

Alignability maps for the prediction and mitigation of localization error

Manuel Castellano-Quero,¹ Tomasz Piotr Kucner,^{2,3} and Martin Magnusson¹

Abstract—Ensuring precise localization is key to most mobile-robot systems, not least those that are deployed in industrial settings. However, even state-of-the-art lidar-based systems may fail or lose accuracy, in particular in feature-sparse environments (e.g. transport corridors and fully stacked warehouse aisles). In this paper, we introduce *alignability maps* that spatially capture localization risk and predict where in a map localization may be less accurate. Our proposal is based on the concept of *alignability*, which represents the capacity of a given range scan to be aligned with subsequent ones. We explore different measures of alignability and evaluate them in terms of their ability to predict localization error. Our experiments show that alignability maps, which we deploy in different industrially relevant environments, serve to predict the quality of localization. Finally, we also integrate our approach in a motion planning use case. Our results demonstrate that a planner can incorporate our alignability maps as cost maps to generate paths with lower risk of localization error.

I. INTRODUCTION

Ensuring accurate localization is essential to most real-world applications of mobile robots. Localization in a map is typically performed by a Bayesian filter; e.g., Monte Carlo localization (MCL) or Kalman filters (KF). Despite demonstrated success, even state-of-the-art localization methods may work poorly or fail; e.g., due to inadequate or insufficient features [1, 2, 3]. Our aim is to be able to predict localization risk (i.e., the risk of generating inaccurate pose estimates) and account for it by taking preemptive measures; e.g., such that a planner can generate “risk-aware” paths that takes both the risk of inaccurate localization and the path length into account.

Some recent methods attempt to determine whether the current sensor view contains sufficient features to be matched to a map [3, 4]. However, the potential loss of localization may be overestimated by such measures as they do not consider the temporal filtering aspects inherent in localization methods like Monte Carlo localization. There is also a body of literature on assessing the quality of localization post-hoc [5, 6, 7], but such methods are not useful for proactively avoiding inaccurate localization.

In this work we evaluate a recent alignability metric [4] that captures the variety of surface directions computed from plane patches, and compare it to two metrics based on an estimate of the covariance of the D2D-NDT scan alignment function [8]. Our experiments show that alignability can be used as an indicator of the expected localization error and we

validate (with Granger causality tests [9]) that it also serves to *anticipate* the occurrence of errors. To enable spatial prediction, we introduce an *alignability map* in which each cell represents the expected alignability that can be obtained from different scans within that area. As a use case, we also propose to include alignability maps as cost maps for a motion planner, which generally leads to better localization.

Regarding related works, the key focus is usually put on online monitoring of the positioning error. A key example is the work by Akai et al. [10], which proposes a model of scan misalignment based on Markov Random Fields (MRFs) to measure localization error. However, this method does not anticipate risk, thus allowing only reactive behavior.

There also exist well-established approaches that measure the scan overlap [11, 12]. However, they are not directly applicable to anticipate localization quality either, since it is not generally possible to set a single threshold that can be applied in different environments [5]. Other works are aimed at a more elaborate fault detection of scan alignment. Bogoslavskyi and Stachniss [7] propose a metric based on free-space information for assessing scan matching quality, although they restrict to objects segmented in a scene. From a more generic perspective, Yin et al. [13] introduce a method based on logistic regression and metrics related to point cloud overlap and similarity for the classification of alignment errors. Also, Almqvist et al. [5] study a number of “weak” alignment classifiers and combine them with AdaBoost.

In contrast to the above, there exist recent approaches that directly estimate the probability of localization error based on the environment’s characteristics. Nubert et al. [3] propose a neural network-based method for estimating the risk of localization failure caused by a lack of environmental features. An alternative approach is the one by Nobili et al. [4], which predicts alignability based on an analysis of the geometry of point clouds. Although both of these methods aim to predict the risk of mislocalization, none are based on a spatial representation of such risk, which is our focus.

