DR 1.3: Sensing, mapping and low-level memory III

Tomáš Svoboda[*], Abel Gawel[†], Hartmut Surmann[‡], Rainer Worst[‡], Renaud Dubé[†], Cesar Cadena[†], . . . , and the TRADR consortium

[*]*CTU in Prague, Czech Republic*
[†]*ETH, Zürich, Switzerland*
[‡]*Fraunhofer IAIS, Sankt Augustin, Germany*
⟨svobodat@fel.cvut.cz⟩

Project, project Id:	EU FP7 TRADR / ICT-60963
Project start date:	Nov 1 2013 (50 months)
Due date of deliverable:	Month 38
Actual submission date:	
Lead partner:	CTU
Revision:	draft
Dissemination level:	PU

We report progress achieved in Year 3 of the TRADR project in WP1: *Persistent models for perception*. It describes the essential robot (UGV) perception functionalities and and new algorithms for realtime 3D mapping with loop closure and merging UGV and UAV maps.

Executive Summary

The key objective of WP1 is to provide sensory data from all involved robots registered in space and time, to keep creating and updating robot centric representations, and ground them into the world coordinate frame. The obtained representations are furnished to other WPs, which maintain higher level situation awareness.

During Year 3 we concentrated on upgrading the 3D mapping algorithm with the crucial ability of finding loop closures and of optimizing the full pose-graph in real-time. We developed two methods for merging UGV maps (lidar) and UAV maps (from vision). One uses sparse vision maps and aims at a very general solution for global localization. The second one is more problem specific it does dense visual reconstruction and exploits a relative pose prior. We extended the adaptive traversal to continuous flipper control asymmetric for front and rear flippers. A new active search algorithm was developed for 3D scene labeling. We designed a bi-directional connection between high and low-level databases for managing multiple object detections.

Role of robot perception and metrical mapping in TRADR

The robot perception means the robot is able to analyze its neighborhood and act accordingly. Terrain recognition is essential for robot locomotion regardless whether the robot is teleoperated or moves autonomously. It is desirable the robot overcomes obstacles in a reasonable way - fast, safe, consuming less power and reducing cognitive load of a human operator. Automatic victim detection is important for many search and rescue scenarios. A human operator may provide final decision however, robots, when crawling through a disaster site should provide warning about possible victim locations.

The metrical mapping serves as the very basis for modeling the world. It is also the basis for sharing information between robots and also among several sorties and even missions.

Persistence

Persistence in WP1 is addressed mainly by re-using the data in creating an enviroment model. The 3D metric map serves as the main basis for multimodal (data), multi-source (robots), multi-level (abstraction, decisions) registration. In WP1, we are working on robust methods for merging partial 3D maps. The merging challenges include weak data overlap, dynamic changes in scenes, large displacement of local coordinate systems. The terrain perception and robot control algorithms use machine learning techniques in a quest of gaining experience from operator-robot interactions.

Contribution to the TRADR scenarios and prototypes

The new 3D mapping algoritm is an important step towards multi-robot collaboration (WP4) and models for acting (WP2), The algorithms for sensing in sensory deprived environments together with related hardware advancements contribute to the models for acting (WP2), multi-robot collaboration (WP4) and also support the human-robot teaming (WP5).

1 Tasks, objectives, results

1.1 Planned work

In Year3, WP1 planned to investigate "Sensing, mapping and low-level memory III – Multi-robot perception" (Milestone MS1.3). The work was divided into two tasks:

- Essential sensing and UGV control functionality (T1.5)
- Robot centric metrical maps and models storage (T1.6)

Both tasks emphasize multi-robot sensory data and also closing feedback loops from semantic higher level and alos from the world-as-a-memory representation

1.2 Addressing reviewers' comments

Comment:

> The Continuous Trajectory Scan Matching mapping framework was demonstrated to be able to merge offline laser maps obtained in different sorties, nevertheless doing it online is still not possible, in spite of being a fundamental feature required by the TRADR system in operational scenarios involving multiple robots.

Response: We developed a tool to interactively align maps of previous sorties that were merged offline. This summary map can then be used online and extended by the UGVs. Furthermore, this enables the UGVs to do both path-planning and multi-robot patrolling on these maps.

Comment:

> Moreover, the innovation beyond the SoA of the SLAM framework was not demonstrated and it was hard to understand in the demonstration why the proposed framework was presented in parallel with an ICP-based map merging code running on a separate computer. Essentially, a clear evaluation of the framework advantages and disadvantages with respect to the SoA was missing from the presentation and demonstration.

Response:

1. The proposed pose-graph SLAM framework [15] was fully integrated in the TRADR system, running on all UGVs.

2. Several iterations for 3D laser-based map merging based on the *OctoMap* representation have been investigated, with increasing complexity, i.e., static map merging with known locations, static map merging with unknown locations and manual alignment, dynamic map merging using *SegMatch* [16].

3. Laser-based loop closure and multi-robot registration has been implemented with the *SegMatch* technique. The previous static ICP SLAM framework is not able to propagate any retrospective updates to the map, rendering the drift correction imposed by loop closure impossible.

Comment:

> The deliverable D1.2 mentions that the consortium is still in the process of developing change detection techniques, therefore the work progress in the WP is slightly behind schedule ...

Response:

1. A dynamic laser-scan filtering has been adopted using the *OctoMap* framework to filter dynamics during a mission and allow robot navigation under dynamic conditions.

2. A change detection algorithm was developed to identify and segment changes between missions and highlight points of interest.

3. We report this progress under WP2, T2.6.

Comment:

> The consortium should review carefully recent work on navigation in smoky environments in the scope of other EU-funded projects (e.g. SmokeBot, http://aass.oru.se/Research/mro/smokebot/)

Response: Research areas of SmokeBot and TRADR partially overlap; for us, interesting results come from their interest in thermal images, especially the effects of cold and hot smoke on the images. Moreover in localization, utilization of radar fused with lidar seems promising, at least for coarse mapping and localization in smoky areas. Unfortunately, their radar setup is specific to their platform and would be difficult to install to the TRADR platform. Nevertheless, we keep in mind these alternative approaches to localization. Finally, the gas sensors the SmokeBot project develops are a clear choice for TRADR scenario, if they become available.

Comment:

> How the traversability information will be integrated into the system, as promised in the presentation/demo, should be clarified. We additionally look forward to seeing non-symmetric flipper control and clear evaluation of the victim detection technology already developed in Y1 and Y2
>
> The use of the flippers to assist in traversal is commended and should continue to develop. The proposed development directions towards continuous, asymmetric control look promising.

Response:

1. In Year3 we concentrated on flipper continous control where front and rear flippers are controlled independently (assymetrically) [56] (Annex Overview 2.4). This is for both from uncomplete data as for the blind mode exploration (detailed in WP2 - Deli2.3).
2. Victim (object) detection were extended in two ways. We advanced toward active search where the result is a complete 3D scene segmentation and an optimal sensor control [58] (Annex Overview 2.7). We also closed the gap between low and high level TRADR databases by reasoning over multiple detections [37] (Annex Overview 2.6).
3. Integrating the traversability information into the multi-robot planning is still ongoing.

1.3 Actual work performed

1.3.1 Task T1.5 (Essential sensing and UGV control functionality III – Multi-robot perception

Terrain perception for robot control

We extended our work on autonomous flipper control [57] (Annex Overview 2.3) towards continous tilting and effectivity of training [56] (Annex Overview 2.4). Last years, we discretized the space of flipper morphology into 5 hand-cratfed configurations and autonomous traversal algorithm selected (by learning) the proper configuration from both interoceptive and exteroceptive measurements. The discretization eased the learning, effectively reducing number of real robot roll-outs. On the other hand, the discretization inevitably led to a suboptimal control and frequent switching between configurations slowed down the traversal. In Year3, we extended the approach towards continuous control of the front and rear flippers.

Policy Gradient methods require many real-world trials. Some of the trials may endanger the robot system and cause its rapid wear. Therefore, a

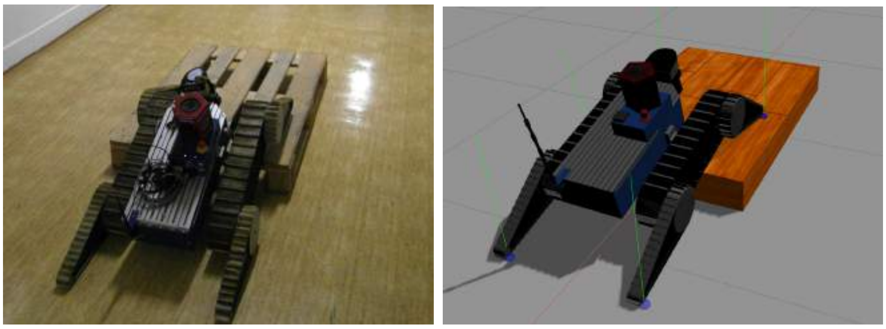

Figure 1: Left: Real robot executing a safe policy. Right: Visualization of the simulated robot.

safe or at least gentle-to-wear exploration is a desired property. We incorporate bounds on the probability of unwanted trials into the recent Contextual Relative Entropy Policy Search method. The task of Reinforcement Learning (RL) [35] is to search through the space of policies $\pi : S \to A$, which map agent states S to possible actions A; the actions are then applied either in reality or using a transition model, and the agent reaches a new state (this description is usually known as Markov Decision Process, MDP).

Model-free policy search algorithms usually follow these steps: (i) generate trajectories from the real-world system, (ii) compute a policy maximizing the expected sum of rewards on the so-far-generated trajectories, (iii) use the policy to generate a new real-world trajectory, (iv) repeat from (ii). Contextual Relative Entropy Policy Seatch (REPS) [40] adds a task-dependent context. We extend Contextual REPS with additional constraints. In particular, for systems which are not inherently safe, or which are prone to wear, we use a cautious physics-based simulator to determine a rollout safety.

