



Intra-Logistics with Integrated Automatic Deployment:
Safe and Scalable Fleets in Shared Spaces

H2020-ICT-2016-2017
Grant agreement no: 732737

DELIVERABLE 1.2

Precise localisation and mapping in dynamic environments

Due date: month 18 (June 2018)
Deliverable type: R
Lead beneficiary: ORU

Dissemination Level: PUBLIC

Main author: Martin Magnusson (ORU)

Contents

1 Introduction	3
2 Structural mapping	4
2.1 Map representation	4
2.2 Scan registration	5
2.3 Heterogeneous mapping	8
2.3.1 Integrating prior map data while mapping	8
2.3.2 Sensor-map alignment with region decomposition	10
3 Localisation	12
4 Summary	12

1 Introduction

This report presents a summary of the work on localisation and mapping (tasks T1.2 and T1.3) in the EU H2020 project ILIAD during its first 18 months. The system has been jointly developed by partners Örebro University (ORU) and University of Lincoln (UoL), with ORU being the main beneficiary.

In particular, we demonstrate results from a prototype implementation that is capable of creating consistent 3D maps and accurately localising in them, with a mean absolute trajectory error (ATE) of well below 30 mm, which is required in industrial practice for picking up pallets using global positioning. Specific challenges of localisation and mapping in a warehouse setting is that the environment is gradually changing, and has many semi-static objects, which cannot easily be filtered out the same way as fully dynamic objects, such as moving people. Challenges also include perceptual aliasing from long racks of shelves that are self-similar, and the fact that long aisles provide little geometric features for use in scan matching for mapping and localisation.

Our mapping and localisation framework is also capable of creating maps using heterogeneous data. Two maps of the same environment, created with a 2D lidar at a specific height, or a 3D lidar, may look quite different, and merging them (especially if they are not completely overlapping) is not a trivial task. Layout maps of the warehouse is another form of map data that is relevant for automatic deployment of a fleet of warehouse trucks.

The main objective of WP1 is to facilitate easy deployment of an autonomous truck fleet in a new site, requiring little assistance from human operators. The core components that are the focus of this deliverable are robust structural mapping and localisation in a dynamic environment, using a team of robots with heterogeneous sensors.

The underlying map representation is presented in Section 2.1, and the scan registration methods used when constructing the maps in Section 2.2. The incorporation of heterogeneous data is presented in Section 2.3. A demonstration of the localisation accuracy attainable by the ILIAD framework is provided in Section 3. Figure 1 shows a sample point-cloud map from one of the ILIAD end user sites, for illustration.

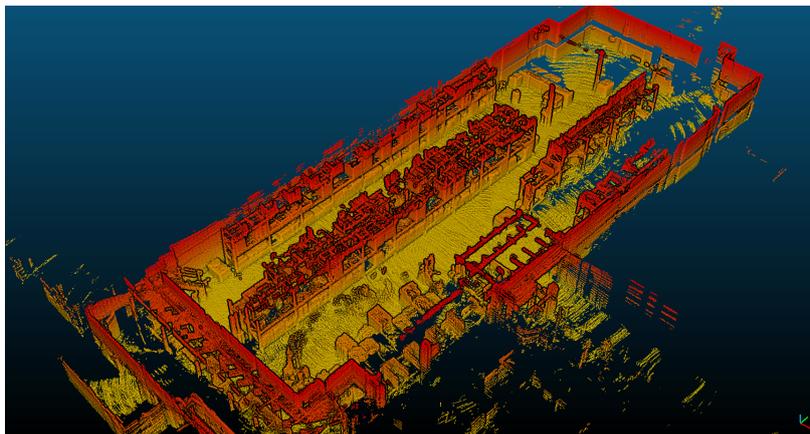


Figure 1: A 3D point-cloud map from one of the warehouses of partner Orkla Foods.

2 Structural mapping

2.1 Map representation

ILIAD's task T1.3 is concerned with the creation and maintenance of consistent maps. The mapping toolkit in ILIAD is primarily based on the NDT-OM (normal-distributions transform occupancy map) representation [12], which combines the NDT (normal-distributions transform) representation [9] with occupancy grids for consistent structural mapping in dynamic environments. As demonstrated in our previous work, NDT-OM also works well as a basis for robust and accurate localisation using NDT-MCL (NDT Monte-Carlo localisation) [13].

The NDT representation was originally developed in the context of 2D lidar registration [3]. The central idea is to represent surfaces by a set of Gaussian probability distributions, distributed in a grid structure. NDT has later been extended to three dimensions in the context of 3D scan registration [9]. However, NDT in itself does not explicitly represent free space, and has no mechanism to update the recorded Gaussians when faced with new observations; e.g., for removing obstacles that are no longer in the map. Therefore it is not well suited for navigation.

NDT-OM, on the other hand, supports consistent probabilistic updates of re-observed portions of the environment, as well as the generation of multi-resolution maps for navigation. Because each NDT cell in the map holds more information than that of a regular occupancy grid, NDT-OM can have comparably lower resolution, thereby improving memory and CPU requirements without sacrificing the accuracy of localisation or the representation of surfaces [12, 13].

