

CFEAR Radarodometry - Conservative Filtering for Efficient and Accurate Radar Odometry

Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, Henrik Andreasson

Abstract—This paper presents the accurate, highly efficient, and learning-free method CFEAR Radarodometry for large-scale radar odometry estimation. By using a filtering technique that keeps the k strongest returns per azimuth and by additionally filtering the radar data in Cartesian space, we are able to compute a sparse set of oriented surface points for efficient and accurate scan matching. Registration is carried out by minimizing a point-to-line metric and robustness to outliers is achieved using a Huber loss. We were able to additionally reduce drift by jointly registering the latest scan to a history of keyframes and found that our odometry method generalizes to different sensor models and datasets without changing a single parameter. We evaluate our method in three widely different environments and demonstrate an improvement over spatially cross-validated state-of-the-art with an overall translation error of 1.76% in a public urban radar odometry benchmark, running at 55 Hz merely on a single laptop CPU thread.

I. INTRODUCTION

Estimating odometry in large-scale GNSS-denied environments is an essential part in any robust autonomous driving or robotics system. Popular systems are based on lidar [1]–[4], vision [5]–[7], interoceptive sensors such as IMU or wheel odometry, or a combination of various sensors [8]–[10].

Lidar is an accurate and fast long-range sensor but struggles to perceive the environment in presence of dust or in harsh weather [11]. Vision is moderately accurate and inexpensive but operates at shorter range and is sensitive to illumination changes or weather conditions. In contrast, radar is a long-range sensing modality that operates well in all weather conditions and is hardly affected by smoke and dust. As radars have recently become compact enough for autonomous systems, it is desired to utilize their excellent properties for robust localization.

State-of-the-art radar odometry systems put great effort into overcoming sensor-specific challenges such as speckle noise, ghost objects and false positives using deep learning methods to remove noise [12], predict robust key points [13], and through robust data association algorithms [14]–[16]. Scan matching is then based on sparse and efficient key-points [13]–[16] or dense radar images [12], [17].

While learning-based methods perform well on locations included in the training set, they provide limited generalization and, accordingly, their accuracy deteriorates at

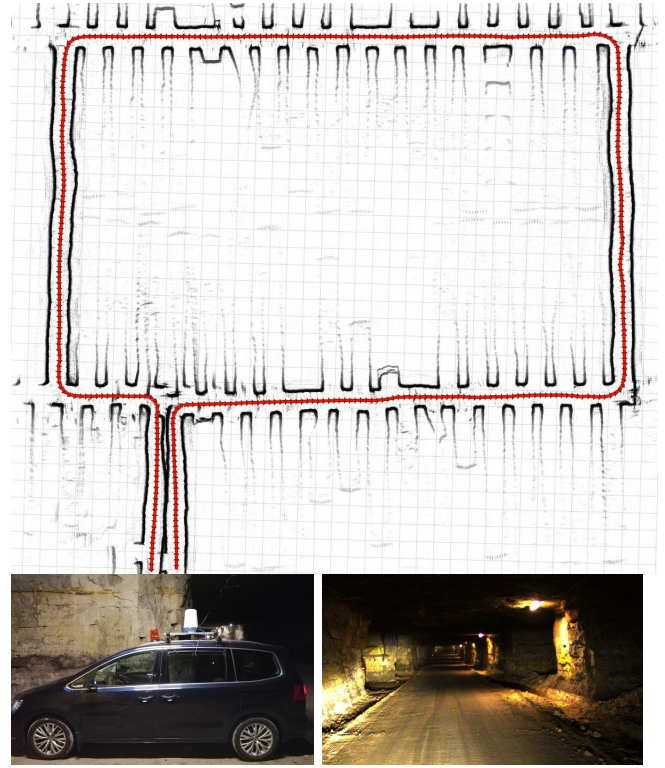


Fig. 1: Example output using the proposed method. Incremental odometry estimated from a roof-mounted radar in an underground mine using the proposed method with no additional sensor. Parameters are tuned for a different sensor in an urban environment but generalize well. The accumulated error is roughly 15m over a distance of 1235m driving (1.2%). Our system is demonstrated at: <https://youtu.be/np-ceNrAcNY>.

previously unseen locations [12] (despite being trained with data acquired from over 200 km of driving) [12], [13].