In summary, the overall contributions of the paper are: (1) a definition of alignability maps based on different measures, (2) a statistical assessment of these measures as predictors of localization error and (3), a validation of the utility of alignability maps for motion planning.

II. ALIGNABILITY MAP

In this section we introduce the different measures we have considered to represent alignability and then show how to use them to build a complete alignability map.

A. Nobili’s alignability metric

We implement the definition of alignability proposed in Nobili et al. [4] as follows. Given an input point cloud

¹AASS MRO lab, Örebro University, Sweden

²Aalto University, Finland

³Finnish Center for Artificial Intelligence (FCAI), Finland

Corresponding author manuel.castellano-quero@oru.se

This work has received funding from the Swedish Knowledge Foundation (KKS) project “NiCE” under agreement number 20200247 01 H, and the European Union’s Horizon 2020 research and innovation programme under grant agreement number 101017274 (DARKO).

\mathcal{P} , we first segment it into a set of n planar surfaces $\mathbf{P} = \{P_1, P_2, \dots, P_n\}$. For that, we apply the region growing algorithm in Rabbani et al. [14]. Each surface will be considered only if its planarity and size are adequate, according to appropriate thresholds p_{th} and s_{th} . For each $P_i \in \mathbf{P}$, we compute a normalized covariance matrix Σ . Then, we calculate planarity as $p = \lambda_s/\lambda_l$, where λ_s and λ_l are the smallest and largest eigenvalues of Σ , respectively. If $p < p_{th}$, we consider P_i as planar and then compute its area s . Finally, the plane will be accepted if $s > s_{th}$. The values for the thresholds are of 0.05 for p_{th} and of 0.01 m² for s_{th} .

From this process, we get a subset of selected planes $\mathbf{P}_s \subseteq \mathbf{P}$ that will be considered to compute alignability, as follows. First, a normal direction $\mathbf{n}_j = (x_j, y_j, z_j)$ is obtained for each point j of all the planes in \mathbf{P}_s , with N the total number of points. Then, an $N \times 3$ matrix \mathbf{M} in which each row represents a normal \mathbf{n}_j is defined. After that, a principal component analysis (PCA) is performed on \mathbf{M} , from which we get three eigenvalues, $\lambda_a \geq \lambda_b \geq \lambda_c \geq 0$. Finally, Nobili's alignability is defined as $\alpha_N = \lambda_c/\lambda_a$, where $\alpha_N \in [0, 1] \subset \mathbb{R}$. The closer to 1 this value is, the higher the variety of surface normals will be in the given scan, and therefore, the better the alignability.

B. NDT-based alignability metrics

In scan registration it is common to estimate a covariance matrix related to the pose estimate obtained from an optimization problem. Such uncertainty, which should be small for well-aligned point clouds and higher for more troublesome cases, can be seen as a form of alignability. In this work we will use the D2D-NDT [8] as a cost function, which has previously been shown to accurately indicate successful alignment [5]. Alignability can be then defined on the D2D-NDT covariance matrix as follows.

The 3D-NDT representation describes an input point cloud as a set of Gaussian probability distributions, i.e., $\text{NDT}(\mathcal{P}) = \{\mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)\}$ with $i = 1, \dots, n_{\mathcal{P}}$, where $n_{\mathcal{P}}$ is the number of distributions in the model. D2D-NDT scan registration [15] minimizes the L_2 norm between the probability distributions associated with the models of a pair of point clouds \mathcal{P}_1 and \mathcal{P}_2 . The result is a transformation \mathbf{p} , consisting of a translation \mathbf{t} and a rotation matrix R . The objective function is given by

$$f_{\text{D2D}}(\mathbf{p}) = \sum_{i=1}^{n_{\mathcal{P}_1}} \sum_{j=1}^{n_{\mathcal{P}_2}} -d_1 \exp \left(-\frac{d_2}{2} \boldsymbol{\mu}_{ij}^\top (R^\top \boldsymbol{\Sigma}_i R + \boldsymbol{\Sigma}_j)^{-1} \boldsymbol{\mu}_{ij} \right) \quad (1)$$