States of the task are: (i) robot body pitch, and (ii) height of the terrain approximately 20 cm in front of the robot body (read from an octomap built online from laser scans).

A policy controls independently the pairs of front and rear flippers using positional control. Therefore, the *action* space is continuous and 2-dimensional.

In the previous adaptive traversal (AT), *safety* is not modeled separately, and some of the safety features are part of the reward. The *reward* for the AT task is a weighted sum of (i) manually assigned safety penalty, (ii) high pitch/roll angle penalty, (iii) penalty for excessive flipper motion, (iv) robot forward speed reward, and (v) motion roughness penalty measured by accelerometers [76]. In the safe traversal (ST) task, the reward is simply the distance traveled in 30 seconds over the pallet (the choice of policy influences e.g. track slippage and motor stress, which lower the speed). Safety is modeled explicitly by the cautious simulator, which marks as unsafe all rollouts in which the robot tops over, hits hard on the ground or obstacle (measured as deceleration), or hits objects with delicate parts of its body (e.g. sensors).

We use a *policy* that is linear in the states, and that controls front and rear flippers separately.

To run the *simulation* on a computer, we use the Open Dynamics Engine (ODE) which aims at fast approximate simulation. In ODE, we use a simple, yet reasonably plausible way to simulate the mobile robot with non-deformable tracks – a method typically used for simulating conveyor belts. The collision model used in the simulator consists of simple-shape collision links (boxes, cylinders) approximating the CAD model of the robot. Weights, centers of mass, inertias, friction coefficients and other dynamics coefficients are estimated manually. It is important to estimate these parameters precisely enough, so that the simulator can be assumed cautious (which can be easily done using the CAD model). Or, in case of very influential unknown parameters, the simulator has to be run with multiple possible values and the worst-case outcome treated as the simulation result.

We also continued our work on blind terrain inspection by exploring the free space by the robotic arm equipped with a newly developed handle for tactile sensing, see Figure 2. The exploration algorithm is described in Deliverable 2.3 (WP2) in a more detail. The tactile sensing by arm shall support terrain perception and robot control in both standard and blind mode. The arm can touch the terrain part wich is in front of and below the robot level and is thus invisible for lidar and camera sensing.

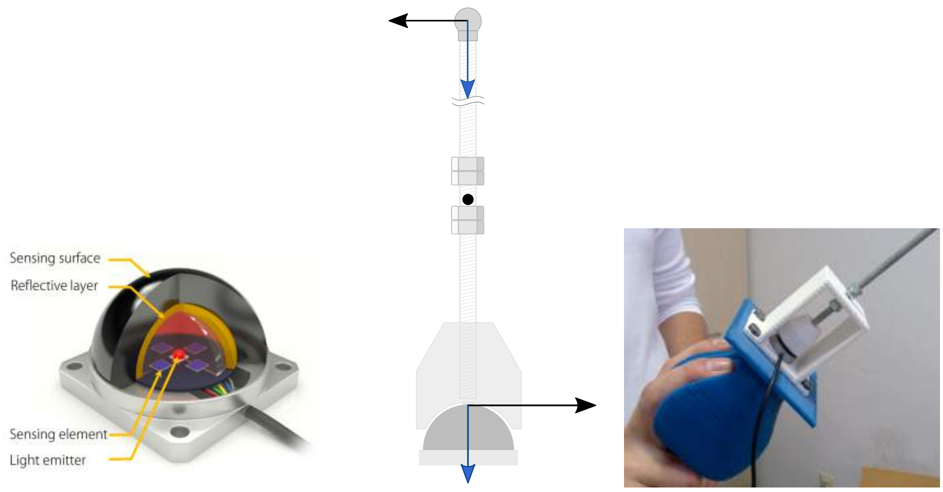

Figure 2: A device for free space exploration. The opto-force sensor (left) is mounted at the end of the stick and the construction allows to measure both lateral (black arrows) and axial (blue arrows) forces, see the scheme at the middle. The blue handle is 3D printed and specially designed for the Jaco arm.

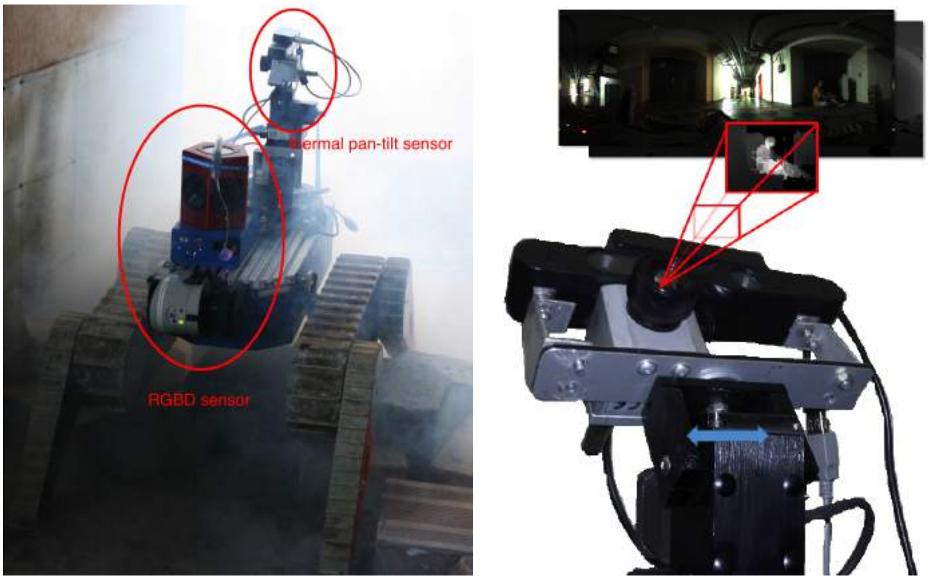

Figure 3: Skid-steer Search&Rescue robot with (i) a panoramic `RGBD` sensor consisting of an omnidirectional camera and a rotation laser scanner, (ii) a narrow-FOV thermal (`T`) sensor, mounted on a pan-tilt unit.

Object detection and scene analysis

We ehanced our work on victim detection towards a more general active 3D-scene labeling algorithm [58] (Annex Overview 2.7). We closed the loop between high and low level database by reasoning over multiple victim detections [37] (Annex Overview 2.6).

We consider the problem of pan-tilt sensor control for active segmentation of incomplete multi-modal data. Since demanding optimal control does not allow for online replanning, we rather employ the optimal planner offline to provide guiding samples for learning of a CNN-based control policy in a guided Q-learning framework. The proposed policy initialization and guided Q-learning avoids poor local optima and yields reasonable results from hundreds of roll-outs.

To make further reading easier, we summarize the algorithm now, using the simplified notation `input`→`output`, which denotes deep convolutional network with defined `input` and `output` layers. Given the human/background annotated dataset, we learn segmentation and control networks:
Segmentation networks: Use annotated dataset to train multi-modal segmentation networks `RGBD`→`H` and `RGBDT`→`H` from pretrained Long's network `RGB`→`H`. Here `V` stands for the layer with the same resolution as the input panoramic data, which assign confidence for each particular pixel to be a human `H`. **Control network:** Learn Q-value network `RGBDT`→`Q`, which maps captured `RGBD`-data directly on Q-values connected with discrete control sig-

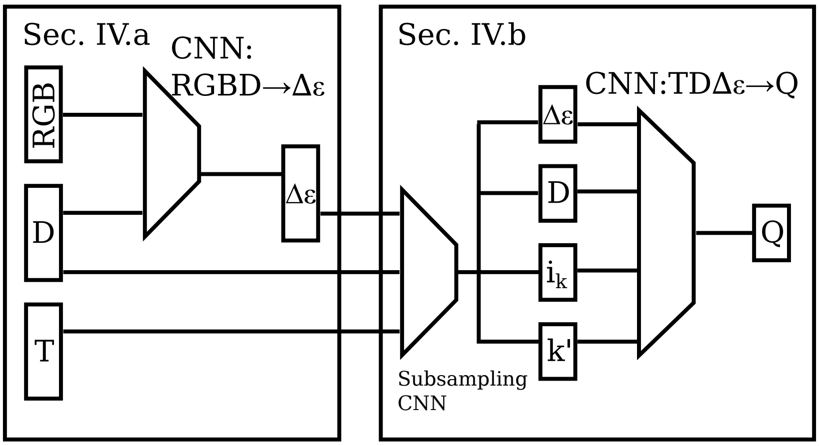

Figure 4: The `RGBDT`→`Q` network is divided into two sub-networks: (i) the `RGBD`→$\Delta\epsilon$ part predicts $\Delta\epsilon$ from `RGBD` data, and (ii) the `TD`$\Delta\epsilon$→`Q` part predicts the Q-values from the depth `D` and the reduction in the classification error $\Delta\epsilon$.

nals $u \in \mathcal{U}$. To avoid learning the huge Q-value network from the scratch, we first use previously learned segmentation networks for its initialization.

When all networks are available, we use them online in the following algorithm: Examples of 3D segmentation results are depicted on Figure 5.

1: Capture `RGBD` data from panoramic sensor.
2: Capture and accumulate `T` data from the current thermal camera viewpoint i.
3: Use the `RGBD`→`H` and `RGBDT`→`H` networks to obtain pixel-wise human confidence.
4: Project and accumulate confidence into corresponding voxels.
5: Use the `RGBDT`→`Q` network to estimate new control signal u for the thermal camera.
6: Move simultaneously: the robot towards the next position in the exploration path and the thermal camera towards the next viewpoint.
7: Repeat from beginning

Algorithm 1: The active segmentation.

We implemented a semantic connection between low-level and high-level TRADR database [37] (Annex Overview 2.6). A 2D visual detector [61] crawls the low-level database for new images and detect object (victims in this case). All available information, including robot position, time stamps, is included in the reasoning about the identity of the detected object. Two approaches were experimentally compared i) similarity based on image descriptors, ii) geometrical verification - essentially, only one object can be on one 3D position at the same time. A victim detection and association example is depicted on Figure 6. Currently, we are extending the semantic connection for a 2D fire detector.