We present Conservative Filtering for Efficient and Accurate Radar Odometry (CFEAR Radarodometry), a learning free method for efficient and accurate large-scale odometry estimation. CFEAR Radarodometry filter radar data that creates a sparse representation in which the environment is modeled as a set of oriented surface points as depicted in Fig. 4b. By keeping only the k strongest power returns per measurement azimuth above the expected noise level, a conservative but sufficient set of observation points can be maintained while removing most radar noise. The filtered radar returns can be used to create a sparse set of surface points and normals that models the underlying geometry of important surface features. Scan registration is efficiently carried out by minimizing a point-to-line metric between correspondences, using Huber loss to achieve robustness to outliers.

The authors are with the MRO lab of the AASS research centre at Örebro University, Sweden. E-mail: Daniel.Adolfsson@oru.se

This work has received funding from the Swedish Knowledge Foundation (KKS) project “Semantic Robots” and European Union’s Horizon 2020 research and innovation programme under grant agreement No 732737 (ILIAD).

The central contribution of this paper is an efficient, accurate and robust method tailored for Frequency-Modulated Continuous Wave (FMCW) radar in large-scale environments, reaching state-of-the-art accuracy on a public benchmark data-set. Our results show that CFEAR Radarodometry method performs well in widely different environments (urban, forest, tunnel/mine) and can be used with different sensor models without changing a single parameter. It requires no training and runs on a single-core laptop CPU at 55 Hz.

II. RELATED WORK

A. Filtering radar data

Raw radar scans are often converted to point clouds using filtering techniques that distinguish between targets and background noise. Due to the challenging noise characteristics of radar data, various approaches have been investigated to remove noise. The traditionally most widely used filtering technique, Constant False-Alarm Rate (CFAR) [18], cannot be applied directly, as it assumes to know the noise distribution *a priori*, which is not the case with FMCW radar, especially when used in different environments. Moreover, CFAR requires manual tuning of multiple parameters [15], [16].

Another method is to simply set a fixed noise/signal threshold [19]. However, doing so comes with the risk to remove important information since the noise background level varies greatly in different settings.

Deep learning methods [12], [20], [21] have also been proposed for filtering but they require offline training and labeled data-sets as in [12]. Wetson et al. [20] used lidar data for labeling applying a static world assumption. Aldera et al. [21] used sensor coherence and wheel odometry for training. In Marck et al. [22], only the strongest power return per azimuth is kept. In Fig. 7b we show that keeping only the strongest returns per azimuth makes the odometry estimation inaccurate.

In this work, we extend the work of Marck et al. [22] by including only the strongest k returns in a given beam that exceed the expected noise level z_{\min} . Section V-A presents an experimental evaluation of the effect of k on the performance of radar odometry. The difference between the proposed CFEAR filtering and the classical CFAR method can be seen in Fig. 2.

B. Radar odometry

Over the years, various methods have been proposed for radar SLAM [23]–[25]. Recent radar odometry methods can be divided into sparse methods that attempt to estimate key points or compute corresponding features from radar images, and dense methods that match a denser set of radar returns. Holder et al. used ICP to align point clouds acquired from a front-facing radar [26] however the method relies on additional sensors.

Mielle et al. [19] demonstrated radar odometry with off-the-shelf frameworks for range-based mapping in an indoor environment, using a fixed-threshold filtering method. However, the validation was done in a single, small-scale




Fig. 2: Typical power/range plot for a single azimuth reading. Here we use $k = 12$ for conservative filtering that reduces false positives at the cost of missing some landmarks by simply limiting the maximum number of detections per azimuth. Accordingly, z_{\min} can be selected low make the filter sensitive without drastically increase the number of false detections. In contrast, CFAR returns more spurious detections and is hard to tune.

environment, and the fixed-threshold filtering has been found to be difficult to generalise to other environments.

Cen et al. presented sparse methods for ego-motion [15], [16]. A landmark extraction method detects peaks while removing multipath reflections. Data association and outlier rejection is carried out by finding the largest common subset of landmarks between subsequent scans. Landmarks are detected only in the polar space using 1d (power/range) signals. In contrast, our method additionally filters range data in Cartesian space and we obtain robustness to outliers via a Huber loss. Hong et al. [14] proposed a full radar SLAM framework. The odometry method extract SURF feature descriptors without filtering the raw radar data and uses robust data association similarly to Cen [15]. Both Cen and Hong estimate the motion by minimizing the (weighted) squared error between correspondences using SVD (Single Value Decomposition), the performance of their methods is included in our evaluation.