where $\boldsymbol{\mu}_{ij} = R\boldsymbol{\mu}_i + \mathbf{t} - \boldsymbol{\mu}_j$ and d_1 and d_2 are parameters with fixed values (see [8]). However, the estimated optimum $\hat{\mathbf{p}}$ is not perfect, and a related covariance matrix $\text{cov}(\hat{\mathbf{p}})$ is defined to account for errors. In this work we use equation (1) for a single point cloud at $\mathbf{p} = \mathbf{0}$, since we are only interested in the alignability of a single scan. Then, we approximate the covariance matrix as in [8]

$$\text{cov}(\hat{\mathbf{p}}) \approx \left(\frac{\partial^2 f_{\text{D2D}}}{\partial \mathbf{p}^2} \right)^{-1} \frac{\partial^2 f_{\text{D2D}}}{\partial \mathbf{p} \partial \mathbf{z}} \text{cov}(\mathbf{z}) \frac{\partial^2 f_{\text{D2D}}}{\partial \mathbf{p} \partial \mathbf{z}} \left(\frac{\partial^2 f_{\text{D2D}}}{\partial \mathbf{p}^2} \right)^{-1} \quad (2)$$




Fig. 1. Environment used in our simulations. (Left) Real warehouse aisles. (Right) Virtual environment with the simulated truck robot and its 3D lidar.

where \mathbf{z} represents the sensory measurements used for the computation of (1). Finally, we define the NDT-based alignability metrics as the trace and maximum eigenvalue of matrix $\text{cov}(\hat{\mathbf{p}})$, which we denote α_T and α_M , respectively.

C. Definition of alignability map

An *alignability map* \mathcal{A} is a 2D matrix-based representation of the environment in which each cell is computed as $\mathcal{A}(i, j) = \text{median}(\mathbf{a})$, where $\mathbf{a} = (\alpha_1, \alpha_2, \dots, \alpha_n)$ is a vector of n samples of alignability values (i.e., each α_i could be represented by either α_N , α_T or α_M) obtained from point clouds produced by a sensor placed within the region of the scene corresponding to cell (i, j) . To prevent biasing of the alignability score, we use robust statistics, i.e., median. We will also assume that the sensor has a 360° field of view, thus, each cell only denotes position. For the case of 3D sensors with a restricted field of view, different layers of alignability maps could be used for different orientations, and, in the case of 2D sensors, lines would be considered instead of planes. Since the alignability value usually changes gradually, we use a lower resolution for the alignability map (1 cell/m) than for the corresponding occupancy grid (10 cells/m). Finally, alignability is only computed for obstacle-free cells.

III. VALIDATION

We experimentally demonstrate the utility of the alignability map in both virtual and real environments. The virtual one is based on Gazebo and ROS [16] and it incorporates a 3D mesh of a real food warehouse in Sweden (see Fig. 1), which has been obtained on-site by using a Toyota BT SAE200 stacker truck, endowed with a Velodyne HDL-32E 3D lidar. Our virtual environment also includes this robot with the 3D lidar (see Fig. 1). We also provide qualitative results from a different use case, i.e., transport robots, which we implemented live in underground transport corridors by using a custom mobile manipulation platform Robotnik Kairos+, endowed with an Ouster OS0-128 lidar.

A. Building of alignability maps

We have produced a different alignability map of the warehouse environment for each metric by generating 3D laser scans from several points in the Gazebo model (Fig. 1). Since we aim to be able to provoke noticeable localization errors, we have cropped the maximum range of the simulated lidar to only 6.5 m, a much shorter distance than the average length of the aisles in the warehouse (approximately 30 m). The resulting maps, with a resolution of 1 meter per cell, are shown in Fig. 2(a) and (b). For space reasons we have