More specifically, the base ILIAD map representation is a *set* of metric NDT-OM submaps, connected by topological graph edges that represent inter-map traversability. Technically, we have substantially revised ORU's existing NDT-OM implementation during ILIAD's first 18 months.

While the concept of hybrid metric-topological maps (i.e., metric submaps such as occupancy grids embedded in a topological structure) is not new, there has been a very limited amount of work that focuses on *how* the submaps should be created in order to provide high localisation accuracy. Scientifically, we have studied different approaches to determine when a new submap should be created, and which submaps should be updated when new sensor measurements arrive. We plan to submit this work for publication during ILIAD's second period.

A challenge in developing useful map representations for localisation is that an object in the environment may appear quite different depending on where it is observed from. For example, it is easier to recognise an object when it is observed from two reasonably close locations rather than from two very different points. Similarly, when a navigation map (e.g., an occupancy grid map or NDT map) is generated from point cloud observations, the measurement process noise from fusing disparate observations will be larger than when using a single observation. As a result, a better localisation estimate can possibly be provided using local submaps that are built from poses within a small region. But what is the best submap size? The spectrum ranges from having one map per sensor input as one extreme, to having a single monolithic map aggregated from all measurements as the other.

In particular, we have evaluated two methods for allocating submaps. One uses a fixed coarse grid and assigns a local map to each cell (*submap grid indexing*). The other creates a new submap whenever the robot has moved past an *interchange radius* threshold from the origin of the closest existing submap. Our experiments indicate that the latter method (*closest node position*) is better for localisation. Section 3 reports on our study of localisation accuracy attainable using these maps.

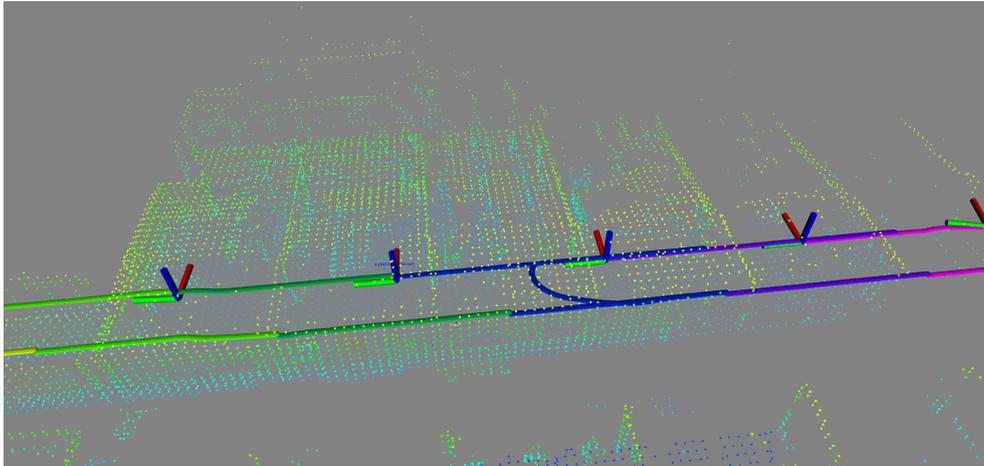


Figure 2: Illustration of the NDT-OM submap structure used for structural mapping in ILIAD. The origin of each submap is marked by a red/green/blue coordinate frame. The NDT-OM submap that belongs to the second node from the left is shown as a point cloud, where each point marks the mean of the Gaussian in one map cell. (The covariances are not displayed.) As can be seen from this image, the submaps are layed out so that they have a large amount of overlap. The colours of the trajectory lines denote which submap has been updated by scans from each point on the trajectory. Since new maps are added incrementally when using the closest node position strategy, the closest map to update from each point may change during multiple runs through the environment, as can be seen from the dual colours of parts of the trajectory.

Figure 2 illustrates an example of the HMT NDT-OM structure, when using the closest node position strategy.

Our framework facilitates mapping and localisation both in 2D and 3D. Figure 3 shows an example 2D map. Please note that this 2D map was created with the safety lidar mounted at floor height. Despite being faced with high levels of clutter (walking people, moving trucks), the map correctly represents only the stationary parts of the environment, with high fidelity. Figure 4 shows the truck platform used in this data collection, and its sensors.

2.2 Scan registration

While constructing each submap, incoming point clouds from lidar sensors are incrementally registered to the corresponding local submap. As a starting point, we have used the D2D-NDT registration method [15] (distribution-to-distribution NDT registration) for this purpose.

One common failure case for scan registration is in feature-sparse environments such as corridors (or warehouse aisles), where there is not enough structure to fully determine where along the corridor’s direction that a scan fits best. In order to address this problem, in ILIAD we have developed a scan matching algorithm that incorporates ego-motion information from the robot’s wheel sensors in a novel way [2]. By including an uncertainty-aware odometry ego-motion estimate in the objective function of registration, solutions are “softly” constrained to poses that match the expected motion of the vehicle (as estimated by odometry), which leads to better mapping and localisation performance.