Barnes et al. presented two learning based methods for radar odometry. In [12] a mask is learned for filtering dynamic objects and noise and matching is carried out with a dense correlation-based approach. Barnes and Posner later proposed a sparse method that learns keypoints, scores and descriptors for odometry and loop closure detection [13]. Keypoints are associated via a dense search and matched by minimizing the squared distance using SVD. Both methods [12], [13] require supervision from a previously existing positioning system and an extensive amount of data collected over 200 km driving. Spatial cross-validation shows that the accuracy drops by a factor of 2.4 when estimating odometry in new streets not included in the training set, yet in nearby locations within the same city [12]. It is currently not known how well [13] generalizes to new locations. Our method does not rely on training or a previously existing positioning system. In the evaluation we compare these methods to ours




Fig. 3: Conversion between polar (left) $Z_{m \times n}$ and Cartesian (right) space. Each row and column in the polar image correspond to a range and angle (a, d) in the Cartesian image.

and qualitatively demonstrate how our method generalizes to widely different environments.

Park et al. [17] used dense matching by applying the Fourier Mellin transform (FMT) to Cartesian and log-polar radar images and estimating the relative motion from phase correlation in a course-to-fine fashion. Their dense method yields poor run time performance and has not been compared to the state-of-the-art.

Radar odometry has also been improved by accounting for motion distortion [27], using multiple Doppler radars [28] and by introspectively detecting registration failures [29].

Finally, radar based localization with respect to a map have also been proposed in [10], [30]–[33]. However, these methods are out of scope as we focus on the radar odometry problem without a prior map.

III. ROTATING RADAR SENSING

Our method CFEAR Radarodometry is tailored for, but not specific to, a rotating FMCW radar without Doppler information. At the end of each full 360° rotation, the sensor outputs a full power reading $Z_{m \times n}$ in polar form, where m is the amount of azimuths per sweep and n is the amount of range bins per azimuth.

Each azimuth and range bin (a, d) in $Z_{m \times n}$ stores a power return and can be converted from polar to Cartesian space as seen in Fig. 3 using Eq. (1).

$$\mathbf{p} = \begin{bmatrix} d\gamma \cos(\theta) \\ d\gamma \sin(\theta) \end{bmatrix}, \quad (1)$$

where $\theta = 2\pi a/m$ and γ is the range resolution set to 4.38cm or 17.5cm in our experiments.

IV. CFEAR RADARODOMETRY

We argue that it is hard to find stable keypoints from peaks per azimuth due to large beam width, and radar sensing artifacts. In CFEAR Radarodometry, we compute a scan representation in two steps. First, only the k strongest returns with power exceeding z_{\min} are selected at each azimuth and converted into Cartesian space. Second, oriented surface points are computed in Cartesian space by analyzing local point distributions at discrete intervals. This creates a sparse surface representation from the originally noisy radar data, which rejects artifacts such as multi-path reflections in favor of important landmarks.

A. k -strongest filtering

The typical behavior of k -strongest filtering compared to the classical method CFAR can be seen in Fig. 2. Keeping only the k strongest returns exceeding the power z_{\min} reduces false positives at the expense of eventually missing landmarks with weaker returns. However as shown by our evaluation depicted in Fig. 7b, missing weaker landmarks has a negligible impact on localization accuracy. This indicates that a conservative filtering technique is preferable for efficient radar odometry estimation.

B. Estimating oriented surface points


Given the filtered point set \mathcal{P}_f we aim to model the underlying surfaces in the environment by a set of oriented points; i.e., surface points and normals $\{\mu_i, \mathbf{n}_i\}$. This is done by first downsampling the point cloud \mathcal{P}_f to \mathcal{P}_d using a grid with side length r/f m, with r the distance in which we expect to find sufficient amount of points to estimate the normal, and f a re-sampling factor that adjusts the density of the final representation. For each grid cell only the centroid is kept. For each point in the downsampled point cloud $p_i \in \mathcal{P}_d$, all neighbors of p_i in \mathcal{P}_f within a radius r (same r as above), are used to estimate the point distribution by the sample mean and covariance $\{\mu_i, \Sigma_i\}$. Ill-defined distributions where the condition number $\kappa(\Sigma_i) = \lambda_{\max}(\Sigma_i)/\lambda_{\min}(\Sigma_i) > 10^5$, or where distributions are estimated from less than 6 points are discarded. This works as an additional filtering step that removes remaining false detections after k -strongest filtering and ensures that surface points are estimated from a sufficient number of points and from multiple azimuths. The surface normal \mathbf{n}_i is obtained from the smallest of the two eigenvectors of Σ_i . The full scan representation \mathcal{M} is obtained by pairing all surface points and normals $\mathcal{M} = \{\mu_i, \mathbf{n}_i\}$. An example of reconstructed surfaces after applying the k -strongest filter in comparison to CFAR can be seen in Fig. 4.