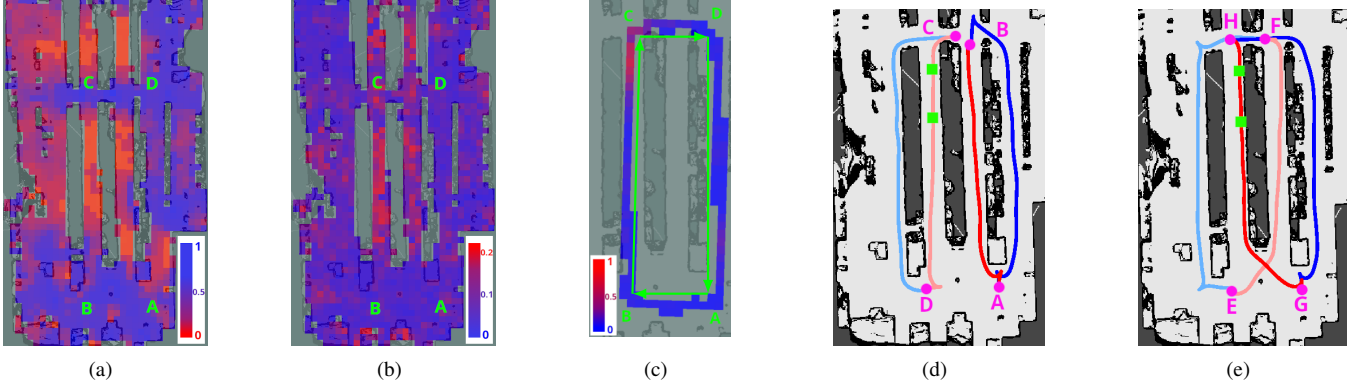


Fig. 2. (a) Alignability map based on Nobili’s metric for the food warehouse, with waypoints. (b) Alignability map of the same warehouse, based on the NDT maximum eigenvalue metric. In both maps, colors indicate values of alignability (see legends placed bottom left). (c) Error map obtained during the driving experiment. Here, colors indicate the ratio of error obtained w.r.t. the maximum one (0.8 m) according to the top left legend. (d) and (e) Motion planning experiments in Section III-C. Blue trajectories represent paths generated by considering alignability costs, while red ones do not consider them. Green squares delimit a region with extremely low alignability. (d) Experiments A-B and C-D. (e) Experiments E-F and G-H.

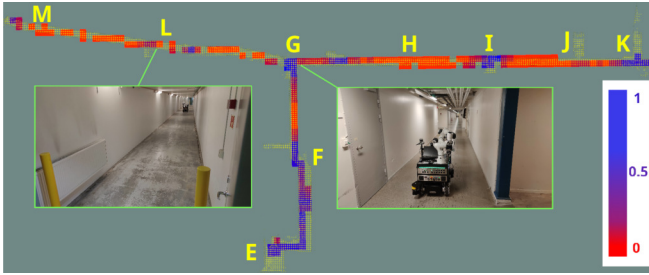


Fig. 3. Alignability map based on Nobili’s metric for the corridor environment, with waypoints. (A 3D-NDT map is shown overlaid, with yellow points.) As expected, alignability is low in the featureless corridors, but higher near the end of corridors and when passing other passages.

omitted the map based on the trace of the covariance (it looks similar to the one in Fig. 2(b)).

Regions with both a *low* Nobili alignability and a *high* value of the NDT-based metrics usually correspond to feature-sparse parts of the scene, like the long corridor B-C. Conversely, regions with a higher amount of features (e.g., corners) tend to have a higher value of the Nobili metric and a lower value of the NDT-based ones, as in the case of corridor A-D or the area between points A and B. Thus, we can affirm that the proposed maps correctly capture the variety of features in the environment, as expected.

We have also produced an alignability map based on Nobili’s metric for the real experiment (see Fig. 3), which correctly captures both feature-sparse and feature-rich regions. High values of the metric correspond to places where there are corners (e.g., around waypoints E, F, I, K or M) or junctions (waypoint G), while low values correspond to the mid parts of long corridors, like in waypoints H, J or L.