Formally, the goal of scan registration is to minimise some objective function, $\min_{\mathbf{p}} f(\mathbf{p})$,

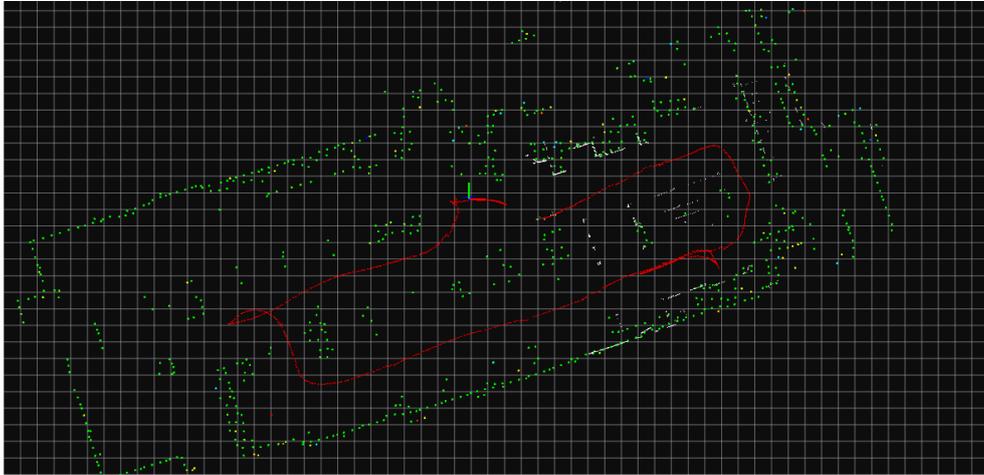
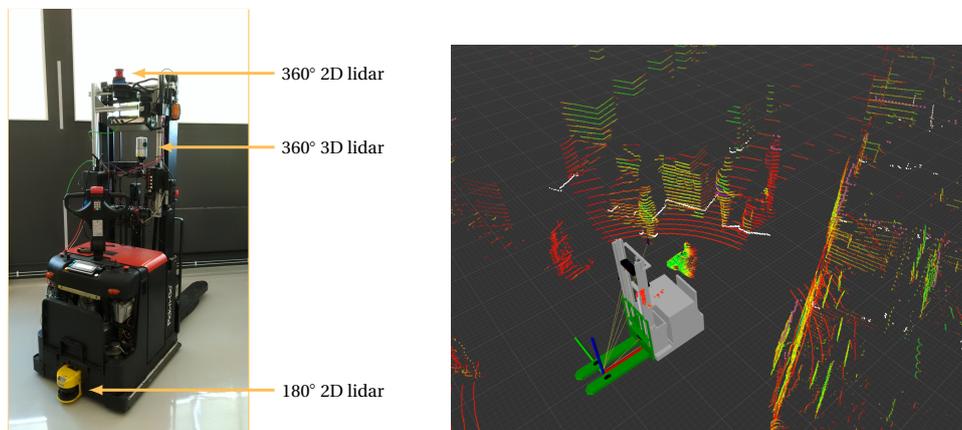


Figure 3: Sample 2D map from one of ILIAD's target end-user environments, created by the mapping toolkit described in Section 2.1. Please note that the 2D map was created with the safety lidar mounted at floor height. Despite the presence of clutter (walking people, moving trucks), the map correctly represents only the stationary parts of the environment.



(a) The sensor-equipped truck platform for data collection at end user sites, displaying the sensors used for navigation.

(b) View from data collection in a cluttered, dynamic environment. Data from the navigation sensors are shown with coloured points (360° 3D lidar data) and white points (180° 2D safety lidar data at floor height). The sensor data is obstructed by walking people and manually operated trucks.

Figure 4: Data collection in intralogistics operations.

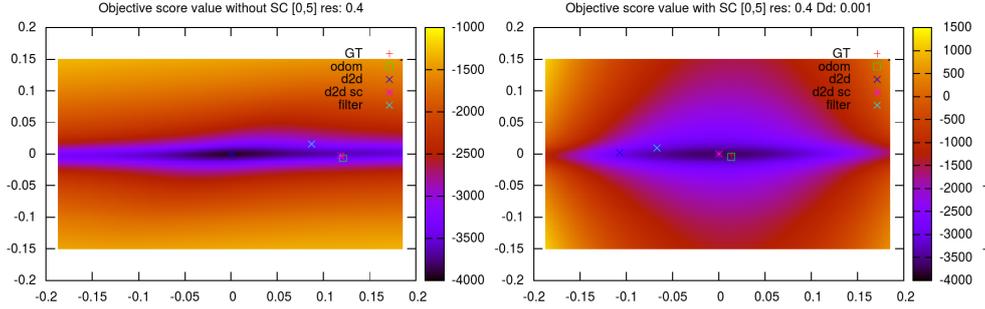


Figure 5: Illustrating the change of shape of the registration objective function with and without soft constraints [2]. *Left*: in a corridor environment, the lidar-only objective function f_{D2D} (marked d2d in the plot) does not have an optimum at the correct pose (marked GT). *Right*: adding soft constraints changes the shape of the objective function, resulting in an optimum (d2d sc in the plot) near the ground truth.