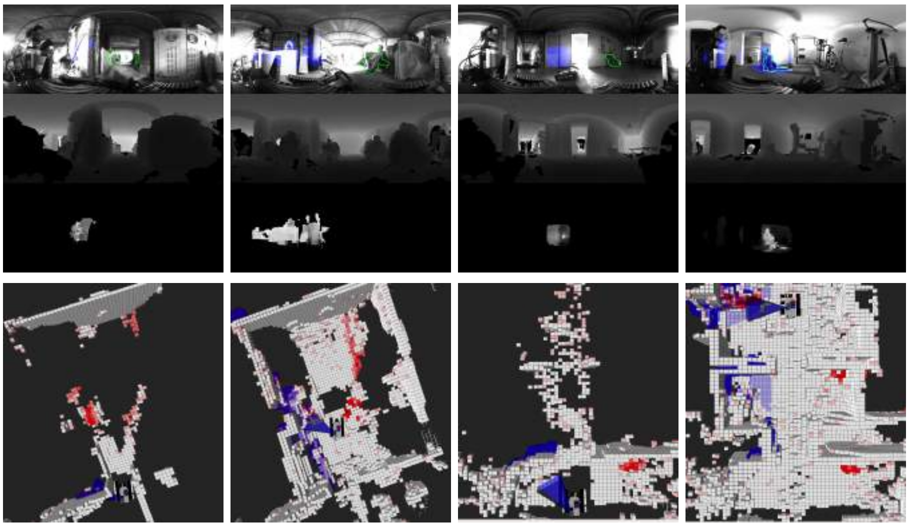

Figure 5: Panoramic images (first row) and corresponding voxels maps (second row) from two experiments with the mobile Search&Rescue platform. First row consists of three panoramic images: (i) grayscale image (generated from `RGB` image) with segmentation of humans outlined by green borders and accumulated temperature rendered from the voxel map emphasized by blue color, (ii) depth `D` image and (iii) thermal `T` image both rendered from the voxel map. Second row shows successively built voxel maps, with robots position and orientation, color of voxels corresponds to the accumulated segmentation confidence (red-human, white-background).

UGV-Mapping

Building on top of the progress the TRADR consortium has made up to the second year, we have extended for the third year the functionalities of the UGV mapping system with the crucial ability of finding loop closures and of optimizing the full pose-graph in real-time.

Current strategies for detecting loop-closures in 3D laser data are primarily based on local keypoint detection and matching, a few others are based on global scene descriptors. The state-of-the-art approaches have been revised and analyzed in [16] (Annex Overview 2.1). Given the drawbacks of the local and global descriptors, especially in the search-and-rescue scenarios where the TRADR system works, we have developed a novel technique that works at the middle ground level without the assumption of being able to detect full objects.

a) Victim 3 - Detection 1

b) Victim 3 - Detection 2

c) Victim 3 - Detection 3

d) Victim 4 - Detection 1

Figure 6: Victim 3 correctly associated in three different observations.

The technique closes loops by matching segments that belong to partial or full objects, or to parts of larger structures (windows, arcs, façades). Our system first extracts and describes segments from a 3D point cloud, matches them to segments from already visited places and uses a geometric-verification step to propose loop-closures candidates. The full pipeline is depicted in Fig. 8 and described in full detail in [16] (Annex Overview 2.1).

This segment-based technique is able to considerably compress the point cloud into a set of distinct and discriminative elements for loop-closure detection. We show that this not only reduces the time needed for matching, but also decreases the likelihood of obtaining false matches. An example of a loop

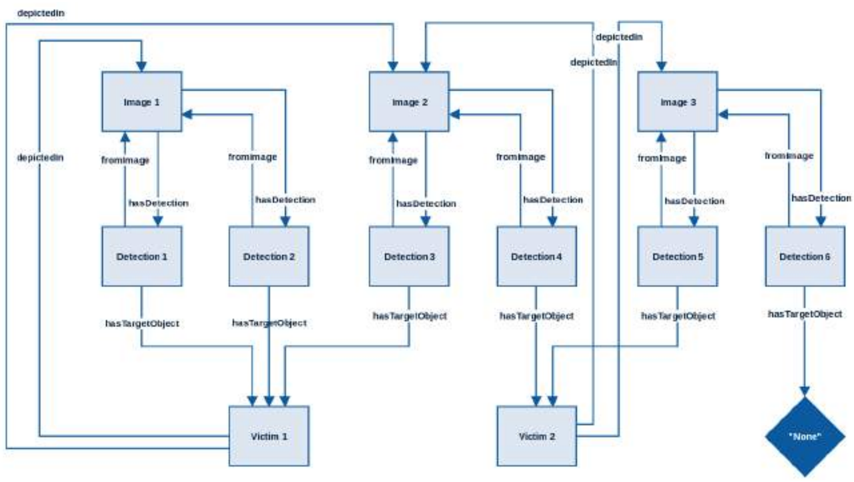

Figure 7: Associations between multiple detections and multiple observations.

closure using our segment based approach is depicted in Fig. 9. The code has been open-sourced at `https://github.com/ethz-asl/segmatch` along with the laser SLAM back-end at `https://github.com/ethz-asl/laser_slam`.

UGV-UAV Map Merging

The fleet of robots in TRADR is composed of UGVs and UAVs. These robots are equipped with different sensing modalities due to their inherent pay-load constraints, e.g., the UGVs can carry on heavy 3D laser scans and the UAVs cannot, instead the UAVs are equipped with RGB cameras. While each robot can build a metric map with their specific sensors, the multi-modal nature of the sensing poses challenges when merging these maps due to the differences in point-of-view (ground VS. aerial) and density of the metric representation.

An evaluation on the state-of-the-art methods for merging maps obtained from the same modality, e.g. laser-laser or vision-vision, has been carried out on our maps from different modalities. This evaluation motivated the development of a technique for merging vision maps and laser maps in the absence of further prior registration, see [29] (Annex Overview 2.2).

The researched technique is based on finding similar structural features in vision and laser maps and reporting back the transformation between matching map elements. The full processing pipeline is depicted in Fig. 10 and described in full detail in [29] (Annex Overview 2.2). Metric maps are independently built from laser sensory data and vision data. The laser maps are constructed using the technique described in [TRADR Year2 Deliverable 2.2, general description of the mapper]. Vision maps are constructed us-

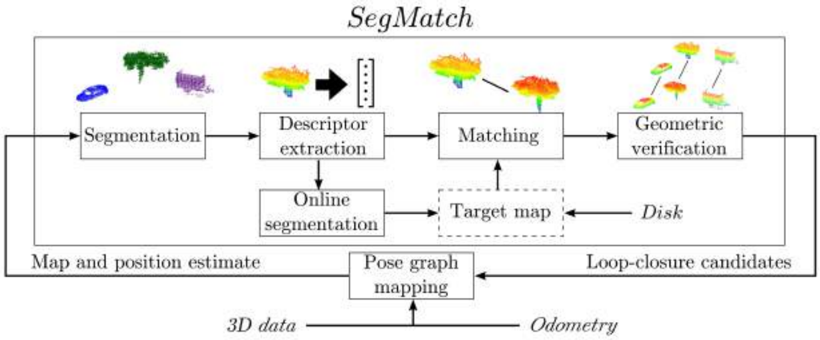

Figure 8: Block diagram of SegMatch, a modular loop detection algorithm. The target map can either be loaded from disk (for localization) or computed online (for loop-closure).

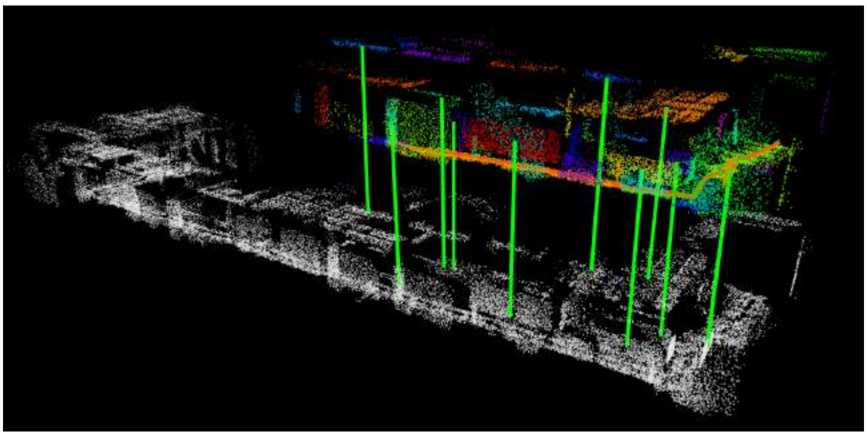

Figure 9: An illustration of the presented loop-closure framework on 3D laser data collected during the TRADR Evaluation 2016 at the Gustav Knepper powerplant. The reference point cloud is shown below (in white), and the local point cloud is aligned above using the loop-closure information. Colours are used to show the point cloud segmentation, and segment matches are indicated with green lines.

