied to the coarsely registered models, yielding an improved registration transformation $T_{i,j}^{fine}$. We apply a threshold e_{ICP}^* to the RMSE of ICP to reject unsuccessful registrations. The final transformation between the models s_i and s_j is then obtained as follows:

$$T_{i,j} = T_{i,j}^{fine} T_{i,j}^{coarse} . \quad (1)$$

Fig. 3 shows two examples of the descriptor matching and the intermediate states of the model registration process.

4) *Merging and Verification*: After obtaining a registration for an object model candidate, we verify the match using a two-stage geometric consistency check, only if it succeeds, we merge model s_i into model s_j . We perform these verification steps based on the models' TSDF grids t_i and t_j . First, we transform the TSDF grid t_i of model s_i into the TSDF grid t_j of model s_j using trilinear interpolation

$$t'_i = \Xi(T_{j,i} \cdot t_i) . \quad (2)$$

The first model to model verification step is performed by counting the overlapping voxels (o_{TSDF}) and calculating the RMSE of all voxel pairs (e_{TSDF}) of the aligned voxel grids. We enforce a maximum RMSE (e_{TSDF}^*) and a minimum overlap (o_{TSDF}^*) between t'_i and t_j :

$$e_{TSDF}^* > \sqrt{(t_j - t'_i)^2} , \quad (3)$$

$$o_{TSDF}^* < |t_j \cap t'_i| . \quad (4)$$

If the match is accepted, the aligned grids are merged by taking the weighted average of each voxel pair:

$$\hat{t}_j = t_j \oplus t'_i . \quad (5)$$

In a final model to scene verification step, the merged TSDF grid $\hat{t}_j$, is transformed into all observed locations in the GSM and their geometric consistency with the scene TSDF t_{scene} is verified based on the same metrics as above. Both verification steps are depicted in Fig. 4. The whole process of inserting a new segment into the database is furthermore outlined in Algorithm 1.

IV. EXPERIMENTS

To evaluate our system, we perform experiments on sequence 66 of the SceneNN dataset [34], where multiple identical objects are present, shown in Fig. 5, and on four

Algorithm 1 Incremental 3D object database

```

1: procedure match_and_register()
2:    $T_{i,j}^{coarse} \leftarrow \text{matching\_and\_RANSAC}(k_i, d_j, k_j, d_j)$ 
3:    $T_{i,j}, e_{ICP} \leftarrow \text{ICP}(c_j, T_{i,j}^{coarse} c_i)$ 
4: end procedure
5: procedure insert_segment_in_database( $s_i$ )
6:   candidates  $\leftarrow \emptyset$ 
7:   while  $s_j \leftarrow \text{load\_next\_segment}(\Omega)$  do
8:      $T_{i,j}, e_{ICP} \leftarrow \text{match\_and\_register}(s_i, s_j)$ 
9:     if  $e_{ICP} > e_{ICP}^*$  then
10:      skip  $s_j$ 
11:     end if
12:      $t'_i \leftarrow \Xi(T_{j,i} \cdot t_i)$ 
13:      $e_{TSDF}, o_{TSDF} \leftarrow \text{geom\_consistency}(t'_i, t_j)$ 
14:     if  $e_{TSDF} > e_{TSDF}^* \vee o_{TSDF} < o_{TSDF}^*$  then
15:       skip  $s_j$ 
16:     end if
17:      $\hat{t}_j \leftarrow t_j \oplus t'_i$ 
18:     for all  $T_k \in \mathbb{T}_i \cup \mathbb{T}_j$  do
19:        $\hat{t}'_j \leftarrow \Xi(T_k \cdot \hat{t}_j)$ 
20:        $e_{TSDF}, o_{TSDF} \leftarrow \text{geom\_consistency}(\hat{t}'_j, t_{scene})$ 
21:       if  $e_{TSDF} > e_{TSDF}^* \vee o_{TSDF} < o_{TSDF}^*$  then
22:         skip  $s_j$ 
23:       end if
24:     end for
25:      $\hat{s}_j \leftarrow \text{compute\_object\_model}(\hat{t}_j)$ 
26:     candidates  $\leftarrow (\hat{s}_j, e_{TSDF})$ 
27:   end while
28:    $\hat{s} \leftarrow \text{get\_lowest\_rmse}(\text{candidates})$ 
29:    $\Omega \leftarrow \text{insert}(\hat{s})$ 
30: end procedure

```

Fig. 5: *top*: The output of GSM on the SceneNN [34], scene 066. Different colors indicate different segments. *bottom*: Ground truth segmentation reported by the authors of the dataset.

indoor datasets collected with a Tango phone, and released with this paper², see Fig. 6.