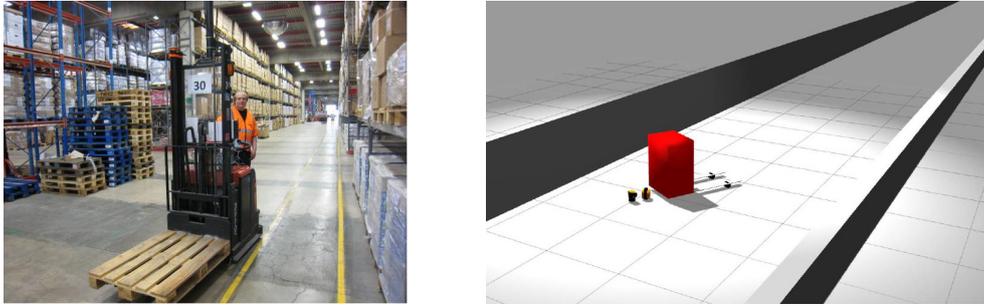


Figure 6: Real-world and simulated data from feature-sparse environments, used in the evaluation of D2D-NDT registration with soft constraints from egomotion [2].

to find the pose parameters \mathbf{p} that optimises the fit between two scans. We reformulate the D2D-NDT objective function $f_{D2D}(\mathbf{p})$ as

$$\min_{\mathbf{p}} f_{D2D}(\mathbf{p}) + \lambda(\mathbf{p} - \mathbf{p}_0)^T \Sigma^{-1}(\mathbf{p} - \mathbf{p}_0),$$

where \mathbf{p}_0 is the mean of the egomotion estimate provided by odometry, Σ is its covariance estimate, and λ is a penalty coefficient. Figure 5 illustrates the effect of adding soft constraints in this manner.

We have evaluated D2D-NDT with soft constraints on relevant real-world warehouse data as well as simulated corridor data (see Figure 6). The evaluation has shown that our method is especially beneficial in cases that due to the environment topology, limitations of the sensor range, or the mapping resolution, do not provide enough features to reliably estimate relative translation from lidar data alone. Figure 7 illustrates the performance of D2D-NDT SC in a simulated endless corridor environment, compared to odometry alone and D2D-NDT registration where the odometry is used as an initial estimate to the registration algorithm. For further technical details, please refer to Andreasson et al. [2].

The soft constraints technique described above can be used not only with D2D-NDT registration. In fact, it can be included in the objective function of most scan registration algorithms. One promising future development is to use this type of registration with *semantic-assisted* registration, dubbed SE-NDT, that has also been developed as part of ILIAD's Work Package 1. During Year 1, we demonstrated a scan matching algorithm that

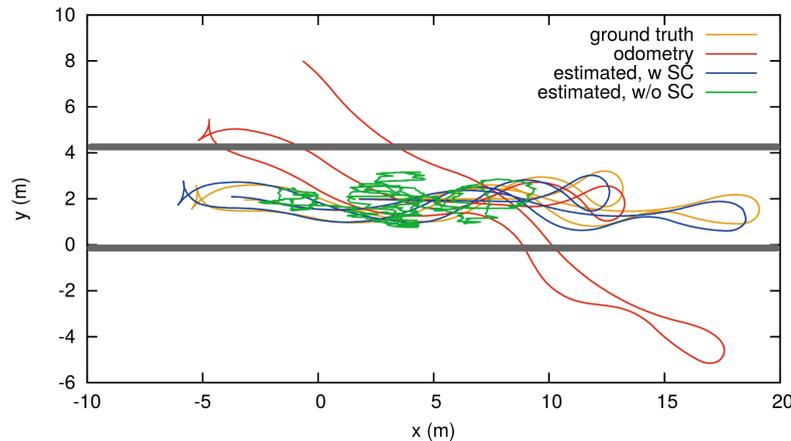


Figure 7: Pose estimated from from driving back and forth in the simulated endless corridor environment from Figure 6; with soft constraints (w SC) and without (w/o SC), odometry and ground truth. [2]

partitions 3D scans from the environment into flat and non-flat regions and then matches the regions in combination, also based on the D2D-NDT formulation [17]. Specifically, the objective function is split into two terms: one matching only flat regions, and one matching only edge/corner regions. We showed that doing so improves the robustness of 3D scan matching in certain important use cases, at the cost of longer execution times. Most notably for ILIAD’s objectives is an improved success rate from very feature-sparse environments, as long as the overlap between consecutive scans is large; which corresponds well to an autonomous navigation scenario in warehouse environments with limited geometric features. More recently [16], we have trained a modified version of the *PointNet* deep-learning architecture to include more descriptive semantic labels in SE-NDT, and a more fine-grained division than flat vs. non-flat.

2.3 Heterogeneous mapping

ILIAD is also concerned with heterogeneous fleets of warehouse robots, not all of which may have the same sensor setup. Consequently, heterogeneous map information needs to be integrated and aligned in a common reference frame. To accommodate this, we intend to integrate two approaches for multi-modal map merging, described in the following. The first method (Section 2.3.1) integrates a pre-existing layout map of the environment in the mapping process (e.g., while constructing a local NDT-OM submap), while the second (Section 2.3.2) can be used in a post-processing step in order to align two overlapping robot maps made by different robots — e.g., one from a floor-mounted 2D safety scanner, and one with a 3D lidar.

