

Dr-BA[🐑]: Separable Optimization for Direct Radar Bundle Adjustment & Localization

Daniil Lisus, Cedric Le Gentil, Timothy D. Barfoot
Robotics Institute, University of Toronto, Toronto, Canada
Email: {daniil.lisus; cedric.legentil; tim.barfoot}@robotics.utoronto.ca

Abstract—This paper introduces Dr-BA, a first-of-its-kind radar bundle adjustment (BA) framework that operates directly on 2D spinning radar intensity images. Unlike camera or lidar sensors, radar is largely unaffected by precipitation, making it a critical modality for autonomous systems that require all-weather robustness. Existing state estimation approaches using spinning radar typically extract sparse point clouds from range-azimuth-intensity measurements and apply point cloud alignment techniques to estimate vehicle motion, scene structure, or to localize within an existing map. In contrast, Dr-BA uses the full radar returns from multiple scans to jointly estimate dense maps and sensor poses. By formulating the problem as a separable optimization, we derive an efficient and general solution that decouples pose estimation from mapping. In addition to solving the BA problem, this formulation naturally extends to direct radar-only localization (DRL) within a previously built map. Dr-BA achieves state-of-the-art radar-based BA and cross-session localization performance, demonstrated on more than 200 km of on-road data across five distinct routes. Our implementation is publicly available at https://github.com/utiasASRL/dr_ba.

I. INTRODUCTION

For a robot to function autonomously, it must have an accurate representation of the environment in which it operates. Consequently, a critical challenge in modern robotic deployment is constructing a reliable map of the environment. This is typically achieved using simultaneous localization and mapping (SLAM) algorithms, in which sensor measurements are used to jointly estimate both the sensor’s trajectory and the surrounding map. SLAM algorithm performance is commonly evaluated by measuring the accuracy of the estimated sensor trajectory using metrics such as the absolute trajectory error (ATE). However, the primary output of a SLAM algorithm is the map itself, and good trajectory accuracy does not necessarily imply high-quality maps for downstream applications. For instance, a trajectory with zero-mean jitter may yield a low ATE while producing a locally inconsistent map that is difficult to use in practice. In this paper, we present Dr-BA, a radar-based bundle adjustment (BA) framework designed to produce maps that are both locally and globally consistent. We evaluate the proposed method through quantitative and qualitative BA metrics, and




Fig. 1: Dr-BA is a method for direct bundle adjustment of 2D radar data for simultaneous localization and mapping. Unlike pose-graph-based approaches, Dr-BA builds locally & globally consistent maps directly for downstream tasks such as localization. The figure shows a Dr-BA-constructed map from a clear day with a localized (red) and ground truth (black) trajectory from a snowy day.

further demonstrate its effectiveness on the downstream task of localization.

Dr-BA is a direct radar method. Direct methods, as opposed to feature-based ones that extract sparse features from measurements, use all of the information provided by a sensor. This is particularly valuable for map construction, since feature-based methods tend to produce relatively sparse maps that typically require downstream applications to use the same feature extraction as used

for map construction. Spinning radars lend themselves well to direct methods as they output dense range-azimuth-intensity polar images. Previous radar mapping approaches (e.g., [6]) extracted point clouds from these images and grouped them using alignment techniques into point maps of the environment. However, these point maps are generally very sparse, difficult to interpret, and require continued use of point cloud methods for localization. In contrast, Dr-BA produces dense intensity maps that can be sampled at any frequency or location, facilitating downstream applications that may prefer maps of different densities. Figure 1 presents an example map produced by Dr-BA on a clear day, with a trajectory localized to that map using data collected during a snowstorm, demonstrating both the practicality of Dr-BA and radar’s robustness to adverse weather conditions.

One drawback of direct methods is that they can quickly become intractable as the number of measurements grows. This is particularly problematic for BA in autonomous driving, where maps span kilometres and are built from thousands of poses. Leveraging the objective function from [34] for medial imaging BA, we formulate our problem in a way that allows us to separately estimate the sensor trajectory and the map. In contrast to [34], we cast the problem as a general separable optimization and employ variable projection (VarPro) to decouple the objectives, avoiding the need for Jacobian-level analysis. Consequently, Dr-BA complexity scales linearly with map size.