B. Alignability as a predictor of localization errors

To illustrate the relationship between alignability and localization error, we have conducted an experiment in which the robot drives along a rectangular trajectory (see Fig. 2) that is repeated three times. We have estimated the robot position \mathbf{p}_{est} by using NDT-MCL localization [17] and have calculated the error \mathbf{e} with respect to the ground-

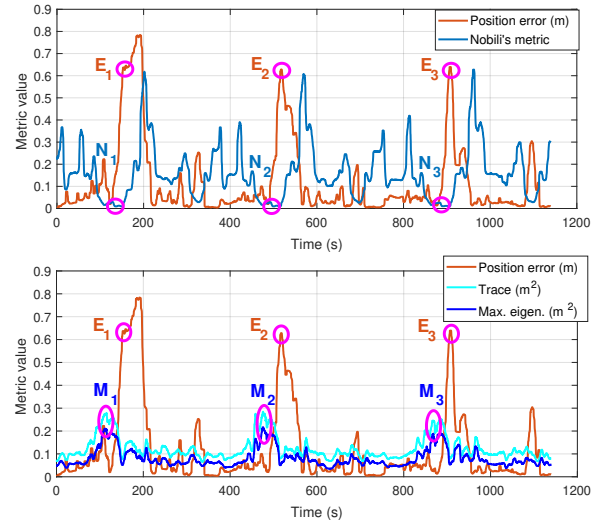


Fig. 4. Evolution of localization error and alignability metrics during the simulated experiment of section III-B. (Top) Localization error and Nobili’s metric. (Bottom) Localization error and NDT-based metrics. The three events marked as E correspond to localization errors higher than 0.6 m, the ones labeled as N to Nobili’s alignability values lower than 0.01 and the ones marked as M to NDT-based alignability values higher than 0.24 m^2 for the trace and 0.18 m^2 for the maximum eigenvalue. The time series have been filtered by using a moving median filter with 6 seconds of window size.

truth provided by Gazebo \mathbf{p}_{gt} . We have also built a map of localization errors. For each cell (i, j) of this map, the error is $\mathbf{e}(i, j) = (1/N_s) \sum_{k=1}^{N_s} \|\mathbf{p}_{est}(k) - \mathbf{p}_{gt}(k)\|$, where k denotes a sample of the estimated and ground-truth positions, out of a total number of N_s samples. The obtained map is depicted in Fig. 2(c). Despite the low alignability in corridor B-C, notice that the highest localization errors (around 0.8 m) are only obtained towards the end of it, due to the robustness of MCL over time. Thus, alignability serves to identify feature-sparse regions that might lead to higher errors in the future.

In the following we demonstrate the utility of alignability as a predictor of localization error (see Fig. 4). We have marked every time the error gets noticeably high and also all events of extremely poor alignability. In Fig. 4 we can observe that each peak of localization error always takes




Fig. 5. (Left) Boxplots for the localization errors obtained in the motion planning experiments in Figs. 2(d) and 2(e). Blue plots show the errors from the planner that uses alignability information, and red plots show the uninformed planner. (Right) View of the largest errors.

place some time after the events of poor alignability (around 18 seconds for Nobili’s metric and 39 seconds for the NDT-based metrics, on average). Thus, we can visually assess that the alignability metrics are useful to predict localization errors. We have also demonstrated this more formally through statistical Granger causality tests. The main parameter of a Granger test is number of past values (number of lags) to consider. Specifically, we have tested from 160 to 200 lags (i.e., 16–20 s) for Nobili’s metric and from 370 to 410 lags (i.e., 37–41 s) for the NDT-based ones, according to the average reaction times observed. In all Granger tests, each based on a Pearson’s χ^2 with a significance level of 0.05, the null hypothesis H_0 was confidently rejected (p-values are around 0.02 for Nobili’s metric and below $2 \cdot 10^{-18}$ for the NDT-based ones) meaning that all the alignability metrics *Granger-cause* the localization error.