The indoor datasets were recorded with a Lenovo Phab 2 Pro, which is equipped with the Google Tango sensor suite, i.e., an RGB-D sensor and a grayscale fisheye camera for motion tracking. The Google Tango system provides accurate camera pose estimates using a keyframe-based

²The datasets are available at <https://projects.asl.ethz.ch/datasets/>.

Fig. 6: Labeled surface mesh of the GSM for the three large Google Tango dataset scenes. Every scene contains eleven distinguishable objects, some of which appear multiple times, e.g. the round chair.

visual inertial SLAM algorithm, including post-processing with loop closure detection and full bundle-adjustment. The datasets were recorded in a lounge area, with different arrangements of objects, i.e., 17 objects from 11 different object categories (*cow* (1), *round chair* (3), *cube chair* (1), *fire extinguisher* (3), *gas cylinder* (2), *helmet* (2), *round table* (1), *square table* (1), *large box* (1), *small container* (1), *floor* (1)).

Firstly, we show the segmented reconstruction output of the GSM. Secondly, we evaluate the benefits of an incremental object database with its merging capabilities. Finally, we demonstrate the system for single- and multi-session scene completion. We report the most relevant parameters of the system in Table I.

A. Depth Segmentation and GSM

The depth segmentation of a single depth image on the Tango dataset is shown in Fig. 2. After the integration of multiple segmented frames into the GSM, we obtain the initial segmented scene, depicted in Fig. 6. In this scene, multiple instances of the same object can be observed, e.g., multiple chairs of the same kind. From the GSM output, it is visible that some objects are segmented into multiple parts, showing the need to merge segments and observation

TABLE I: Parameters of ISS/Harris3D keypoint detectors, FPFH descriptor, Kd-tree Matching, and TSDF merging.

Harris3D	ISS	FPFH	Matching	Merging
$r_H = 3 \text{ cm}$	$r_b = 2 \text{ cm}$	$r = 10 \text{ cm}$	$t_{sim} = 0.4$	$e_{TSDF}^* = 2 \text{ cm}$
$k_{NN}^{min} = 10$	$r_s = 6 \text{ cm}$		$k_{NN} = 5$	$o_{TSDF}^* = 10^4$
$t_{curv} = 0.07$	$r_{nm} = 8 \text{ cm}$		$e_{ICP}^* = 2 \text{ cm}$	
$t_{resp} = 10^{-6}$	$k_{NN}^{min} = 5$			

TABLE II: Computation times for processing the large three scenes of the Tango dataset: Point cloud extraction (**P**, III-C.1) from the TSDF grid, the Keypoint extraction (**K**, III-C.2), the Descriptor computation (**D**, III-C.2), the Matching & Registration (**MR**, III-C.3), the Merging & model to model Verification (**MV**, III-C.4), and model to Scene Verification (**SV**, III-C.4).

	P	K	D	MR	MV	SV	Σ
Time [s]	11.9	22.4	24.4	1267.6	21.6	24.7	1372.6

of repetitive objects. Similarly, the SceneNN dataset contains multiple instances of one chair and, hence, akin observations can be made on the GSM output, shown in Fig. 5.

The implementation of depth segmentation is currently the bottleneck of the GSM pipeline in terms of processing speed, able to process VGA resolution depth maps at only $\sim 3 \text{ Hz}$. The GSM can integrate segmented depth maps at VGA resolution at $\sim 8 \text{ Hz}$. Hence, for real-time performance, VGA depth maps need to be down-sampled. The Tango devices, on the other hand, provide depth maps at a lower resolution ($224 \times 172 \text{ px}$) at 5 Hz and can be processed in real-time.

B. Incremental Object Database

The main goal of our incremental object database is to improve the quality and completeness of our object models by merging the knowledge of multiple observed object instances. We show this and the reduction of raw segments, extracted from the GSM, by matching and merging them to object models previously inserted into the database if the verification steps indicate equal object instances, in Fig. 7.

From an initial set of 330 raw segments, recorded over three sessions, we recognize 48 raw mergeable segments that are reduced to 27 database object models. Six of these objects consist of two or more merged segments.

In Table II, we report the computational times of the individual steps of the database from point cloud extraction from TSDF grids up to the verification steps among segments and to the scene, computed on a Intel Xeon CPU @ 2.80 GHz (8 cores) running in a single thread. On average we spend 32.0 ms on extracting the point cloud, 95.4 ms on keypoint extraction and 107.3 ms on computing descriptors for a single segment. The timings show that the most time consuming step is the object model matching, more specifically the ICP refinement step. Please note that the timings reported for matching, merging and verification are highly dependent on the dataset and the resulting database.