2.3.1 Integrating prior map data while mapping

In many cases, a layout map is already available of warehouses, which can be used both to aid robot mapping and as a means of human–robot interaction; e.g., by drawing regions in the layout map where the robots should not go. If robots can match the sensor data that they perceive to the prior layout map, they can “auto-complete” their map even before having seen the whole site. However, layout maps are typically not metrically accurate (scales of rooms are not consistent) and many objects that are present in the environ-

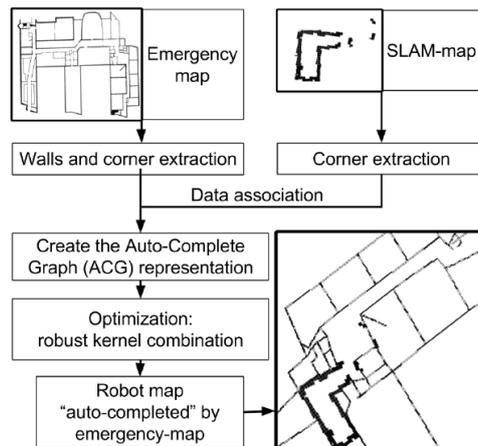


Figure 8: Flow chart depicting the process of integrating prior map data from layout maps (e.g., emergency evacuation maps) with robot maps created with online sensor data. This figure is from Mielle et al. [10].

ment, and seen by the robot, are not present in the layout map. Therefore, automatically incorporating this kind of prior map data is a very challenging task.

We intend to integrate a method for incorporating prior data from layout maps that has been developed by partner ORU in the EU H2020 project SmokeBot.¹ This method merges a layout map and the robot’s sensor map by using corners from both maps as landmarks in a graph-SLAM formulation, combined with two robust SLAM back-ends in tandem in order to deal with the highly uncertain data association between the corners that are observed in the sensory data and the ones that are present in the layout map. The procedure is outlined in Figure 8. In the scenarios targeted by the SmokeBot project, prior layout maps typically come from the emergency evacuation maps that are posted in buildings. In ILIAD, they may come from CAD-style drawings. Neither source of prior map is accurate enough for direct alignment to a sensor map.

The method is outlined in the following. For a full technical description, see Mielle et al. [10]. Corners in the layout map are found using a line-following algorithm. Corners in the sensor map (the current NDT-OM submap) are found by comparing the orientations of Gaussians in neighbouring map cells. The corner features from both map representations are included as feature nodes in a factor graph. The corners of the prior map are connected with graph edges initiated by the length of the wall. The information matrix associated with these edges is designed to allow the walls to extend or shrink (to account for non-uniform scale differences), but penalise rotation w.r.t. walls that are connected at a corner (to avoid corners being bent). The graph also contains robot pose nodes given by odometry or scan registration (as in Section 2.2). Each pose node has an edge to all the sensor-map corners that are visible from that pose. Finally, each sensor-map corner is associated with a graph edge to all layout-map corners within a certain radius.

In principle, an off-the-shelf SLAM back-end could be used to optimise the poses of the nodes in the graph in order to maximise the alignment between the layout and the sensor map. However, given the difference in scale, and the uncertainty in corner-to-corner data association, even state-of-the-art robust back-ends such as dynamic covariance scaling [1] (DCS) are not able to correctly align the maps under realistic conditions. Our solution is to use two robust back-ends. We first optimise the graph using a Huber kernel. When

¹<http://smokebot.eu>

the first optimisation has converged, we run a second optimisation step using dynamic covariance scaling. The Huber kernel guarantees unicity of the solution, but the result is still influenced by incorrect corner associations. In order to improve alignment, we then optimise again using DCS to scale down the information matrices in edges that introduce a large error in the graph. Using DCS only tends to “switch off” edges too early, resulting in poor alignment.

2.3.2 Sensor-map alignment with region decomposition

As ILIAD addresses operation of multiple-actor fleets, possibly with heterogeneous sensor setups, it is of particular interest to align robot sensor maps created using different sensing modalities: primarily maps made from 3D or 2D range data, with sensors at different heights.

In joint work with Halmstad University, Sweden, we have developed two related methods that can be used for this purpose (see Gholami Shahbandi and Magnusson [6] and Gholami Shahbandi, Magnusson, and Iagnemma [7]). The first is a decomposition-based map alignment technique [6] to estimate an initial alignment; which is optionally followed by a non-linear map-to-map registration [7].

In order to align maps of different modalities, they first need to use the same representation. To that effect, the alignment is performed in 2D space, which is appropriate in a warehouse setting with planar floors. In other words, 3D maps are first converted to a 2D grid map; and the resulting transformation can then be applied to the full 3D map.

Since 2D grid maps can also be seen as images, one might expect that existing general-purpose image registration algorithms could be used to align them. However, as demonstrated by our experimental results [7], image registration algorithms typically do not produce satisfactory results. The reason, we believe, is that non-linear image registration requires a higher level of local information than is present in grid maps, which mostly consist of homogeneous patches of low information. There have also been several prior works on map matching that use a Hough/Radon transform, finding an alignment by decomposing it into a translation and a rotation estimation. While fast, we have shown [6] that these methods [4, 14] have trouble aligning maps that are of different types or partially distorted due to mapping inaccuracies.