We have also performed a single-query computation time study on a PC with an Intel i7-11850H at 2.5GHz and 32 GB DDR4. Mean time results for the cropped lidar range are of 0.1013 s for Nobili’s metric and 0.0232 s for the NDT-based ones. For the full lidar range, mean times are of 0.1255 s and 0.0451 s, respectively; thus, NDT-based metrics always perform better and their use is preferred in time-critical applications.

C. Alignability in motion planning

Finally, we have also carried out simulations to illustrate the impact of considering alignability in motion planning, which we have implemented by using the standard `move_base` ROS package [16]. Then, we have produced a binary version of the alignability map in Fig. 2(a) (by using a threshold of 0.02 for Nobili’s metric, considered very poor alignability) and employed it as a cost map for the motion planner. This is done to force the motion planner to avoid regions with extremely low alignability.

We have performed four experiments in which the robot follows a path between two waypoints. Each experiment has been repeated two times, leading into different paths (see Figs. 2(d) and 2(e)). Those ignoring alignability costs (red) are the shortest possible, while the other ones (blue) tend to avoid the two central corridors, which exhibit lower alignability (see Figs. 2(a) and 2(b)). We have also measured the localization errors for each path (see Fig. 5), finding that those generated by considering alignability costs always suffer from lower errors.

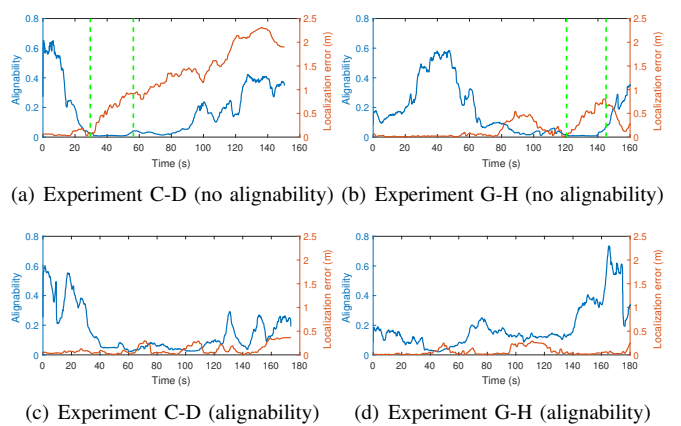


Fig. 6. Alignability and localization error measured over time for motion planning experiments C-D and G-H. Results are shown for paths that consider alignability as a cost as well as for those that do not. Green dashed lines in plots (a) and (b) delimit the time interval corresponding to navigation within the region with extremely low alignability marked in Figs. 2(d) and 2(e).

Finally, we have also measured alignability for all the experiments. For the sake of brevity, we only show the results for C-D and G-H (see Fig. 6). In general, paths made by the plain motion planner (Figs. 6(a) and 6(b)) suffer from lower values of alignability and much higher errors compared to the others (Figs. 6(c) and 6(d)), as expected. However, it is interesting to analyze why localization is much worse in experiment C-D compared to G-H for paths that do not consider alignability, even when they follow the same corridor but in opposite directions.

In experiment C-D, the robot first faces an extremely low alignability region for approximately 25 s (see green dashed lines in Fig. 6(a)) and localization error increases from 0.07 to 0.92 m. Later on, although alignability slightly increases, it is still low for approximately 20 more seconds, and the error gets much worse. In contrast, when the path is followed the other way round (Fig. 6(b)), localization error can be recovered before entering the region with extreme low alignability (green lines in Fig. 6(b)) and, although such error increases up to 0.8 m, it is low enough to recover when exiting the corridor.

IV. CONCLUSIONS AND FUTURE WORK

We have introduced the notion of alignability map, which aims to represent the risk of localization error by considering the geometry of the environment. We have validated this map in both virtual and real scenarios, finding that it correctly captures the amount and variety of features present. We have also shown, through statistical testing, that the alignability metrics considered can be used to anticipate localization error and therefore that our map can be used to predict it. Finally, we have demonstrated that using alignability maps in motion planning serves to decrease localization error in navigation.