C. Multiple keyframe scan registration


We employ a sliding window keyframe approach where the latest scan \mathcal{M}^t is registered to the s most recent keyframes $\mathcal{M}^{k_{1..s}}$. When the estimated relative pose exceeds a minimum rotation or translation to the latest keyframe, a new keyframe is stored from the current pose and point cloud $\{\mathcal{M}^{k_s}, \mathbf{T}^{k_s}\} \leftarrow \{\mathcal{M}^t, \mathbf{T}^t\}$ and the oldest keyframe is removed. This ensures that drift is not introduced when the robot is stationary. Simultaneously registering to multiple keyframes can improve registration accuracy by imposing additional constraints, and improve robustness in cases when the registration is under-constrained or scan overlap drastically changes. Large changes in overlap can occur due to occlusions from stationary or dynamic objects or due to non-planar road conditions [29].

We formulate the registration problem by minimizing a scan-to-keyframes ($s2ks$) cost function

$$\arg \min_{\mathbf{x}} f_{s2ks}(\mathcal{M}^{k_{1..s}}, \mathcal{M}^t, \mathbf{x}), \quad (2)$$



(a) CFAR keeps an excessive amount of points (black), including false positives, and surface points are poorly estimated.



(b) CFEAR Radarodometry: environment is represented by oriented surface points (red). The k -strongest filtered points (fewer false positives) are shown in black.

Fig. 4: Comparison of filtering methods, and oriented surface points.




Fig. 5: The latest scan \mathcal{P}^t is registered to the $s = 2$ previous key-frames. This can improve accuracy due to noise and robustness due to occlusions.

where we seek the transformation parameters $\mathbf{x} = (x \ y \ \theta)$ that best align the current scan \mathcal{M}^t with the key-frames $\mathcal{M}^{k1..s}$ jointly

$$f_{s2k}(\mathcal{M}^{k1..s}, \mathcal{M}^t, \mathbf{x}) = \sum_{n=1}^s f_{p2l}(\mathcal{M}^{k_n}, \mathcal{M}^t, \mathbf{x}). \quad (3)$$

The cost of aligning the latest scan \mathcal{M}^t to a single key-frame \mathcal{M}^k is given by the *point-to-line* (p2l) function f_{p2l} in Eq. (4).

$$f_{p2l}(\mathcal{M}^k, \mathcal{M}^t, \mathbf{x}) = \sum_{\forall i \in \mathcal{M}^t} \mathcal{L}_\delta(\mathbf{n}_j^k \cdot (\mathbf{R}_\theta \boldsymbol{\mu}_i^t + \mathbf{t}_{x,y} - \boldsymbol{\mu}_j^k)), \quad (4)$$

where each surface point $i \in \mathcal{M}^t$ is associated with the closest surface point $j \in \mathcal{M}^k$ within a radius r if the angle between the surface normals is within a tolerance θ_{\max} . \mathbf{R}_θ and $\mathbf{t}_{x,y}$ is the rotation matrix and translation vector created from the optimization parameters. An intuition for the score is visualized in Fig. 6.

Using the Huber loss \mathcal{L}_δ , defined in Eq. (5), makes the cost less sensitive to outliers [34]. This is done by piecewise reshaping the cost function to increase quadratically for small




Fig. 6: \mathcal{M}^t (blue) registered with \mathcal{M}^k (red). Most blue surface points (except one) have a correspondence with normals within θ_{\max} tolerance. The cost is given by the point-to-line distance visualized by black arrows.

values and linearly for larger values.

$$\mathcal{L}_\delta(s) = \begin{cases} \frac{1}{2}s^2, & \text{if } s \leq \delta \\ \delta(|s| - \frac{1}{2}\delta), & \text{otherwise.} \end{cases} \quad (5)$$

We solve (2) using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) line search method with gradients obtained through automatic differentiation in the Ceres non-linear optimization framework [35].

D. Accounting for motion

At each frame the previous pose \mathbf{T}^{t-1} and velocity estimate was used to predict a starting point \mathbf{x}_{init} for the optimization by linear extrapolation. The velocity was also used to compensate point cloud \mathcal{P}_f from motion distortion. This is done by transforming all points into the time half way through the sweep, assuming that velocity is constant.


V. EVALUATION

In this section we investigate the performance of our method quantitatively on the public Oxford Radar RobotCar Dataset [36], in an urban environment, and qualitatively on two data sets that we collected in non-urban environments (*Kvarntorp mine* and *Volvo test track*), to demonstrate how the method generalises without site and sensor specific parameter tuning. All experiments were carried out on an i7-7700HQ 2.80 GHz laptop CPU, running in a single thread.