In order to align two sensor maps, the method decomposes the maps into multiple regions, represented as a set of edges and faces in a doubly-connected edge list data structure. The best affine alignment is found by an exhaustive search over hypotheses generated from matching each region in one map to all regions in the other map. The decomposition procedure is illustrated in Figure 9.

To account also for cases where nonlinear alignment is necessary, In the second step, a set of control points is selected in the source map, and a gradient map is computed from the target map. The positions of the control points are then optimised using the gradient map as a fitness function together with a coherency condition that constrains the motion of each control point to be consistent with its neighbours. An example of the outcome of this alignment procedure is shown in Figure 10.

Multi-modal map alignment of partly overlapping maps is a challenging problem, as can be seen from our experimental evaluation. Still, for one indicative data set from an office environments (in the order of 30 rooms, with 14 different sensor maps), our method [6] has a 68 % success rate, compared to 3–29 % for the baseline methods that we have compared to [4, 5, 8, 11, 14]. In particular, the alignment methods are sensitive to clutter. In future work, we intend to further evaluate the consequences and possible solutions specifically in warehouse environments.

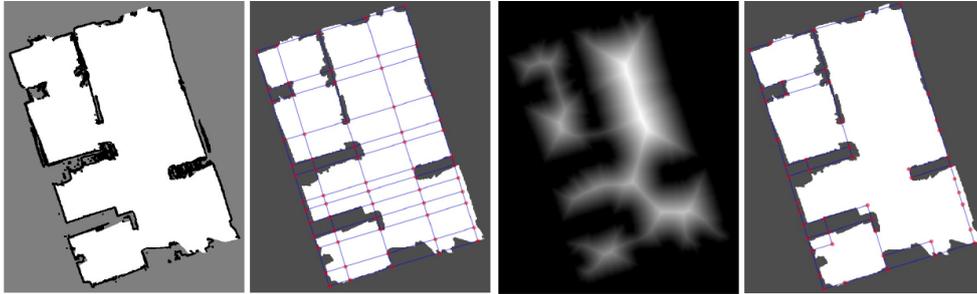


Figure 9: Map decomposition [6]. *From left to right*: an occupancy map, its original decomposition, a distance image used for pruning edges in the decomposition, and the final cleaned-up version of the arrangement.

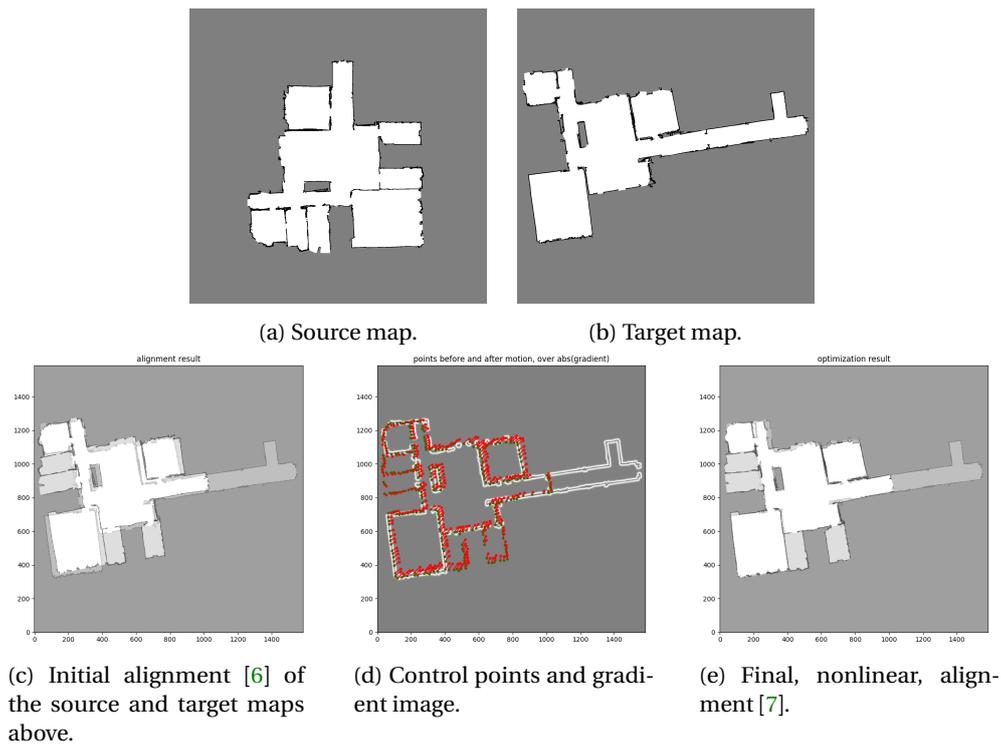


Figure 10: Example of rigid and non-rigid map alignment (see Section 2.3.2) of two sensor maps from the same indoor environment (see <https://github.com/saeedghsh/Halmstad-Robot-Maps>), but covering different parts of the environment, and each with slight mapping errors. Figure (c) shows the best rigid alignment found in step 1 of the process. Figure (e) shows the result after nonlinear alignment.