The key contributions of this paper are as follows:

- Dr-BA: the first framework for direct radar BA.
- The derivation of this framework as a separable optimization problem, separating the trajectory and map optimizations for efficiency.
- The demonstration of a downstream application with the implementation of a first direct radar localization (DRL) method using Dr-BA-generated maps.
- A publicly available ROS2-enabled implementation¹ of Dr-BA and DRL, validated on over 200 km of on-road data collected across diverse routes, including highly challenging environments.

II. RELATED WORK

This section provides a brief overview of both feature-based and direct methods for state estimation with 2D spinning (imaging) radars. We refer the reader to [1, 16, 38] for a wider overview of radar techniques in robotics.

The term *direct method* for state estimation originates from vision literature, where it is used to distinguish approaches that operate on the full set of sensor returns from feature-based methods that function on a subset of measurements. This distinction also applies to 2D spinning radar-based methods. Spinning radars typically return raw data in polar images with each row corresponding to an azimuth, each column to a range bin, and each pixel

value to the recorded intensity. The intensity represents the strength of radio-wave reflections from objects in the environment. These images can be transformed into Cartesian coordinates for easier interpretability.

A. Radar odometry methods

The majority of spinning radar odometry methods to-date have been feature-based. The authors of [10] present a radar feature extraction and association method tailored for odometry. ORORA [30] introduces a graph-based pruning algorithm to remove outlier associations between features from two scans. Other frameworks use normal distribution transform (NDT) to perform odometry [26]. Modern radar-based methods often extract point clouds from the raw data before performing some iterative closest point (ICP)-like scan-to-scan or scan-to-map registration. An extensive analysis of different point cloud extraction strategies for radar-based odometry is provided in [37]. Some of the best-performing point-based radar odometry frameworks are [3] that leverage point-to-distribution distances, and [8] with point-to-point registration. In addition to differing in the distance metrics used for registration, these methods rely on distinct state representations: discrete states in [3] and continuous-time states in [8]. The choice of state representation leads to two strategies for handling motion and Doppler distortion in the radar data: [3] performs a ‘one-off’ undistortion step based on previous state estimates, while [8] tightly couples the problem of data undistortion and ego-motion estimation in a single formulation.

While efficient, feature-based methods overlook a non-negligible amount of information from the original data. In contrast, direct methods leverage the full radar return in Cartesian or polar form. A common direct approach is to compute cross-correlation scores between consecutive radar scans using a discrete set of motion hypotheses and select the one with the highest score [5, 11, 36]. It is important to note that many direct methods do not address the issues of motion and Doppler distortion of the raw data. Using a radar unit with a particular frequency modulation pattern, it is possible to estimate the ego-velocity of the system with direct cross-correlation between consecutive collection azimuths [31]. This method provides a very efficient source of odometry when coupled with a gyroscope. DRO [27] introduces a continuous cross-correlation optimization between a scan and a local map to directly address motion and Doppler distortion as part of odometry.

B. Radar SLAM and localization methods

While odometry is a required component in an autonomous system’s stack, it is not sufficient for most applications. For many tasks, and especially in self-driving, a critical feature is localization within previously built maps. Localization along repeated routes does not require globally consistent maps, as demonstrated in [6, 13]. However,

¹Code available at https://github.com/utiasASRL/dr_ba.

the availability of global maps can enable a wider range of localization scenarios with more diverse trajectories. In these scenarios, odometry alone cannot be used to generate maps of the environment due to its inherent drift. SLAM and place recognition methods are required to generate globally consistent maps. Early work in SLAM used radar for landmark-based state estimation [12] by extracting distinctive points and performing data association. A similar feature-based approach, leveraging SIFT [33] from the vision-based research community, can be found in [9]. Pose-graph optimization (PGO) is the base of most existing radar SLAM frameworks [2, 19, 21]. These rely on loop-closure-detection approaches adapted from lidar applications [17, 18, 24] or specifically designed for radar place-recognition [22, 23]. The most relevant PGO approach to this paper is Dr-PoGO [29], which employs a direct radar-based odometry method (DRO [27]), RaPlace [22] for loop-closure detection, and a modified version of DRO for loop-closure registration. While PGO-based methods can provide fairly accurate trajectory estimates, the resulting maps, created from ‘stitching together’ radar data at trajectory estimates, are not guaranteed to be locally consistent. To-date, every SLAM method except Dr-PoGO has been feature-base.