A. Parameters

The full set of parameters can be seen in Table I. We tuned the parameters to an urban environment using the sequence 16-13-42 in the Oxford dataset. This sequence is not included in the evaluation set, with the intention to avoid fitting parameters to the specific traffic conditions in one of the evaluated sequences. After the parameter tuning, the exact same parameters were used through all experiments, despite widely different environments and although different radar models with different range resolutions were used.

A brief evaluation of the most important parameters (performed on the Oxford dataset) can be seen in Fig. 7. The results in Fig. 7a show that the accuracy can be improved by increasing s (submap keyframes) at the expense of runtime performance. Fig. 7b shows the effect of varying the number k of returns in our filtering method. Using too few ($k < 12$) or too many ($k > 35$) points negatively impact accuracy, but within the wide range $12 \leq k \leq 35$ the method is largely insensitive to k .



(a) Increasing the amount of key-frames s reduces the translation error but requires additional processing time. Resample factor is set to $f = 3$.



(b) Within a wide range of values of k , CFEAR Radarodometry is insensitive to the exact value of k . A too small or too large value, however, reduces translation and rotation accuracy.

Fig. 7: Accuracy/computation time using different parameters.

Parameters	Value
k	12
Resample factor f	1
z_{\min}	55
Resolution r	3.5
θ_{\max}	30
Huber loss \mathcal{L}_δ	0.1
Key-frame (dist[m]/rot[deg])	1.5/5°
s	3
min. sensor distance [m]	5
max. sensor distance [m]	100

TABLE I: Full list of parameters. These are used for all datasets and sensors, except when explicitly mentioned in Sec. V-C.1.

B. Quantitative evaluation in urban environment

We provide a quantitative evaluation of the odometry accuracy and performance in a large-scale environment with ground truth positioning. The odometry accuracy was measured as proposed in the KITTI odometry benchmark [37] using the tool [38] that computes the average translation error (%) and rotation error (deg/m) over all sub sequences between $\{100, 200, \dots, 800\}$ m.

We used the public Oxford Radar RobotCar Dataset which was collected with a roof-mounted Navtech CTS350-X FMCW radar running at 4 Hz, configured with a range resolution of $\gamma = 4.38$ cm. The dataset contains 32 traver-



Fig. 8: Incremental radar odometry using the proposed method over a 10 km sequence from the Oxford dataset. The first and final pose estimate is highlighted and gives an indication of the total aggregated drift.

sals through a 10-km urban environment in varied weather, lighting and traffic conditions.

In order to compare our results, we selected the same sequences that were used in the evaluation by Hong [14], both for quantitative evaluation as presented in Tab. II and the visualization in Fig. 9. Learning-based methods are compared using their SCV (Spatial Cross Validation). With SCV, test and training datasets are taken from different parts of the environment, rather than selecting the training data from different traverses in the same locations. Without SCV, test and training datasets are not independent since test and training points may be neighbors [39].

As can be seen in Tab. II, CFEAR Radarodometry produces the lowest per-sequence and mean SCV error with a translation and rotation error of 1.76% and 0.5 deg/100m, followed by Barnes Dual Cart [12] (2.78%), Hong odometry [14] (3.11%), and Cen [15] (3.72%). *Barnes Dual Cart* [12] achieves greater mean accuracy (1.16% error) but only when training and evaluating on the same spatial location. When the authors evaluated their methods on places not seen in the training set, the error increased to from 1.76% to 2.78%. CFEAR Radarodometry provides a 14% lower mean odometry error compared to [13] without SCV. The method does not report SCV error in the original publication and it is unclear how well the method generalizes to new nearby locations or environments.

We therefore argue that CFEAR Radarodometry method improves state-of-the-art accuracy for sparse methods, but also improves state-of-the-art among all radar odometry methods conditioned on that the evaluation is not biased with spatially overlapping training and evaluation data.

C. Qualitative evaluation

In order to assess the generality of the method, additional non-urban data was collected. We perform a qualitative eval-




Fig. 9: Evaluated Oxford sequences using CFEAR Radarodometry. For qualitative comparison to other methods we refer to [14].