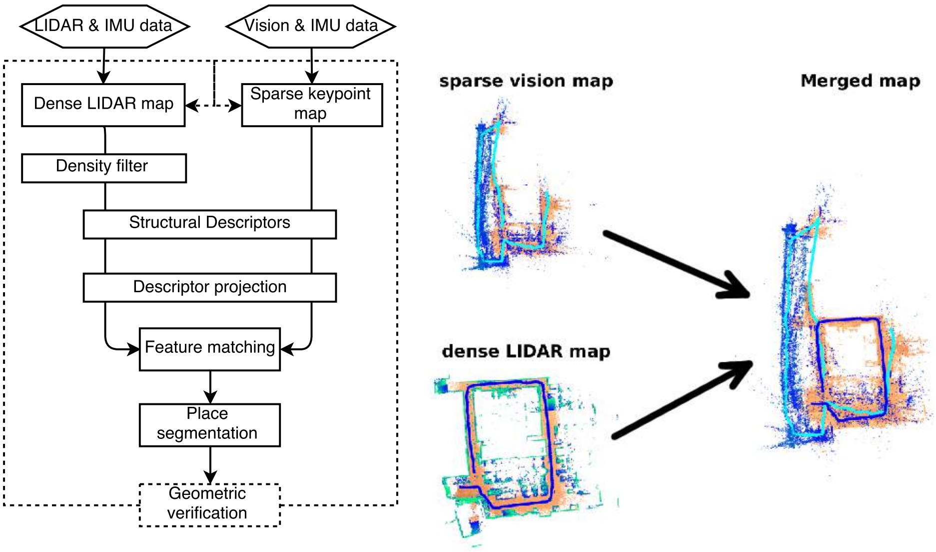

Figure 10: Left: System diagram of the approach for matching and fusing vision-laser maps. Right: The figure illustrates two different point-cloud maps obtained with vision and laser and the final alignment resulting from our approach.

ing a monocular SLAM system [47], facilitating a similar pose-graph-based pipeline.

The two point clouds are of different density; the laser-map is denser than the vision-map. Then, a density filter is applied to the laser point cloud to approximate the sparsity of the visual keypoint map. Having two point-cloud candidates of equal densities, structural descriptors are calculated on a subset of the points. Structural descriptors capture the local neighbourhood around a chosen keypoint using a combination of a binning shape and a descriptor paradigm. Simple occupancy-based and binary density-comparing descriptors offer the best performances amongst all evaluated descriptors.

In a next step the descriptor dimensionality is drastically reduced by performing a projection to the most expressive dimensions. This can be achieved by assessing artificial noisy matches between nearby keypoints and observing the robust descriptor dimensions. This functionality was inspired by [6]. A reduced descriptor dimensionality significantly accelerates the following feature matching step.

A subsequent k-nearest-neighbour search between the two keypoint-feature maps results in the matching matrix. The densest regions of the matching matrix are then segmented into place matching candidates using the approach by [44].

Before finally feeding the place matching candidates for map merging back into the map, a geometric verification is performed. The verification step assesses the individual feature matching candidates of a place match for a common transformation between the places.

Ultimately, the place matching candidates are added as constraints between the vision and laser map and can be optimized within the pose-graph framework. The result of such approach can be visualized in Fig. 10 (Right).

1.3.2 Task T1.6 (Robot centric metrical maps and models storage III – Multi-robot models grounding: into one frame, into one representation)

Vision-Lidar point cloud registration

We developed a method for vision-laser point cloud registration [4] (Annex Overview 2.8). The method allows the localization of the UGV in a vision based map. The approach uses surfaces, that abstract from the underlying data structure and therefore can compensate for minor disturbances while still containing sufficient information for the motion estimation. The resulting map combines the information from both sensors and thus has a higher information content. Because the UAV operates from the air, it can collect data faster than the UGV and create a map of the environment in advance. For this purpose the collected data of the UAV must be processed by an structure from motion method independent of the UGV. Subsequently, the results of the processing can be provided to the UGV for a first localization.

The global map is provided by the UAV, which records the environment during a first flight over the environment and generates a point cloud by means of a vision-based SLAM-algorithm. If the absolute pose of the UGV is known in the global map, the map can be extended by the information of the laser scan and a more detailed map can be built step by step. This is useful, on the one hand, in low-textur regions, which can not be covered by most camera-based methods. On the other hand, map areas such as interiors which are not accessible to the UAV or which are not visible in the event of a flyover due to occlusions can also be included in the global map. All processing steps involved are explained below. For an overview of the whole process see Figure 11.

Merging of metrical maps

The progress done on online loop-closure using the technique described in [16] (Annex Overview 2.1) allows to handle several metric maps from the UGVs to be merged in a single and consistent metric map, see Fig. 13 for an example.

Due to the efficiency and low memory consumption of the proposed segment-based loop closure, the framework can handle multiple sorties on the same environment, to improve and complete the metric model. As the TRADR laser SLAM framework uses the iSAM2 [36] algorithm for optimizing the pose-graph, it can efficiently handle several robots simultaneously

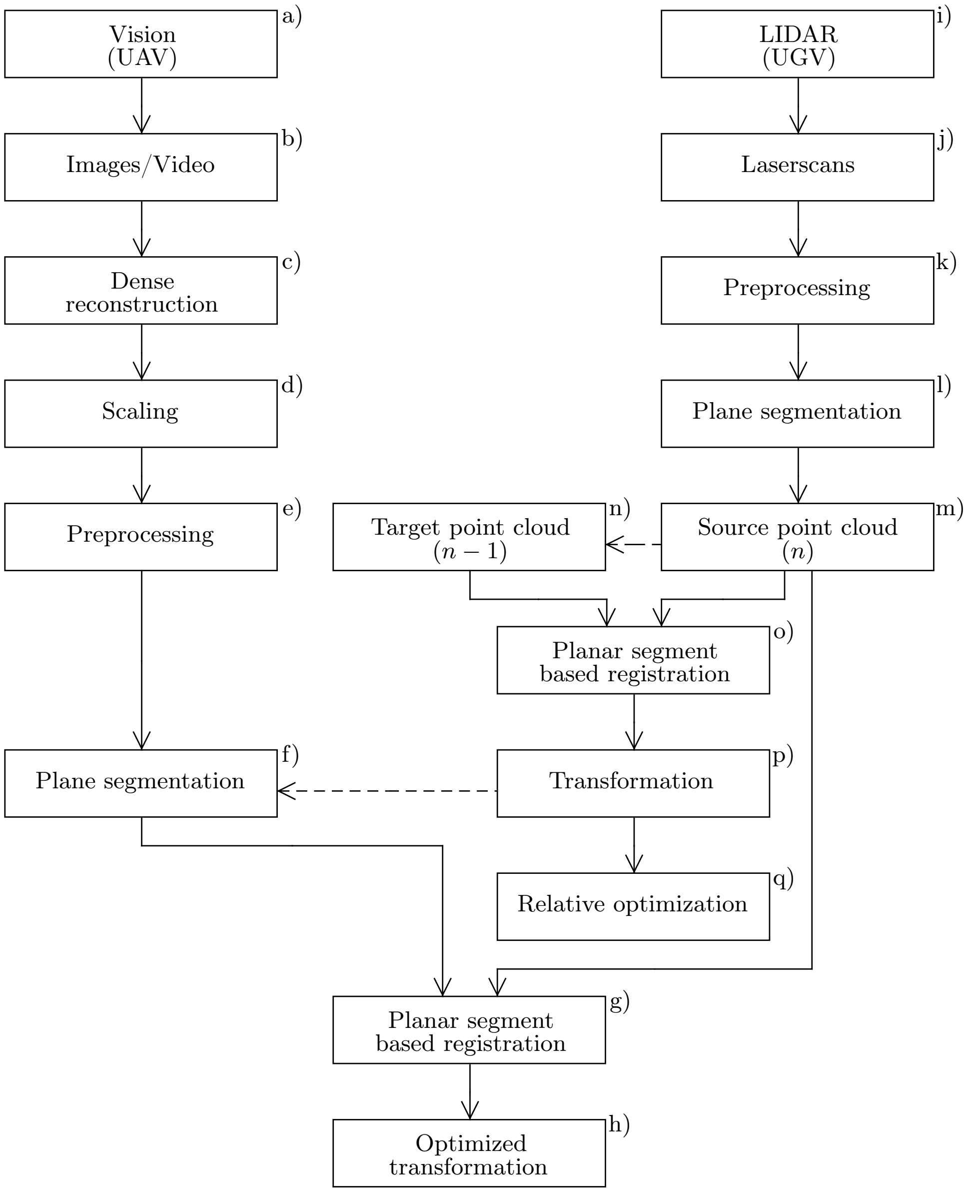

Figure 11: Localization pipeline for laser point clouds with following global optimization step.

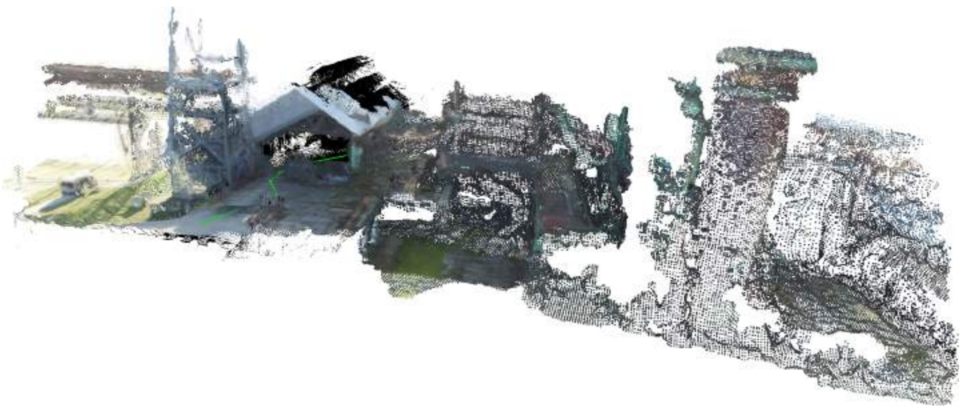

Figure 12: Registered point clouds of the globally optimized localization. Points of the laser point clouds were dyed with the color information of the global point cloud. Within the factory no color information could be extracted. The estimated trajectory is shown in green.

and in real-time by linearizing and recomputing pose variables only when required.

1.4 Relation to the state-of-the-art

1.4.1 Policy search for flipper control

Policy gradient (PG) methods usually require many trials which endanger the real system or cause its excessive wear. Therefore, they are usually not used directly on the real system, but on data-driven models. For example, Kupcsik et al. [40] demonstrate data-driven PG learning of the ball throwing problem with a robotic arm, and Tedrake et al. [67] argues that Policy Gradient learning for aerial maneuvers with an ornithopter may be very efficient, in fact. Transeth et al. [69] show that for snake-like robots with significant side-slip, no closed form expression of the snake's motion exists, therefore policy learning must resort to simulation.