Dr-BA is the first direct radar BA framework to produce maps that are both locally and globally consistent, optimizing direct registration across all collected scans. We also introduce DRL, the first framework to leverage these maps for localization.

C. Separable optimization

In this paper, we leverage the theory of separable optimization to decouple the BA problem. Separable optimization problems admit an efficient closed-form solution for a subset of variables when the remaining variables are fixed. One particularly effective approach for solving such problems is VarPro, which analytically eliminates the non-fixed variables and solves a reduced problem over the remaining ones [14, 15]. Recently, VarPro has been applied to several robotics algorithms, yielding improvements in both computational efficiency and convergence [20, 35].

Although the formulation in [34] is not explicitly posed as a separable problem, the authors showed that direct BA can be solved independently for the states and the map. In this paper, we extend and generalize their result by explicitly casting the problem as a separable optimization. This allows us to decouple the variables directly at the objective-function level, avoiding the need to inspect the Jacobian structure or rely on the Schur complement as done in [34]. Additionally, our formulation supports weighted error terms and simplifies the incorporation of additional error types. We apply direct BA to spinning radar data for the first time and demonstrate a simple, yet effective, localization framework that validates the resulting maps on a real downstream task.

III. METHODOLOGY

The key idea of Dr-BA is to formulate the BA error such that the map variables can be eliminated from the objective function. This is done by considering the map variables to be intensities at predetermined sample locations, instead of the locations themselves as is common in feature-based SLAM. This formulation turns the problem into a separable optimization in which the poses and map intensities can be estimated independently through the use of VarPro, making trajectory estimation linear in the number of map states. Furthermore, the entire map estimation problem can be simplified to a per-map-location estimation task, which is also linear in the number of map states. A final practical benefit is that we are able to easily generate maps in a principled manner given poses from any source, be they the solution of the BA problem or acquired by other means.

A. Direct bundle adjustment

We wish to simultaneously estimate the intensity values $\mathbf{i} \in \mathbb{R}^P$ of a set of d -dimensional map points, $\mathcal{M} = \{\mathbf{m}_v \in \mathbb{R}^d\}_{v=1}^P$, and the trajectory of a sensor within that map. The intensity value at each map position, $\mathbf{m}_v$, is given by $i(\mathbf{m}_v) \equiv i_v$. Note that we are not optimizing for the positions $\mathbf{m}_v$ themselves, only for the intensities at those positions. We intentionally keep the map sampling general, since it is application- and compute-resource-dependent. In this paper, we use a uniform sampling approach, but one could employ more complex information-density-based sampling depending on the task at hand. The trajectory is represented as a set of d -dimensional discrete states, $\mathbf{T} = \{\mathbf{T}_n\}_{n=1}^N$, at which intensity measurements, called *scans*, are rooted. We can get the measured intensity value from a given scan at $\mathbf{m}_v$ by using a measurement sampling function $\beta(\cdot)$,

$$\tilde{i}_{v,n} = \beta(\mathbf{m}_v, \mathbf{T}_n). \quad (1)$$

Depending on the scan representation, $\beta(\cdot)$ can directly query or interpolate a measurement at $\mathbf{m}_v$ in the scan.

We form a general loss on intensity differences across the map as a weighted least-squares problem

$$\mathbf{i}^*, \mathbf{T}^* = \arg \min_{\mathbf{i}, \mathbf{T}} \sum_{v=1}^P \sum_{n \in \mathcal{V}_v} w_{v,n} (i_v - \tilde{i}_{v,n})^2, \quad (2)$$

where $\mathcal{V}_v$ is the set of indices of discrete states that observe a given $\mathbf{m}_v$, and $w_{v,n} > 0$ is the weight of each residual.