In future works, we plan to extend our approach so that it includes different sources of information for estimating localization risk (e.g., dynamics of obstacles). Also, we plan to validate our map for a wider variety of settings, e.g., for outdoor and unstructured environments.

REFERENCES

- [1] J. Kim, J. Park, and W. Chung. “Self-Diagnosis of Localization Status for Autonomous Mobile Robots”. In: *Sensors* 18.9 (Sept. 2018), p. 3168.
- [2] S. A. Kurz Gerhard and Scherer, P. Biber, and D. Fler. “When Geometry Is Not Enough: Using Reflector Markers in Lidar SLAM”. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2022.
- [3] J. Nubert, E. Walther, S. Khattak, and M. Hutter. “Learning-based Localizability Estimation for Robust LiDAR Localization”. In: *2022 IEEE International Conference on Intelligent Robots and Systems (IROS)*.
- [4] S. Nobili, G. Tinchev, and M. Fallon. “Predicting Alignment Risk to Prevent Localization Failure”. In: *2018 IEEE International Conference on Robotics and Automation (ICRA)*. 2018, pp. 1003–1010.
- [5] H. Almqvist, M. Magnusson, T. P. Kucner, and A. J. Lilienthal. “Learning to detect misaligned point clouds”. In: 35.5 (Aug. 2018), pp. 662–677.
- [6] D. Adolfsson, M. Castellano-Quero, M. Magnusson, A. J. Lilienthal, and H. Andreasson. “CorAl: Introspection for robust radar and lidar perception in diverse environments using differential entropy”. In: *Robotics and Autonomous Systems* (2022), pp. 104–136.
- [7] I. Bogoslavskyi and C. Stachniss. “Analyzing the quality of matched 3D point clouds of objects”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2017, pp. 6685–6690.
- [8] T. Stoyanov. “Reliable Autonomous Navigation in Semi-Structured Environments using the Three-Dimensional Normal Distributions Transform (3D-NDT)”. PhD thesis. Örebro University, 2012.
- [9] C. W. J. Granger. “Investigating Causal Relations by Econometric Models and Cross-spectral Methods”. In: *Econometrica* 37.3 (1969), pp. 424–438.
- [10] N. Akai, Y. Akagi, T. Hirayama, T. Morikawa, and H. Murase. “Detection of Localization Failures Using Markov Random Fields With Fully Connected Latent Variables for Safe LiDAR-Based Automated Driving”. In: *IEEE Transactions on Intelligent Transportation Systems* (2022), pp. 1–13.
- [11] A. Segal, D. Haehnel, and S. Thrun. “Generalized-ICP”. In: *Robotics: Science and Systems V*. Robotics: Science and Systems Foundation, June 2009.
- [12] Q. Liao, D. Sun, and H. Andreasson. “Point Set Registration for 3D Range Scans Using Fuzzy Cluster-Based Metric and Efficient Global Optimization”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43.9 (Sept. 2021), pp. 3229–3246.
- [13] H. Yin, L. Tang, X. Ding, Y. Wang, and R. Xiong. “A failure detection method for 3D LiDAR based localization”. In: *2019 Chinese Automation Congress (CAC)*. 2019, pp. 4559–4563.
- [14] T. Rabbani, F. A. van den Heuvel, and G. Vosselman. “Segmentation of Point Clouds using Smoothness Constraints”. In: *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 36.5 (2006), pp. 248–253.
- [15] T. Stoyanov, M. Magnusson, and A. J. Lilienthal. “Point Set Registration through Minimization of the L2 Distance between 3D-NDT Models”. In: St. Paul, MI, USA, May 2012, pp. 5196–5201.
- [16] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng. “ROS: an open-source Robot Operating System”. In: *ICRA workshop on open source software*. 2009.
- [17] J. Saarinen, H. Andreasson, T. Stoyanov, and A. Lilienthal. “Normal Distribution Transform Monte-Carlo Localization (NDT-MCL)”. In: *Proc. IEEE/RSJ Int. Conf. on Intell. Robots and Syst.* 2013, pp. 382–389.