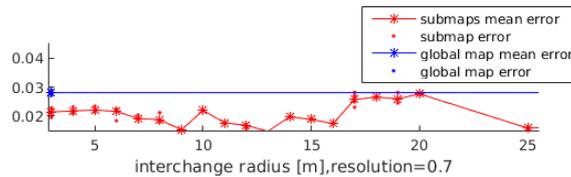


Figure 11: Mean absolute trajectory error (vertical axis, in metres) for different values of the interchange radius between submaps, at an NDT-OM grid resolution of 0.7 m. Using a single global map, the mean error is just below 3 cm. With the submap approach used in ILIAD, the mean error is approximately 2 cm, depending on the choice of interchange radius.

3 Localisation

This section presents an evaluation of localisation accuracy obtainable using the NDT-OM submapping approach presented in Section 2.1, compared to using monolithic NDT maps.

When localising in a hybrid metric–topological (HMT) map, it makes sense to carefully select *which* submap to use for localisation. Following the same line of argument as in Section 2.1, localising in a local submap generated from poses within a small region should intuitively allow for better accuracy than when using submaps created from widely different locations, or monolithic maps. A novel contribution from the work in T1.2 is a quantitative assessment of how localisation accuracy can be affected by using local (HMT) submaps compared to monolithic maps. In addition to the two methods mentioned in Section 2.1, we have evaluated a third method for selecting submaps during localisation in a pre-built map. Instead of selecting a submap based on the position in a predefined map grid or choosing the closest node position with a translation/rotation distance metric, we propose to use the recollected information of the sensor poses from where the map has been updated in order to find the most *densely updated* map at the location of the robot’s sensor. In our experimental validation, we have shown that this method improves performance by up to 40 % over a monolithic NDT-OM representation. Using data from a dairy production facility and a warehouse, the mean absolute trajectory error (ATE) is between 22 mm and 33 mm when localising by means of D2D-NDT registration to the map, compared to a reference reflector-based positioning system (see Figure 11). We plan to submit this work for publication during ILIAD’s second period. This accuracy is sufficient for the needs of the project, and is within the margins required for picking up EUR pallets from known positions using a global reference frame.

For on-line localisation we use NDT-MCL (NDT Monte-Carlo Localisation) for additional robustness (multi-hypothesis tracking). Figure 12 shows an example run through the map from Figures 1 and 3, at the point where the truck in question picks up a pallet.

4 Summary

In summary, we have developed a prototype implementation of a mapping framework that is capable of creating consistent maps from heterogeneous data — 2D and 3D lidar data, as well as rough priors from layout maps.

The structural maps used for navigation all use the hybrid metric–topological (HMT) NDT-OM maps described in Section 2.1. Maps can be created using 2D or 3D range sensors, and can also be combined with prior map data, such as layout maps. Given the map merging techniques described in Section 2.3, the different map modalities can be

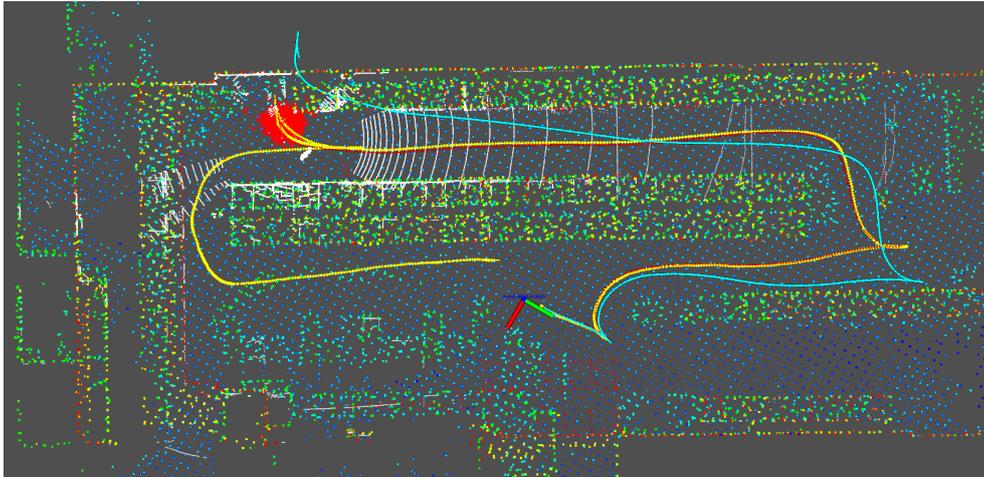


Figure 12: Illustrative example of NDT-MCL localisation. The MCL particles are shown in red (here in a circular arrangement). Although the distribution of particles may appear wide, the weighted mean that is returned as the truck's pose estimate (red line) is consistently very close to the ground truth position (yellow line). The pose estimate from odometry (blue line) quickly drifts off. At the point of this snapshot, the truck picks up a pallet.

aligned in a common reference frame, so that the fleet can share localisation information, even when using maps that have been created in isolation. Maps created from different sensors are still kept separate, although matched to the same coordinate frame, and non-rigidly aligned so that common features overlap, in case the maps have local differences in scale or appearance.