C. Single- & Multi-Session Scene Completion

Finally, we demonstrate the performance of our system in completing scenes by using detected merged instances

Fig. 7: Extracted GSM raw segments from three Google Tango recordings of a scene that contained eleven distinguishable objects, with one or more instances of each object. From originally 330 raw segments, 48 mergeable (non-planar) raw segments remained (*top*). They are processed by the database and combined into 27 object models, of which 8 (*bottom*) are the result of merging several raw segments: cow (3), round chair (10), fire-extinguisher (4), gas cylinder (4), box(2), cube chair (2). The algorithm produces only a single erroneously merged object model (2nd object bottom row).

of objects in single- and multi-session applications. The difference between these applications is that in the first case, our system starts with an empty database, while the second setup operates with an existing database (from the previous sessions, which gets extended). In Fig. 8, we depict two scenes in which multiple object models were detected and inserted, one from the SceneNN66 and one from the Tango dataset. During inclusion, only models that align well with the TSDF grid of the scene are included. Note, how the object model inclusion succeeds to accurately align with the partial views in the scene (*blue*) and manages to fill gaps of unobserved parts (*red*).

D. Limitations

We have shown that our system can perform well on different datasets, however some challenges still remain:

- *Planar structures and geometrically uninteresting objects*, such as cylinders and cubes are not descriptive

Fig. 8: Inserting the completed object model instances into the reconstructions results in a more complete and accurate model of the scenes. The images show distances of the points from the completed reconstruction to the points on the originally observed scene (*red*: high to *green*: low). The original reconstructions are shown at the top of Fig. 5 and the bottom scene of Fig. 6.

enough when partially observed and are, therefore, very hard to merge. Thus, we exclude planar objects from matching. One potential solution to this problem could be to include the visual appearance into the description.

- Objects with *thin surfaces* are very hard to represent by a TSDF because of the truncation distance. This is an inherent problem of the representation and can be avoided by using a different one, such as point clouds, meshes, etc.
- In the segmentation step, we make that assumption that object can be represented with convex regions, although this is not always the case. In order to deal with *non-convex shapes*, different segmentation techniques could be applied.
- For some object models, *descriptors* and *keypoints* were not descriptive enough, failing to properly match even if the same instance of the object in the scene was observed. An alternative would be to use more descriptive features, such as global descriptors, TSDF-based descriptors or learning based descriptors.

V. CONCLUSION AND FUTURE WORK

In this work, we have presented a novel database system for incremental 3D object model generation. The presented method is capable of automatically discovering new object models and updating, improving, and completing existing

models with new observations. We showed that the knowledge of merged object models can be used to complete scenes and, hence, improve scene reconstructions. No prior information about the objects is required for the entire database generation and updating, facilitating a completely unsupervised object detection scheme. Finally, we evaluated our methodology on one publicly available and on four newly created RGB-D object datasets. The latter are released with this publication.

The presented object database expedite several emerging applications. In robotic navigation, for instance, object based SLAM approaches can make use of this system without prior knowledge or strong assumptions about the objects present in the mapped environment. This gives rise to broader applicability of such mapping systems. Furthermore, the segment matching and registration procedures within such a mapping system may be simultaneously used for detection of loop closures. In future work, we wish to incorporate semantic information to the object models, for a faster matching procedure. We furthermore believe that systems based on other depth sensing modalities, *e.g.*, laser range finders, can benefit from the proposed object database framework in an equal fashion as RGB-D.

ACKNOWLEDGMENT

This work was partially supported by the Swiss National Science Foundation (SNF), within the National Centre of Competence in Research on Digital Fabrication, and the Swiss Commission for Technology and Innovation (CTI).