Method	Evaluation	resolution	Sequence	10-12-32	16-13-09	17-13-26	18-14-14	18-15-20	10-11-46	16-11-53	18-14-46	Mean	Mean SCV	runtime (ms)
Visual Odometry [40]	[12]	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	3.78/0.01	N/A	
SuMa (Lidar) [4]	[14]	N/A	1.1/0.3	1.2/0.4	1.1/0.3	0.9/0.1	1.0/0.2	1.1/0.3	0.9/0.3	1.0/0.1	1.16/0.3	1.03/0.3		
Cen [15]	[12]	0.175	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	3.72/0.95	3.63/0.96	
Robust Keypoints [13]		0.346	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	2.05/0.67*	N/A	
Barnes Dual Cart [12]		0.043	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	1.16/0.3*	2.784/0.85	
Hong odometry [14]		0.043	2.98/0.8	3.12/0.9	2.92/0.8	3.18/0.9	2.85/0.9	3.26/0.9	3.28/0.9	3.33/1	3.11/0.9	3.11/0.9	3.11/0.9	
CFEAR Radarodometry (ours)		0.043	1.64/0.48	1.86/0.52	1.66/0.48	1.71/0.49	1.75/0.51	1.65/0.48	1.99/0.53	1.79/0.5	1.76/0.50	1.76 ± 0.12/0.50 ± 0.02		18 ± 0.4


TABLE II: Evaluation on 8 sequences with different methods and sensor modalities on the Oxford Radar RobotCar dataset [36]. Results are given in (% translation error / deg/100 m). Note the difference between columns "Mean" and "Mean SCV". In "Mean" column, methods marked * cannot be compared directly as these are trained and evaluated on the same spatial location. The most relevant number for comparison is instead "Mean SCV" which ensures that test and training data are not correlated by using spatial cross validation.

uation of radar odometry in two different industrial scenarios using another Navtech FMCW radar model, CIR154XH. For these datasets, the sensor was configured with range resolution $\gamma = 17.5\text{cm}$, with 400 azimuth bins. Ground truth positioning is missing in these datasets. However, we estimate the error between the first and final scan by comparing the offset between duplicated landmarks in the map which gives a rough indication of odometry drift.

1) *Evaluation in semi-structured outdoor environment – Volvo CE*: The first dataset was collected from a test track for wheel loaders and dump trucks in a partly forest environment driven at an approximate speed of 10km/h. The track consists mainly of gravel roads at various slopes and the radar sensor was mounted on the driver cabin of a Volvo wheel loader (see Fig. 10). The larger loop in Fig. 10 is 1150

meters long. After completing the larger gravel-road loop, two smaller loops were traversed on uneven paths inside the forest reaching a total trajectory length of 1605 meters. We expect this highly cluttered environment to be challenging for CFEAR Radarodometry that attempts to reconstruct and utilize surface normals.

The estimated odometry and a corresponding map, obtained without changing parameters from the urban dataset, is depicted in Fig. 11a. The maximum translation error is 15 m (1.3%) over the large loop. We notice that $z_{\min} = 55$ is too low in combination with the sensor model and cluttered forest environment, however, *k-strongest filtering* limits the amount of noise included in the filtered point cloud. We also note that motion compensation can have a negative impact in this dataset. For that reason, we made an attempt to improve



(a) Vehicle test track. A wheel loader with radar starts inside the white tent (right) and is driven in a large loop on a gravel road and then in two smaller loops on an uneven path in the forest. The full driven trajectory (seen in red) is imported from Fig. 11b. The locations of the photos in figure (b) and (c) are indicated along the trajectory.




(b) Driving on a narrow uneven path in the forest. (c) Driving on a gravel road uphill.


Fig. 10: Overview of semi-structured VolvoCE dataset.

the odometry by increasing the power threshold $z_{\min} = 85$, by switching off motion compensation, and by increasing surface point density to $f = 3$. These parameters reduced the estimated error to 2.5 m (0.2%) as depicted in Fig. 11b. For comparison, we plot the odometry with updated parameters but with a single submap keyframe $s = 1$, this increases the error to 5 m (0.4%) as seen in Fig. 11c. In line with our previous results, increasing the number of submap keyframes reduces the drift. This can be seen by the relative increase in map blur in Fig. 11c ($s = 1$) compared to Fig. 11b ($s = 3$). We believe this is because additionally constraining the registration reduces pose estimate noise and increases robustness due to uneven driving conditions, especially in the forest.


2) *Evaluation in mine – Kvarntorp*: In another experiment we mounted the Navtech CIR154XH radar on a car and drove through a 1235-m long loop in an underground mine, at an approximate speed of 10 km/h. The environment is depicted in Fig. 1. The solid walls are highly radar-reflective and the interaction between radar, mountain walls and the car gives rise to a large amount of strong multi-path reflections. In



(a) Full estimated trajectory (red) and map (black), obtained reusing parameters tuned for the Oxford urban dataset. The maximum estimated large loop error is 15 m (1.3%). Grid size (10×10) m.