Contextual REPS uses a stochastic upper-level policy which generates deterministic lower-level policy samples. The performance of these policies is evaluated by executing them in the real world, and is used to estimate the upper-level policy gradient. The Gaussian Process REPS (GPREPS) method [40] adds a Gaussian Process (GP) in the loop, which learns a representation of the system dynamics. The GP is used for better evaluation of the policies without the need for executing more real-world samples. Several PG methods also take constraints into account: Uchibe and Doya [1] propose constrained policy search for GPOMDPs [3]. However, GPOMDPs belong to early PG algorithms which use the likelihood-ratio trick to compute the gradient of the expected sum of rewards and then update the policy param-

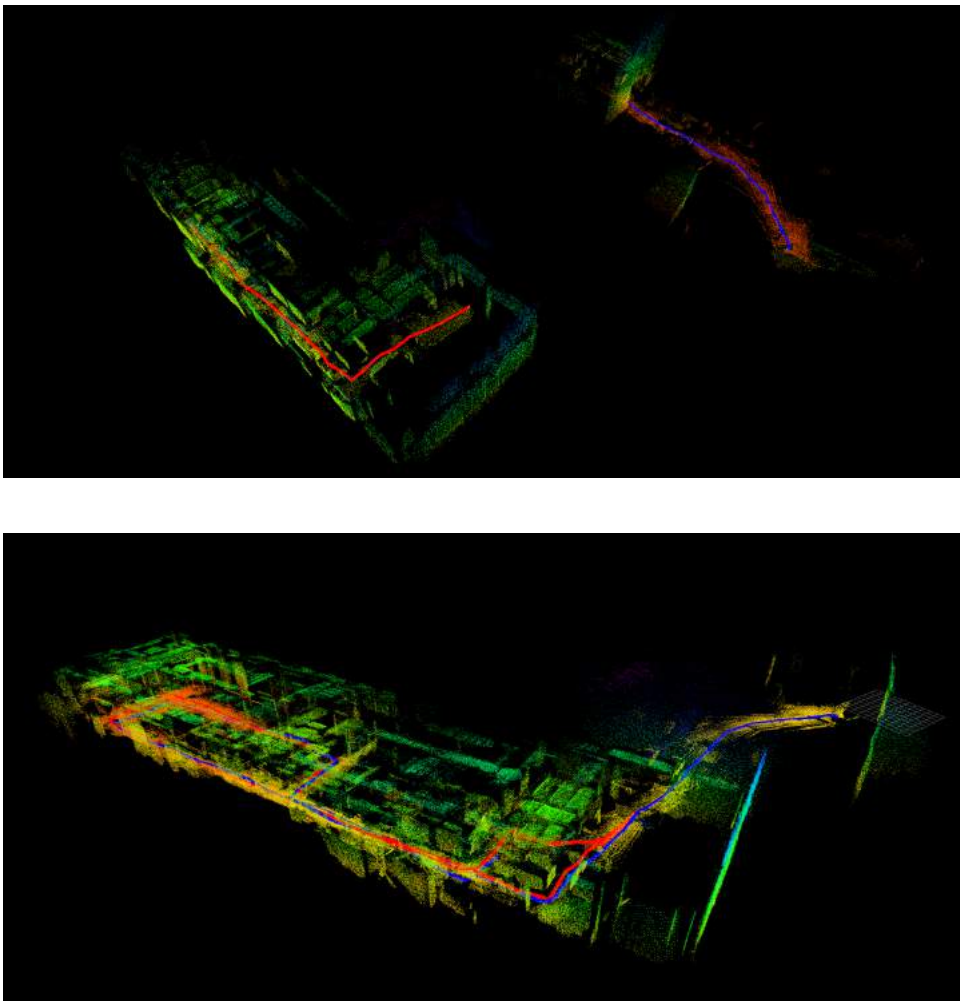

Figure 13: The top figure illustrates the beginning of the multi-robot mapping process with unknown relationship between the two UGVs trajectories. Once robot-associations are made, the metric maps are fused and the trajectories are optimized accordingly, as shown in the bottom figure. This data was collected during the TRADR Evaluation 2016 at the Gustav Knepper powerplant.

eters by a user-defined learning rate. Prashanth [2] propose constrained PG method for Stochastic Shortest Path problem with inequality constraints on Conditional Value-at-Risk (CVaR) as a risk measure. This method does not allow to include implicit constraints and cannot be easily extended for general episodic rewards, such as minimum distance of the trajectory from a target position. We propose a combination of GPREPS with the work of Uchibe and Doya – to extend Contextual REPS with constraints. The proposed method is called *Constrained REPS*, (CREPS). It evaluates the generated policies in a simulator and successively constrains the upper-level policy distribution. This (i) reduces the number of needed samples/iterations and consequently speeds-up the learning process of the model, and (ii) provides a safe policy when used with the real system.

1.4.2 Active scene segmenation

Convolutional Neural Networks (CNNs) have recently been shown to be powerful representation for both classification [39] and control [41]. However, the success of CNNs is usually conditioned either by (i) a large number of labeled training examples [39, 46, 42], or (ii) a careful initialization [43, 41]. We show that in contrast to a general reinforcement learning task, the structure of simultaneous exploration and segmentation with incomplete data (SES) allows for efficient policy initialization.

In particular, we first extend Long's segmentation CNN [43] by depth and depth+thermal modalities and retrain it on our own human/background annotated `RGBDT`-dataset. These segmentation CNNs are further used for self-supervised training of a control sub-network (on a not annotated `RGBDT`-dataset), which estimates potential impact of thermal measurements on the classification error. The control sub-network is further extended by subsampling layers and fully connected layers and trained to predict long-term impact of possible thermal-camera motions on the classification error. To train the control CNN efficiently, we propose a guided Q-learning algorithm, which makes use of optimal trajectories estimated by the Mixed Integer Linear Programming (MILP) planner to guide the exploration of the Q-learning and consequently avoids poor local optima. We show how pretrained segmentation network [43] can be extended by depth and thermal modalities. We propose guided Q-learning and show that it outperforms non-guided Q-learning of Mnih et al. [46]. We suggest self-supervised policy initialization for instances of SES problems.

1.4.3 Vision based SLAM

In order to determine the visual odometry, only keypoints are selected which make a robust correspondence search possible. While some methods are computing complex features ([38, 12]), new developments increasingly using image points directly ([50, 17, 24]). Direct approaches have the advantage that they are not reduced to certain feature points but can exploit all image points to determine the odometry and depth values and thus provide more dense reconstructions of the environment. Depending on how many image points are utilized, approaches can be divided into dense and semi-dense methods.

An example of a semi-dense approach is the Semi-Direct Monocular Visual Odometry (SVO)-algorithm, which is presented in the work of [24]. The method uses point features, but these are not explicitly extracted, but are an implicit result of a direct motion estimation. The initialization of the pose is achieved by minimizing the photometric error. LSD-SLAM [18] provides another direct approach. Based on the odometry method of [17], the algorithm generates globally consistent maps of the environment by means of graph optimization in large-area environments. Similar to the SVO algorithm, a probabilistic representation of the depth map is also used here to model inaccuracies. [48] also uses a probabilistic approach, but the method is based on a feature-based monocular SLAM system ([49]). Furthermore, in contrast to SVO and LSD-SLAM, the depth values of a reference image are not filtered over many individual images, but only key images are used for the reconstruction.

Stuehmer et. al. [66] presents one of the first real-time methods, which provides dense reconstructions with a monocular camera. The tracking of the camera is based on the approach of [38]. The reconstruction is carried out using several key images. By expanding to several images, regions that would be hidden in two images or would be outside the corresponding image can also be reconstructed with a higher probability. DTAM ([50]) also provides dense reconstructions in real-time. In order to estimate the depth values, the method performs a global energy reduction over many individual images. REMODE ([59]) is a method for the reconstruction of dense point clouds, which integrates a Bayesian estimate into the optimization process. By modeling uncertainties of measurement for each pixel, regularization can be controlled precisely and inaccuracies in the localization can be reduced. Real-time capability is achieved through a CUDA-based implementation. For the pose estimation the method of [24] is used. One of the recent developments of dense reconstructions is DPPTAM [11]. The approach reconstructs high textured regions with a semi-dense approach and low textured regions by the approximation of surfaces. Thereby the assumption is made that homogeneously colored image regions form a plane which can be determined by superpixels ([19]).

The procedures described so far fall under the category of online procedures, i.e. they are real-time capable and can deliver first results during camera recording. In contrast, offline methods require all collected recordings in advance and then carry out the corresponding calculations. In [25] Fuhrmann et. al. present a pipeline for reconstruction, which combines all necessary processing steps in a software framework called MVE. The framework is also capable of reconstructing texturized surfaces.

1.4.4 Registration methods

A basic prerequisite for many tasks, such as the navigation, mapping or the cooperation of UAV and UGV, is the robot localization. When working with three-dimensional point clouds, the registration is significantly involved in the success of an exact localization [34]. Since the aim of this work is that the UAV and UGV can be located together in a global map, a registration method must be found, which can handle point clouds from different domains. In this context, it is important that the methods for registration as well as the generation of vision based point clouds can be combined.