Specifically, we have demonstrated the performance of the framework in real-world warehouse environments, which feature challenges in terms of semi-static objects and perceptual aliasing. In initial tests on real-world warehouse data, the localisation accuracy (mean absolute trajectory error) is approximately 20 mm when measured against ground truth position estimates from existing infrastructure-based localisation (see Figure 11). Further tests with long-term data from end-user sites, which will better assess the performance in changing environments, will be performed during the remaining period of ILIAD.

References

- [1] Pratik Agarwal, Giorgio Grisetti, Gian Diego Tipaldi, Luciano Spinello, Wolfram Burgard, and Cyrill Stachniss. "Experimental Analysis of Dynamic Covariance Scaling for Robust Map Optimization Under Bad Initial Estimates". In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 2014. DOI: [10.1109/ICRA.2013.6630557](https://doi.org/10.1109/ICRA.2013.6630557).
- [2] Henrik Andreasson, Daniel Adolphsson, Todor Stoyanov, Martin Magnusson, and Achim J. Lilienthal. "Incorporating Ego-motion Uncertainty Estimates in Range Data Registration". In: *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*. Sept. 2017, pp. 1389–1395.

- [3] Peter Biber and Wolfgang Straßer. “The Normal Distributions Transform: A New Approach to Laser Scan Matching”. In: *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*. Las Vegas, USA, Oct. 2003, pp. 2743–2748.
- [4] S. Carpin. “Fast and accurate map merging for multi-robot systems”. In: *Autonomous Robots* 25.3 (2008), pp. 305–316.
- [5] G. D. Evangelidis and E. Z. Psarakis. “Parametric image alignment using enhanced correlation coefficient maximization”. In: *IEEE PAMI* 30.10 (2008), pp. 1858–1865.
- [6] Saeed Gholami Shahbandi and Martin Magnusson. “2D Map Alignment With Region Decomposition”. In: *Autonomous Robots* (2018). To appear. arXiv: [1709.00309 \[cs.LG\]](https://arxiv.org/abs/1709.00309).
- [7] Saeed Gholami Shahbandi, Martin Magnusson, and Karl Iagnemma. “Nonlinear Optimization of Multimodal Two-Dimensional Map Alignment With Application to Prior Knowledge Transfer”. In: *IEEE Robotics and Automation Letters* 3.3 (July 2018), pp. 2040–2047.
- [8] D. G. Lowe. “Object recognition from local scale-invariant features”. In: *IEEE ICCV*. 1999, pp. 1150–1157.
- [9] Martin Magnusson, Achim J. Lilienthal, and Tom Duckett. “Scan Registration for Autonomous Mining Vehicles Using 3D-NDT”. In: *Journal of Field Robotics* 24.10 (Oct. 2007), pp. 803–827.
- [10] Malcolm Mielle, Martin Magnusson, Henrik Andreasson, and Achim J. Lilienthal. “SLAM auto-complete: completing a robot map using an emergency map”. In: *Proceedings of the International Symposium on Safety, Security and Rescue Robotics (SSRR)*. Best student paper award. 2017.
- [11] Andriy Myronenko, Xubo Song, and Miguel Á. Carreira-Perpiñán. “Non-rigid point set registration: Coherent Point Drift”. In: *Advances in Neural Information Processing Systems*. 2007, pp. 1009–1016.
- [12] Jari Saarinen, Henrik Andreasson, Todor Stoyanov, and Achim J. Lilienthal. “3D normal distributions transform occupancy maps: an efficient representation for mapping in dynamic environments”. In: *The International Journal of Robotics Research* (Sept. 2013). DOI: [10.1177/0278364913499415](https://doi.org/10.1177/0278364913499415).
- [13] Jari Saarinen, Henrik Andreasson, Todor Stoyanov, and Achim J. Lilienthal. “Normal distributions transform Monte-Carlo localization (NDT-MCL)”. In: *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*. 2013, pp. 382–389. DOI: [10.1109/IROS.2013.6696380](https://doi.org/10.1109/IROS.2013.6696380).
- [14] Sajad Saeedi, Liam Paull, Michael Trentini, Mae Seto, and Howard Li. “Efficient map merging using a probabilistic generalized Voronoi diagram”. In: *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, pp. 4419–4424.
- [15] Todor Stoyanov, Martin Magnusson, and Achim J. Lilienthal. “Point Set Registration through Minimization of the L2 Distance between 3D-NDT Models”. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 2012, pp. 5196–5201.
- [16] Anestis Zaganidis, Li Sun, Tom Duckett, and Grzegorz Cielniak. “Integrating Deep Semantic Segmentation into 3D Point Cloud Registration”. In: *IEEE Robotics and Automation Letters* (2018). DOI: [10.1109/LRA.2018.2848308](https://doi.org/10.1109/LRA.2018.2848308).

-
- [17] Anestis Zaganidis, Martin Magnusson, Tom Duckett, and Grzegorz Cielniak. “Semantic-assisted 3D Normal Distributions Transform for scan registration in environments with limited structure”. In: *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*. Sept. 2017, pp. 4064–4069.