(b) Odometry improved by tailoring parameters to the VolvoCE dataset (motion compensation switched off, $z_{\min} = 85$, $f = 3$). The Map is blurred around revisited regions of the large loop due to odometry drift but fairly sharp around the small loop. We estimate the maximum large loop position error to be 2.5 m (0.2%).



(c) Full estimated trajectory with a single key-frame. The maximum large loop error is estimated to be 5 m (0.4% error). The map around the large and small loop is more distorted compared to 11b.

Fig. 11: Qualitative evaluation in VolvoCE dataset using various parameters.

contrast to the other evaluated datasets, the tunnels in parts of this environment have limited features that can constrain the registration in the longitudinal direction.

Reusing the parameters tuned for Oxford dataset, we found that our method was able to reconstruct an accurate map of the environment with an estimated final position error of 15 m (1.2%) translation error. There are no visible errors in the passages between the top and bottom aisles in Fig. 1, which is an indication that the method works well, even in a feature poor environment.

VI. CONCLUSIONS

In this paper we have presented CFEAR Radarodometry, a radar odometry method that conservatively filters radar data using the k strongest returns and represents the environment by a sparse set of oriented surface points. The oriented surface points can be used to efficiently and accurately perform scan registration by minimizing a point-to-line metric, obtaining outlier robustness using a Huber loss. Our method is learning-free, requires no training data and can operate well in widely different environments and on different sensor models without changing a single parameter. Best odometry performance is, however, obtained by parameters tailored to the specific sensor and environment. We found that drift can be further reduced by registering the latest scan to a history of key frames jointly, at the expense of computational time. Despite this, we were able to achieve 55 Hz single-thread performance on a laptop CPU, improving SCV state-of-the-art with 1.76% overall translation error in a public urban dataset. Future work will focus on evaluating different loss functions and weighting residuals by surface uncertainty.