Methods for registration can be divided roughly into point-based or iterative and feature-based methods ([34, 53]). An example of a known iterative method is the ICP-algorithm, which has already been implemented in several variants. According to [5] the transformation is determined by minimizing the Euclidean distance of the found point correspondences. The search for corresponding points and the calculation of the associated transformation for the alignment of these points is finally repeated iteratively until set limits have been reached. A disadvantage of iterative methods, however, is that they can converge to a local minimum under certain assumptions, such as an insufficient overlay of the scenes [34]. In addition, they can be sensitive to outliers and can be very computationally intensive with large amounts of data [71]. If several point clouds have to be registered, the generated scene must also be globally consistent. To achieve better results, it is common that feature-based methods are used for the initial registration and iterative procedures are used for refining the already estimated transformation [34]. Features can be described by feature descriptors, that are incorporating geometric structures. If surfaces are used as a geometric structure, a high compression rate and thus a fast correspondence search can be made possible [53]. The work of [54] introduces a SLAM algorithm based on the registration of planar segments. The algorithm for the extraction of planar-based segments is based on the work of [60], which takes up the region-growing algorithm of [32] and adapts it by optimizations for the use in a SLAM system. For correspondence search and registration, the work of [55] is used. The presented MUMC-algorithm (Minimally Uncertain Maximum Consensus) maximizes geometric consistency while minimizing the resulting uncertainties. As shown in the work of [54], both faster and more robust results can be obtained in comparison to an ICP-alorithm. [73] provides another plane-based registration method, which is based on the work of [54]. An approach, that is also concerned with the registration of point clouds from different sensor groups, is presented in [30]. As a first step the method determines structural descriptors. For faster calculation, the descriptors are then projected into a subspace. A matching scheme is used to compare the descriptors and compute vote scores. The voting space is then used for place segmentation and for registration.

For TRADR, an improved algorithm is developed which is based on the approaches of [73] and [54]. The presented algorithm for surface extraction can be applied to unorganized point clouds and is fast in the calculation. The method of [54] has also proven itself in a test environment which is close to USAR environements.

1.4.5 Loop closures in 3D point clouds

Detecting loop-closures from 3D data is still an open problem in robot localization. The problem has been tackled with different approaches. We have identified three main trends: (i) approaches based on local features, (ii) global descriptors and (iii) based on planes or objects.

The works presented in [9, 75, 64, 65, 28] propose to extract local features from keypoints and perform matches on the basis of these features. Bosse and Zlot [9] extract keypoints directly from the point clouds and describe them with a 3D *Gestalt* descriptor. Keypoints then vote for their nearest neighbours in a *vote matrix* which is finally thresholded for recognizing places. Similar approach has been used in [28]. Apart from such Gestalt descriptors, a number of alternative local feature descriptors exist which can be used in similar frameworks. This includes features such as *fast point feature histogram (FPFH)* [63] which we employed to compare our approach. Alternatively, Zhuang et al. [75] transform the local scans into bearing-angle images and extract Speeded Up Robust Features (SURFs) from these images. A strategy based on 3D spatial information is employed to order the scenes before matching the descriptors. A similar technique by Steder et al. [64] first transforms the local scans into a range image. Local features are extracted and compared to the ones stored in a database, employing the Euclidean distance for matching keypoints. This work is extended in [65] by using Normal-Aligned Radial Features (NARF) descriptors and a bag of words approach for matching. Zhang and Singh [74] are able to estimate odometry in real-time using range data. Loop-closures are mentioned but rely on an offline algorithm.

Using global descriptors of the local point cloud for loop-closures is also proposed [62, 31, 45]. Rohling et al. [62] propose to describe each local point cloud with a 1D histogram of point heights, assuming that the sensor keeps a constant height above the ground. The histograms are then compared using the *Wasserstein* metric for recognizing places. Granström et al. [31] describe point clouds with rotation invariant features such as volume, nominal range, and range histogram. Distances are computed for scalar features and cross-correlation for histogram features, and an AdaBoost classifier is trained to match places. Finally, ICP is used for computing the relative pose between point clouds. In another approach, Magnusson et al. [45] split the cloud into overlapping grids and compute shape properties (spherical, linear, and several type of planar) of each cell and combine them into a matrix of surface shape histograms. Similar to other works, these descriptors are compared for finding loop-closures.

While local keypoint features often lack descriptive power, global descriptors can struggle with invariance. Therefore other works have also proposed to use 3D shapes or objects for the place recognition task. Fernandez-Moral et al. [20], for example, propose to perform place recognition by detecting planes in 3D environments. The planes are accumulated in a graph and an interpretation tree is used to match sub-graphs. A final geometric consistency test is conducted over the planes in the matched sub-graphs. The work is extended in [21] to use the covariance of the plane parameters instead of the number of points in planes for matching. This strategy is only applied to small, indoor environments and assumes a plane model for segments which is no longer valid in unstructured environment. A somewhat analogous, seminal work on object-based loop-closure detection in indoor environments using RGB-D cameras is presented by Finman et al. [22]. Although presenting interesting ideas, their work can only handle a small number of well segmented objects in small scale environments.

We therefore aim for an approach which does not rely on assumptions about the environment being composed of simplistic geometric primitives such as planes, or a rich library of objects. This allows for a more general, scalable solution. Inspiration is taken from Douillard et al. [13] and Nieto et al. [51] which proposed different SLAM techniques based on segments. A strategy for aligning Velodyne scans based on segments is proposed in [13] where the *Symmetric Shape Distance* is used to compare and match segments as defined in [14]. Analogously, [51] proposed an Extended Kalman Filter solution which uses segments as landmarks, rather than point features.

1.4.6 UGV-UAV Map Merging

Relative UAV to UGV localiation is an important part in the TRADR scenario for seamless integration of mission data. For example, if UAVs are used for reconaissance, interest points should be communicated to the UGV for further assessment. Therefore their maps need to be aligned. Vision to vision place recognition can be performed between the robots if both are equipped with cameras, as presented in [27, 44, 47]. In TRADR, the UAV is equipped with a camera due to weight constraints, while the UGV is equipped with both camera and 3D LiDAR sensor. However, tests showed that the drastically different viewpoints of UAV and UGV camera prevent reliable application of vision based place recognition in the TRADR scenarios, rendering the vision to vision matching infeasible. The other case can be LiDAR based place recognition, if both robots are equipped with LiDAR sensors, as presented in [7, 8, 16]. Furthermore, the obvious case of having a robot with both sensors can yield multi-modal maps and localization in these as shown by [33, 70, 52].

Our approach as presented in [29] focuses on the case, in which one robot is equipped with a LiDAR sensor, while the other on is carrying a camera. Former work on such applications focuses on simulating camera pictures using the LiDAR's intensity returns [72] or reconstructing dense maps from vision data [23]. While the first approach has been shown to work in urban environments, the strong viewpoint dependence and its weak generalization to cluttered environments make it infeasible for the application in TRADR. In contrast, reconstructing scenes from camera images has shown to be more general. However, these 3D reconstruction techniques as presented in [26, 68] are computationally expensive and fail in untextured or poorly illuminated areas.

Using a good initial guess, it was shown that sparse vision maps and dense LiDAR maps can be aligned [10]. With our approach we go beyond this and show that sparse visual keypoint locations and LIDAR maps contain sufficient mutual structural information to be merged without an initial guess.

References

[1] Constrained reinforcement learning from intrinsic and extrinsic rewards. In *IEEE 6th International Conference on Development and Learning*, pages 163–168, 2007. ISBN 978-1-4244-1116-0.

[2] Prashanth L. A. Policy gradients for cvar-constrained mdps. In *International Conference Algorithmic Learning Theory*, pages 155–169, 2014.

[3] Jonathan Baxter and Peter L. Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research (JAIR)*, 15:319–350, 2001.

[4] Nils Berninger, Hartmut Surmann, and Rainer Worst. Planar segment based vision-laser point cloud registration. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2017*, 2017. submitted.

[5] Paul J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14 (2):239–256, 1992.

[6] M. Bosse and R. Zlot. Keypoint design and evaluation for place recognition in 2D LIDAR maps. *Robotics and Autonomous Systems (RAS)*, 57(12):1211–1224, 2009.

[7] Michael Bosse. Place recognition using regional point descriptors for 3D mapping. *Field and Service Robotics*, pages 195–204, 2010. doi: 10.1007/978-3-642-13408-1_18.

[8] Michael Bosse and Robert Zlot. Place recognition using keypoint voting in large 3D lidar datasets. *ICRA*, pages 2677–2684, 2013.

[9] Michael Bosse and Robert Zlot. Place recognition using keypoint voting in large 3D lidar datasets. In *IEEE Int. Conf. on Robotics and Automation*, 2013.

[10] Tim Caselitz, Bastian Steder, Michael Ruhnke, and Wolfram Burgard. Matching Geometry for Long-term Monocular Camera Localization. In *ICRA Workshop: AI for long-term Autonomy*, 2016.

[11] Alejo Concha and Javier Civera. Dense Piecewise Planar Tracking and Mapping from a Monocular Sequence. In *Proc. of The International Conference on Intelligent Robots and Systems (IROS)*, 2015.

[12] Andrew J. Davison, Ian D. Reid, Nicholas D. Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *IEEE transactions on pattern analysis and machine intelligence*, 29(6):1052–1067, 2007.

[13] Bertrand Douillard, A Quadros, Peter Morton, James Patrick Underwood, Mark De Deuge, S Hugosson, M Hallström, and Tim Bailey. Scan segments matching for pairwise 3d alignment. In *IEEE Int. Conf. on Robotics and Automation*, 2012.

[14] Bertrand Douillard, J Underwood, Vsevolod Vlaskine, A Quadros, and S Singh. A pipeline for the segmentation and classification of 3d point clouds. In *Experimental Robotics*, pages 585–600. Springer, 2014.

[15] Renaud Dubé, Abel Gawel, Luigi Freda, Mario Gianni, Cesar Cadena, and Roland Siegwart. 3d localization, mapping and path planning for search and rescue operations. In *Proceedings of the IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*, 2016.

[16] Renaud Dubé, Daniel Dugas, Elena Stumm, Juan Nieto, Roland Siegwart, and Cesar Cadena. Segmatch: Segment based loop-closure for 3d point clouds. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (submitted)*, 2017.

[17] J. Engel, J. Sturm, and D. Cremers. Semi-dense visual odometry for a monocular camera. In *IEEE International Conference on Computer Vision (ICCV)*, Sydney, Australia, December 2013.

[18] J. Engel, T. Schöps, and D. Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In *European Conference on Computer Vision (ECCV)*, September 2014.