REFERENCES

- [1] Y. Li, A. Dai, L. Guibas, and M. Nießner, "Database-assisted object retrieval for real-time 3d reconstruction," *Computer Graphics Forum*, vol. 34, no. 2, pp. 435–446, 2015.
- [2] R. Salas-Moreno, R. Newcombe, H. Strasdat, P. Kelly, and A. Davison, "Slam++: Simultaneous localisation and mapping at the level of objects," in *CVPR*, 2013.
- [3] O. Russakovsky, Y. Lin, K. Yu, and L. Fei-Fei, "Object-centric spatial pooling for image classification," in *ECCV*, 2012.
- [4] J. Deng, N. Ding, Y. Jia, A. Frome, K. Murphy, S. Bengio, Y. Li, H. Neven, and H. Adam, "Large-scale object classification using label relation graphs," in *ECCV*, 2014.
- [5] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," *arXiv preprint arXiv:1312.6034*, 2013.
- [6] A. Uekermann, R. Haschke, and H. Ritter, "Realtime 3d segmentation for human-robot interaction," in *IROS*, 2013.
- [7] K. Tateno, F. Tombari, and N. Navab, "Real-time and scalable incremental segmentation on dense slam," in *IROS*, 2015.
- [8] J. Kenney, T. Buckley, and O. Brock, "Interactive segmentation for manipulation in unstructured environments," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. IEEE, 2009, pp. 1377–1382.
- [9] S. Gupta, R. Girshick, P. Arbelaez, and J. Malik, "Learning rich features from RGB-D images for object detection and segmentation," in *ECCV*, 2014.
- [10] A. Karpathy, S. Miller, and F. F. Li, "Object discovery in 3d scenes via shape analysis," in *ICRA*, 2013.
- [11] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, "Mask r-cnn," *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2980–2988, 2017.
- [12] E. Shelhamer, K. Rakelly, J. Hoffman, and T. Darrell, "Clockwork convnets for video semantic segmentation," in *European Conference on Computer Vision*. Springer, 2016, pp. 852–868.
- [13] R. Dubé, M. G. Gollub, H. Sommer, I. Gilitschenski, R. Siegwart, C. Cadena, and J. Nieto, "Incremental-segment-based localization in 3-d point clouds," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1832–1839, 2018.
- [14] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *IJCV*, vol. 59, no. 2, pp. 167–181, 2004.
- [15] M. Fehr, F. Furrer, I. Dryanovski, J. Sturm, I. Gilitschenski, R. Siegwart, and C. Cadena, "TSDF-based change detection for consistent long-term dense reconstruction and dynamic object discovery," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 5237–5244.
- [16] A. Uekermann, C. Elbrechter, R. Haschke, and H. Ritter, "3d scene segmentation for autonomous robot grasping," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 1734–1740.
- [17] A. Singh, J. Sha, K. S. Narayan, T. Achim, and P. Abbeel, "Bigbird: A large-scale 3d database of object instances," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 509–516.
- [18] R. Finman, T. Whelan, M. Kaess, and J. J. Leonard, "Efficient incremental map segmentation in dense rgb-d maps," in *ICRA*, 2014.
- [19] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. ACM, 1996, pp. 303–312.
- [20] K. Tateno, F. Tombari, and N. Navab, "When 2.5 d is not enough: Simultaneous reconstruction, segmentation and recognition on dense slam," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016, pp. 2295–2302.
- [21] A. Dai, C. R. Qi, and M. Nießner, "Shape completion using 3d-encoder-predictor cnns and shape synthesis," in *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2017.
- [22] X. Han, Z. Li, H. Huang, E. Kalogerakis, and Y. Yu, "High-resolution shape completion using deep neural networks for global structure and local geometry inference," *arXiv preprint arXiv:1709.07599*, 2017.
- [23] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision*. Springer, 2016, pp. 21–37.
- [24] J. Redmon and A. Farhadi, "Yolo9000: better, faster, stronger," *arXiv preprint*, vol. 1612, 2016.
- [25] I. Sipiran and B. Bustos, "Harris 3D: A robust extension of the Harris operator for interest point detection on 3D meshes," *Visual Computer*, vol. 27, no. 11, pp. 963–976, 2011.
- [26] Y. Zhong, "Intrinsic shape signatures: A shape descriptor for 3d object recognition," in *ICCV Workshops*, 2009.
- [27] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (fpfh) for 3d registration," in *ICRA*, 2009.
- [28] M. Fischler and R. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *CACM*, vol. 24, no. 6, 1981.
- [29] A. Aldoma, Z.-C. Marton, F. Tombari, W. Wohlkinger, C. Potthast, B. Zeisl, R. B. Rusu, S. Gedikli, and M. Vincze, "Tutorial: Point cloud library: Three-dimensional object recognition and 6 dof pose estimation," *IEEE Robotics & Automation Magazine*, vol. 19, no. 3, pp. 80–91, 2012.
- [30] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [31] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," in *ACM siggraph computer graphics*, vol. 21, no. 4. ACM, 1987, pp. 163–169.
- [32] F. Tombari, S. Salti, and L. DiStefano, "Performance evaluation of 3d keypoint detectors," *International Journal of Computer Vision*, vol. 102, no. 1, pp. 198–220, Mar 2013.
- [33] O. Mattausch, D. Panozzo, C. Mura, O. Sorkine-Hornung, and R. Pajarola, "Object detection and classification from large-scale cluttered indoor scans," *Computer Graphics Forum*, vol. 33, no. 2, pp. 11–21, 2014.
- [34] B.-S. Hua, Q.-H. Pham, D. T. Nguyen, M.-K. Tran, L.-F. Yu, and S.-K. Yeung, "Scenenn: A scene meshes dataset with annotations," in *International Conference on 3D Vision (3DV)*, 2016.