REFERENCES

- [1] T. Stoyanov, M. Magnusson, H. Andreasson, and A. Lilienthal, "Fast and accurate scan registration through minimization of the distance between compact 3D NDT representations," *IJRRRobotics Research*, vol. 31, pp. 1377–1393, 09 2012.
- [2] A. Segal, D. Hähnel, and S. Thrun, "Generalized-ICP," 06 2009.
- [3] J. Zhang and S. Singh, "Low-drift and real-time lidar odometry and mapping," *Autonomous Robots*, vol. 41, pp. 401–416, 02 2017.
- [4] J. Behley and C. Stachniss, "Efficient surfel-based SLAM using 3D laser range data in urban environments," in *Systems (RSS)*, 2018.
- [5] J. Zhu, "Image gradient-based joint direct visual odometry for stereo camera," in *IJCAI*, 2017, pp. 4558–4564.
- [6] M. Buczko and V. Willert, "Flow-decoupled normalized reprojection error for visual odometry," in *2016 IEEE 19th (ITSC)*, 2016, pp. 1161–1167.
- [7] C. Forster, M. Pizzoli, and D. Scaramuzza, "SVO: Fast semi-direct monocular visual odometry," in *2014 (ICRA)*, 2014, pp. 15–22.
- [8] J. Zhang and S. Singh, "Visual-lidar odometry and mapping: low-drift, robust, and fast," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 2174–2181.
- [9] H. Andreasson, D. Adolfsson, T. Stoyanov, M. Magnusson, and A. J. Lilienthal, "Incorporating ego-motion uncertainty estimates in range data registration," in *2017 IEEE/RSJ (IROS)*, 2017, pp. 1389–1395.
- [10] L. Narula, P. A. Iannucci, and T. E. Humphreys, "All-weather sub-50-cm radar-inertial positioning," [Online]. Available: <http://arxiv.org/abs/2009.04814>
- [11] A. Carballo, J. Lambert, A. Monrroy, D. Wong, P. Narksri, Y. Kit-sukawa, E. Takeuchi, S. Kato, and K. Takeda, "LIBRE: The multiple 3d lidar dataset," *arXiv preprint arXiv:2003.06129*, 2020, (accepted for presentation at IV2020).
- [12] D. Barnes, R. Weston, and I. Posner, "Masking by moving: Learning distraction-free radar odometry from pose information," in *CoRL*, ser. CoRL, L. P. Kaelbling, D. Kragic, and K. Sugiura, Eds., vol. 100. PMLR, 30 Oct–01 Nov 2020, pp. 303–316.
- [13] D. Barnes and I. Posner, "Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar," in *(ICRA)*, 2020, pp. 9484–9490.
- [14] Z. Hong, Y. Petillot, and S. Wang, "RadarSLAM: Radar based large-scale SLAM in all weathers," 2020.
- [15] S. H. Cen and P. Newman, "Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions," in *2018 IEEE (ICRA)*, 2018, pp. 6045–6052.
- [16] —, "Radar-only ego-motion estimation in difficult settings via graph matching," in *2019 (ICRA)*, 2019, pp. 298–304.
- [17] Y. S. Park, Y. S. Shin, and A. Kim, "Pharao: Direct radar odometry using phase correlation," in *2020 IEEE (ICRA)*, 2020, pp. 2617–2623.
- [18] D. Vivet, P. Checchin, and R. Chapuis, "Localization and mapping using only a rotating fmcw radar sensor," *Sensors*, vol. 13, no. 4, pp. 4527–4552, 2013.
- [19] M. Mielle, M. Magnusson, and A. J. Lilienthal, "A comparative analysis of radar and lidar sensing for localization and mapping," in *ecmr*, Sept. 2019.
- [20] R. Weston, S. Cen, P. Newman, and I. Posner, "Probably unknown: Deep inverse sensor modelling radar," in *2019 (ICRA)*. IEEE, 2019, pp. 5446–5452.
- [21] R. Aldera, D. D. Martini, M. Gadd, and P. Newman, "Fast radar motion estimation with a learnt focus of attention using weak supervision," in *2019 (ICRA)*, 2019, pp. 1190–1196.
- [22] J. W. Marck, A. Mohamoud, E. vd Houwen, and R. van Heijster, "Indoor radar slam a radar application for vision and gps denied environments," in *2013 European Radar Conference*. IEEE, 2013, pp. 471–474.
- [23] J. W. Marck, A. Mohamoud, E. vd Houwen, and R. van Heijster, "Indoor radar SLAM: A radar application for vision and GPS denied environments," in *2013 European Radar Conference*, 2013, pp. 471–474.
- [24] D. Vivet, P. Checchin, and R. Chapuis, "Localization and mapping using only a rotating FMCW radar sensor," *Sensors*, vol. 13, no. 4, p. 4527–4552, Apr 2013. [Online]. Available: <http://dx.doi.org/10.3390/s130404527>
- [25] F. Schuster, C. G. Keller, M. Rapp, M. Haeus, and C. Curio, "Landmark based radar SLAM using graph optimization," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, 2016, pp. 2559–2564.
- [26] M. Holder, S. Hellwig, and H. Winner, "Real-time pose graph SLAM based on radar," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1145–1151.
- [27] K. Burnett, A. P. Schoellig, and T. D. Barfoot, "Do we need to compensate for motion distortion and doppler effects in spinning radar navigation?" vol. 6, no. 2, pp. 771–778, conference Name: IEEE Robotics and Automation Letters.
- [28] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer, "Instantaneous ego-motion estimation using multiple Doppler radars," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 1592–1597.
- [29] R. Aldera, D. D. Martini, M. Gadd, and P. Newman, "What could go wrong? introspective radar odometry in challenging environments," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019, pp. 2835–2842.
- [30] Ș. Săftescu, M. Gadd, D. De Martini, D. Barnes, and P. Newman, "Kidnapped radar: Topological radar localisation using rotationally-invariant metric learning," [Online]. Available: <http://arxiv.org/abs/2001.09438>
- [31] M. Gadd, D. D. Martini, and P. Newman, "Look around you: Sequence-based radar place recognition with learned rotational invariance," 2020.
- [32] T. Y. Tang, D. De Martini, D. Barnes, and P. Newman, "Rsl-net: Localising in satellite images from a radar on the ground," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1087–1094, 2020.
- [33] T. Y. Tang, D. D. Martini, S. Wu, and P. Newman, "Self-supervised localisation between range sensors and overhead imagery," 2020.
- [34] P. J. Huber, *Robust Estimation of a Location Parameter*. New York, NY: Springer New York, 1992, pp. 492–518. [Online]. Available: https://doi.org/10.1007/978-1-4612-4380-9_35
- [35] S. Agarwal, K. Mierle, and Others, "Ceres solver." [Online]. Available: <http://ceres-solver.org>
- [36] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Paris, 2020. [Online]. Available: <https://arxiv.org/abs/1909.01300>
- [37] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the KITTI vision benchmark suite," in *Pattern Recognition (CVPR)*, 2012.
- [38] H. Zhan, C. S. Weerasekera, J. Bian, and I. Reid, "Visual odometry revisited: What should be learnt?" *arXiv preprint arXiv:1909.09803*, 2019.
- [39] R. Lovelace, J. Nowosad, and J. Muenchow, *Geocomputation with R*. CRC Press, 2019.
- [40] W. S. Churchill, "Experience based navigation : theory, practice and implementation," 2012.