[19] Pedro F Felzenszwalb and Daniel P Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision*, 59(2): 167–181, 2004.

[20] Eduardo Fernandez-Moral, Walterio Mayol-Cuevas, Vicente Arevalo, and Javier Gonzalez-Jimenez. Fast place recognition with plane-based maps. In *IEEE Int. Conf. on Robotics and Automation*, 2013.

[21] Eduardo Fernández-Moral, Patrick Rives, Vicente Arévalo, and Javier González-Jiménez. Scene structure registration for localization and mapping. *Robotics and Autonomous Systems*, 75:649–660, 2016.

[22] Ross Finman, Liam Paull, and John J Leonard. Toward object-based place recognition in dense rgb-d maps. In *ICRA workshop on visual place recognition in changing environments*, 2015.

[23] Christian Forster, Matia Pizzoli, and Davide Scaramuzza. Air-ground localization and map augmentation using monocular dense reconstruction. In *IROS*, pages 3971–3978, 2013.

[24] Christian Forster, Matia Pizzoli, and Davide Scaramuzza. SVO: Fast semi-direct monocular visual odometry. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.

[25] Simon Fuhrmann, Fabian Langguth, and Michael Goesele. Mve-a multiview reconstruction environment. In *Proceedings of the Eurographics Workshop on Graphics and Cultural Heritage (GCH)*, volume 6, page 8, 2014.

[26] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8):1362–1376, 2010.

[27] Dorian Gálvez-López and Juan D Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 28(5):1188–1197, 2012.

[28] Abel Gawel, Titus Cieslewski, Renaud Dubé, Michael Bosse, Roland Siegwart, and Juan Nieto. Structure-based Vision-Laser Matching. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Daejeon, 2016.

[29] Abel Gawel, Titus Cieslewski, Renaud Dubé, Mike Bosse, Roland Siegwart, and Juan Nieto. Structure-based vision-laser matching. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 182–188. IEEE, 2016.

[30] Abel Gawel, Titus Cieslewski, Renaud Dubé, Mike Bosse, Roland Siegwart, and Juan Nieto. Structure-based vision-laser matching. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 182–188. IEEE, 2016.

[31] Karl Granström, Thomas B Schön, Juan I Nieto, and Fabio T Ramos. Learning to close loops from range data. *The Int. Journal of Robotics Research*, 30(14):1728–1754, 2011.

[32] Dirk Hähnel, Wolfram Burgard, and Sebastian Thrun. Learning compact 3d models of indoor and outdoor environments with a mobile robot. *Robotics and Autonomous Systems*, 44(1):15–27, 2003.

[33] Kin Leong Ho and Paul Newman. Loop closure detection in SLAM by combining visual and spatial appearance. *Robotics and Autonomous Systems*, 54(9):740–749, 2006. doi: 10.1016/j.robot.2006.04.016.

[34] Dirk Holz, Alexandru E Ichim, Federico Tombari, Radu B Rusu, and Sven Behnke. Registration with the point cloud library: A modular framework for aligning in 3-d. *IEEE Robotics & Automation Magazine*, 22(4):110–124, 2015.

[35] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

[36] Michael Kaess, Hordur Johannsson, Richard Roberts, Viorela Ila, John Leonard, and Frank Dellaert. isam2: Incremental smoothing and mapping with fluid relinearization and incremental variable reordering. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 3281–3288. IEEE, 2011.

[37] Philipp Florian Käshammmer. A semantic interpreter for multimodal and multirobot data. Master's thesis, Czech Technical University in Prague, Czech Republic, 2016.

[38] Georg Klein and David Murray. Parallel tracking and mapping for small AR workspaces. In *Proc. Sixth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR'07)*, Nara, Japan, November 2007.

[39] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Advances In Neural Information Processing Systems*, pages 1–9, 2012. ISSN 10495258. doi: http://dx.doi.org/10.1016/j.protcy.2014.09.007.

[40] Andras Kupcsik, Marc Peter Deisenroth, Jan Peters, Loh Ai Poh, Prahlad Vadakkepat, and Gerhard Neumann. Model-based contextual policy search for data-efficient generalization of robot skills. *Artificial Intelligence*, 2014.

[41] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-End Training of Deep Visuomotor Policies. *Journal of Machine Learning Research*, 17(1):1334–1373, 2016.

[42] Sergey Levine, Peter Pastor, Alex Krizhevsky, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *CoRR*, abs/1603.02199, 2016.

[43] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. ISBN 9781467369640. doi: 10.1109/CVPR.2015.7298965.

[44] Simon Lynen, Michael Bosse, Paul Furgale, and Roland Siegwart. Placeless place-recognition. In *International Conference on 3D Vision*, pages 303–310, 2014.

[45] Martin Magnusson, Henrik Andreasson, Andreas Nüchter, and Achim J Lilienthal. Automatic appearance-based loop detection from three-dimensional laser data using the normal distributions transform. *Journal of Field Robotics*, 26(11-12):892–914, 2009.

[46] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 02 2015.

[47] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós. ORB-SLAM: a Versatile and Accurate Monocular SLAM System. *IEEE Transactions on Robotics (TRO)*, 31(5):1147–1163, 2015.

[48] Raúl Mur-Artal and Juan D Tardós. Probabilistic semi-dense mapping from highly accurate feature-based monocular slam. *Proceedings of Robotics: Science and Systems, Rome, Italy*, 1, 2015.

[49] Raúl Mur-Artal, J. M. M. Montiel, and Juan D. Tardós. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics*, 31(5):1147–1163, 2015.

[50] Richard A. Newcombe, Steven J. Lovegrove, and Andrew J. Davison. Dtam: Dense tracking and mapping in real-time. In *2011 international conference on computer vision*, pages 2320–2327. IEEE, 2011.

[51] Juan Nieto, Tim Bailey, and Eduardo Nebot. Scan-slam: Combining ekf-slam and scan correlation. pages 167–178. Springer, 2006.

[52] Gaurav Pandey, James R McBride, Silvio Savarese, and Ryan M Eustice. Toward mutual information based automatic registration of 3D point clouds. In *IROS*, pages 2698–2704, 2012.

[53] Kaustubh Pathak, Narunas Vaskevicius, Jann Poppinga, Max Pfingsthorn, Sören Schwertfeger, and Andreas Birk. Fast 3d mapping by matching planes extracted from range sensor point-clouds. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, 2009.

[54] Kaustubh Pathak, Andreas Birk, Narunas Vaskevicius, Max Pfingsthorn, Sören Schwertfeger, and Jann Poppinga. Online three-dimensional SLAM by registration of large planar surface segments and closed-form pose-graph relaxation. *Journal of Field Robotics*, 27(1): 52–84, 2010.

[55] Kaustubh Pathak, Andreas Birk, Narunas Vaskevicius, and Jann Poppinga. Fast registration based on noisy planes with unknown correspondences for 3-d mapping. *IEEE Transactions on Robotics*, 26(3):424–441, 2010.

[56] Martin Pecka, Votjěch Šalanský, Karel Zimmermann, and Tomas Svoboda. Autonomous flipper control with safety constraints. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 2889–2894, Los Alamitos, USA, Oct 2016. IEEE. ISBN 978-1-5090-3762-9. doi: 10.1109/IROS.2016.7759447.

[57] Martin Pecka, Karel Zimmermann, Michal Reinstein, and Tomáš Svoboda. Controlling robot morphology from incomplete measurements. *IEEE Transactions on Industrial Electronics*, accepted for publication 05/2016. Special issue on: on Motion Control for Novel Emerging Robotic Devices and Systems.

[58] T. Petříček, V Šalanský, K. Zimmermann, and T. Svoboda. Guided reinforcement learning for simultaneous exploration and segmentation. *submitted to IEEE Transactions on Robotics*, 2017. unpublished, submitted.

[59] Matia Pizzoli, Christian Forster, and Davide Scaramuzza. REMODE: Probabilistic, monocular dense reconstruction in real time. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.

[60] Jann Poppinga, Narunas Vaskevicius, Andreas Birk, and Kaustubh Pathak. Fast plane detection and polygonalization in noisy 3d range images. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3378–3383. IEEE, 2008.

[61] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Neural Information Processing Systems (NIPS)*, pages 1–10, 2015. URL `http://arxiv.org/abs/1506.01497`.

[62] Timo Rohling, Jennifer Mack, and Dirk Schulz. A fast histogram-based similarity measure for detecting loop closures in 3-d lidar data. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2015.

[63] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. In *IEEE Int. Conf. on Robotics and Automation*, pages 3212–3217, 2009.

[64] Bastian Steder, Giorgio Grisetti, and Wolfram Burgard. Robust place recognition for 3D range data based on point features. In *IEEE Int. Conf. on Robotics and Automation*, 2010.

[65] Bastian Steder, Michael Ruhnke, Slawomir Grzonka, and Wolfram Burgard. Place recognition in 3d scans using a combination of bag of words and point feature based relative pose estimation. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2011.

[66] J. Stühmer, S. Gumhold, and D. Cremers. Real-time dense geometry from a handheld camera. In *Pattern Recognition (Proc. DAGM)*, pages 11–20, Darmstadt, Germany, September 2010.

[67] Russ Tedrake, Zack Jackowski, Rick Cory, John William Roberts, and Warren Hoburg. Learning to fly like a bird. *Under review*, 8, 2009.

[68] C Tomasi and T Kanade. Shape and motion from image streams: a factorization method. *Proceedings of the National Academy of Sciences of the United States of America*, 90(21):9795–9802, 1993.

[69] AA Transeth, KY Pettersen, and P Liljebäck. A survey on snake robot modeling and locomotion. *Robotica*, 2009.

[70] Fredy Tungadi, W. L D Lui, Lindsay Kleeman, and Ray Jarvis. Robust online map merging system using laser scan matching and omnidirectional vision. *IROS*, pages 7–14, 2010.

[71] Diego Viejo and Miguel Cazorla. 3d plane-based egomotion for slam on semi-structured environment. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2761–2766. IEEE, 2007.

[72] Ryan W. Wolcott and Ryan M. Eustice. Visual Localization within LIDAR Maps for Automated Urban Driving. In *IROS*, pages 176 – 183, 2014. ISBN 9781479969340. doi: 10.1109/IROS.2014.6942558.

[73] Junhao Xiao, Benjamin Adler, Jianwei Zhang, and Houxiang Zhang. Planar segment based three-dimensional point cloud registration in outdoor environments. *Journal of Field Robotics*, 30(4):552–582, 2013.

[74] Ji Zhang and Sanjiv Singh. Loam: Lidar odometry and mapping in real-time. In *Robotics: Science and Systems*, 2014.

[75] Yan Zhuang, Nan Jiang, Huosheng Hu, and Fei Yan. 3-d-laser-based scene measurement and place recognition for mobile robots in dynamic indoor environments. *IEEE Transactions on Instrumentation and Measurement*, 62(2):438–450, 2013.

[76] Karel Zimmermann, Petr Zuzanek, Michal Reinstein, and Vaclav Hlavac. Adaptive Traversability of Unknown Complex Terrain with Obstacles for Mobile Robots. In *IEEE International Conference on Robotics and Automation*, pages 5177—-5182, 2014.

2 Annexes

2.1 Dubé (2017), "*SegMatch*: Segment based loop-closure for 3D point clouds"

Bibliography Renaud Dubé, Daniel Dugas, Elena Stumm, Juan Nieto, Roland Siegwart and Cesar Cadena. "*SegMatch*: Segment based loop-closure for 3D point clouds." In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (submitted)*, 2017.

Abstract Loop-closure detection on 3D data is a challenging task that has been commonly approached by adapting image-based solutions. Methods based on local features suffer from ambiguity and from robustness to environment changes while methods based on global features are viewpoint dependent. We propose *SegMatch*, a reliable loop-closure detection algorithm based on the matching of 3D segments. Segments provide a good compromise between local and global descriptions, incorporating their strengths while reducing their individual drawbacks.

SegMatch does not rely on assumptions of 'perfect segmentation', or on the existence of 'objects' in the environment, which allows for reliable execution on large scale, unstructured environments. We quantitatively demonstrate that *SegMatch* can achieve accurate localization at a frequency of 1Hz on the largest sequence of the KITTI odometry dataset. We furthermore show how this algorithm can reliably detect and close loops in real-time, during online operation. In addition, the source code for the *SegMatch* algorithm is made publicly available.

Relation to WP Describes the loop-closure strategy for 3D point-clouds introduced in WP 1, T1.5, Section 1.3.1.

Availablity Unrestricted and available online (`https://arxiv.org/pdf/1609.07720v1.pdf`).

2.2 Gawel (2016), "Structure-based Vision-Laser Matching"

Bibliography Abel Gawel, Titus Cieslewski, Renaud Dubé Mike Bosse, Roland Siegwart and Juan Nieto. "Structure-based Vision-Laser Matching." In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016.

Abstract Persistent merging of maps created by different sensor modalities is an insufficiently addressed problem. Current approaches either rely on appearance-based features which may suffer from lighting and viewpoint changes or require pre-registration between all sensor modalities used. This work presents a framework using structural descriptors for matching LIDAR point-cloud maps and sparse vision keypoint maps. The matching algorithm works independently of the sensors' viewpoint and varying lighting and does not require pre-registration between the sensors used. Furthermore, we employ the approach in a novel vision-laser map-merging algorithm. We analyse a range of structural descriptors and present results of the method integrated within a full mapping framework. Despite the fact that we match between the visual and laser domains, we can successfully perform map-merging using structural descriptors. The effectiveness of the presented structure-based vision-laser matching is evaluated on the public KITTI dataset and furthermore demonstrated on a map merging problem in an industrial site.

Relation to WP Describes a map registration strategy for vision and laser pointclouds introduced in WP 1, T1.5, Section 1.3.1.

Availablity Unrestricted and available online (`http://rpg.ifi.uzh.ch/docs/IROS16_Gawel.pdf`).

2.3 Pecka-TIE2016, "Controlling Robot Morphology from Incomplete Measurements"

Bibliography Pecka, Martin and Zimmermann, Karel and Reinstein, Michal and Svoboda, Tomáš. "Controlling Robot Morphology from Incomplete Measurements" In *IEEE Transactions on Industrial Electronics*, Special issue on: on Motion Control for Novel Emerging Robotic Devices and Systems. Accepted, early access `http://ieeexplore.ieee.org/document/7490379/`.

Abstract Mobile robots with complex morphology are essential for traversing rough terrains in Urban Search & Rescue missions (USAR). Since teleoperation of the complex morphology causes high cognitive load of the operator, the morphology is controlled autonomously. The autonomous control measures the robot state and surrounding terrain which is usually only partially observable, and thus the data are often incomplete. We marginalize the control over the missing measurements and evaluate an explicit safety condition. If the safety condition is violated, tactile terrain exploration by the body-mounted robotic arm gathers the missing data.

Relation to WP Describes an approach for automatic robot control on rough terrain. Contributes to the robot perception suite. T1.5.

Availablity Unrestricted. Included in the public version of this deliverable. `https://arxiv.org/abs/1612.02739`.

2.4 Pecka-IROS-2016, "Autonomous Flipper Control with Safety Constraints"

Bibliography Martin Pecka, Votjech Salansky, Karel Zimmermann, and Tomas Svoboda. "Autonomous flipper control with safety constraints". In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*

Abstract Policy Gradient methods require many real-world trials. Some of the trials may endanger the robot system and cause its rapid wear. Therefore, a safe or at least gentle-to-wear exploration is a desired property. We incorporate bounds on the probability of unwanted trials into the recent Contextual Relative Entropy Policy Search method. The proposed algorithm is evaluated on the task of autonomous flipper control for a real Search and Rescue rover platform.

Relation to WP Describes a novel reinforcement learning approch that effectively combines simulator and real roll-outs of the ground robot. Contributes to the robot perception suite. T1.5.

Availablity Unrestricted. Included in the public version of this deliverable. `http://cmp.felk.cvut.cz/~peckama2/papers/2016_IROS_Autonomous_Flipper_Control_with_Safety_Constraints.pdf`.

2.5 Jasek-bc-thesis-2016, "Detecting Objects for Autonomous System Verification"

Bibliography Jašek, Otakar. Detecting Objects for Autonomous System Verification. Bachelor Thesis, Czech Technical University in Prague.

Abstract In this thesis we created a framework for easy evaluation and training of Faster R-CNN type of networks. We fine-tuned VGG16 and ZFNet networks on our internal Victims dataset as well as standard KITTI dataset. We later showed that VGG16 architecture is far more suitable for fine-tuning on data from slightly different training and target domains. This framework can later serve as a baseline for further improvements in the field.

Relation to WP Contributes to the scene analysis. T1.5.

Availablity Unrestricted. Available for download. `https://dspace.cvut.cz/bitstream/handle/10467/64667/F3-BP-2016-Jasek-Otakar-jasek.pdf`

2.6 Kashammer-ms-thesis-2016, "A semantic interpreter for multimodal and multirobot data"

Bibliography Käshammmer, Philipp Florian. "A semantic interpreter for multimodal and multirobot data". Master thesis, Czech Technical University in Prague and Lulea University of Technology. 2016.

Abstract Huge natural disaster events can be so devastating that they often overwhelm human rescuers and yet, they seem to occur more often. The TRADR (Long-Term Human- Robot Teaming for Robot Assisted Disaster Response) research project aims at devel- oping methodology for heterogeneous teams composed of human rescuers as well as ground and aerial robots. While the robots swarm the disaster sites, equipped with advanced sensors, they collect a huge amount row-data that cannot be processed ef- ficiently by humans. Therefore, in the frame of the here presented work, a semantic interpreter has been developed that crawls through the raw data, using state of the art object detection algorithms to identify victim targets and extracts all kinds of informa- tion that is relevant for rescuers to plan their missions. Subsequently, this information is restructured by a reasoning process and then stored into a high-level database that can be queried accordingly and ensures data constancy.

Relation to WP Contributes to the scene analysis, closing higher and lower loops. T1.5.

Availablity Unrestricted. Available for download. `ftp://cmp.felk.cvut.cz/pub/cmp/articles/svoboda/Kashammer-ms-thesis-2016.pdf`

2.7 Petricek-TRO2017-unpub, "Guided reinforcement learning for simultaneous exploration and segmentation"

Bibliography Tomáš Petříček, Vojtěch Šalanský, Karel Zimmermann, and Tomáš Svoboda. "Guided reinforcement learning for simultaneous exploration and segmentation". Submitted to *IEEE Transactions on Robotics*

Abstract We consider the problem of pan-tilt sensor control for active segmentation of incomplete multi-modal data. Since demanding optimal control does not allow for online replanning, we rather employ the optimal planner offline to provide guiding samples for learning of a CNN-based control policy in a guided Q-learning framework. The proposed policy initialization and guided Q-learning avoids poor local optima and yields reasonable results from hundreds of roll-outs. The results suggest that the proposed policy outperforms the baseline and is suitable for real-time control.

Relation to WP An active sensing algorithm for scene segmentation. Contributes to T1.5.

Availability Restricted. Not included in the public version of this deliverable.

2.8 Berninger-IROS2017-unpub, "Planar segment based vision-laser point cloud registration"

Bibliography Nils Berninger, Hartmut Surmann, Rainer Worst. "Planar segment based vision-laser point cloud registration". Submitted to IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2017.

Abstract This work is concerned with the registration of point clouds, which were generated from laser scans and camera recordings. The registration method is based on the matching of corresponding planar segments, that are extracted from the point clouds. Based on the registration, an approach for a globally optimized localization is presented. Apart from the structural information of the point clouds, no further information is required for the localization. Experiments show the results of the overall registration.

Relaion to WP Contributes to the world modeling, T1.6.

Availability Restricted. Not included in the public version of this deliverable.