We direct the reader’s attention to the fact that the intensities, i_v , enter linearly into the error terms of the cost. This fact allows us to apply the VarPro method to this problem, which can eliminate such linear variables to reduce problem complexity [15]. Treating $\mathbf{T}$ as fixed and optimizing only over $\mathbf{i}$ in (2), we note that the optimal value of each intensity, i_v , is given by

$$i_v(\mathbf{T}) = \frac{1}{w_v} \sum_{n \in \mathcal{V}_v} w_{v,n} \tilde{i}_{v,n} \equiv \bar{i}_v, \quad (3)$$




Fig. 2: An overview of the implementation of Dr-BA, including initialization steps and a downstream localization task.

where $\sum_{n \in \mathcal{V}_v} w_{v,n} \equiv w_v$ for brevity. Thus, if the scan poses are known, the intensity values of the map will simply be the weighted average of the interpolated scan intensities that overlap each map location. Conveniently, (3) enables efficient, easily parallelized per-location map intensity generation, avoiding the need to form a large linear system.

Having solved for $\mathbf{i}$ as a function of $\mathbf{T}$, we plug (3) back into (2) to now optimize only over $\mathbf{T}$ as

$$\mathbf{T}^* = \arg \min_{\mathbf{T}} \sum_{v=1}^P \sum_{n \in \mathcal{V}_v} w_{v,n} (\bar{i}_v - \tilde{i}_{v,n})^2, \quad (4)$$

where the objective function is still dependent on sampling locations, $\mathbf{m}$, but where our substitution has projected out the explicit dependence on the intensities, $\mathbf{i}$. Note that (4) is finding an optimal $\mathbf{T}$ such that the weighted variance of intensity measurements is minimized at all map locations. Since (4) is a fully relative problem, we fix the first pose.

We iteratively solve (4) using Gauss-Newton, alternating between linearization of the objective and solving the resulting linear system as

$$\mathbf{H} \Delta \mathbf{T} = \mathbf{b}, \quad (5)$$

where $\mathbf{H} = \mathbf{J}^\top \mathbf{J} \in \mathbb{R}^{dN \times dN}$, $\mathbf{b} = -\mathbf{J}^\top \mathbf{e} \in \mathbb{R}^{dN}$, $\mathbf{e}$ is the vector of residuals, $\mathbf{J}$ is the Jacobian of the residuals with respect to the state perturbations, and $\Delta \mathbf{T}$ is the pose update for a given iteration. Both $\mathbf{H}$ and $\mathbf{b}$ can be efficiently computed as

$$\mathbf{H} = \sum_{v=1}^P \mathbf{J}_v^\top \mathbf{J}_v, \quad \mathbf{b} = \sum_{v=1}^P \mathbf{J}_v^\top \mathbf{e}_v, \quad (6)$$

where $\mathbf{J}_v \in \mathbb{R}^{|\mathcal{V}_v| \times dN}$ is the Jacobian of error terms involving only $\mathbf{m}_v$, grouped as $\mathbf{e}_v \in \mathbb{R}^{|\mathcal{V}_v|}$, with respect to all states. In the worst case, solving $\mathbf{H}$ has cubic complexity in the number of trajectory poses. However, in practice, $\mathbf{J}_v$ is sparse in many bundle adjustment problems since only $d|\mathcal{V}_v|$ columns are non-zero. In typical mapping scenarios, co-observability between states is much smaller than the total number of states, so $|\mathcal{V}_v| \ll N$. Consequently, $\mathbf{H}$ also becomes relatively sparse with entries only on the block diagonal and in cross-correlation blocks between states observing the same map locations. Notably, the size of the linear problem (5) does not depend on the number of map points sampled. Increasing the sampling rate of the map will affect the number of loops we execute, but not the

amount of memory required to solve for the trajectory. After solving for $\mathbf{T}$, we can use (3) to build the map.

Note that we have yet to specify a map, scan, or state representation, nor have we defined a specific scan sampling function. We apply this efficient BA approach to the task of direct 2D spinning radar mapping in this paper, referring to our method as Dr-BA. However, this formulation can be used with any sensor that produces intensity measurements in any number of dimensions.

The derivation extends naturally to objectives of the form (2) with additional pose- or map-only terms, as these do not affect separability. For instance, state priors for full SLAM can be incorporated directly into (4) without affecting (3).

B. Direct localization

The simplest way to do localization, given our previous formulation, is to re-use (2) to optimize exclusively for a single state given fixed map intensity values $\mathbf{i}$. The objective function to localize pose $\mathbf{T}_n$ in the map is thus

$$\mathbf{T}_n^* = \arg \min_{\mathbf{T}_n} \sum_{v=1}^Q w_{v,n} (i_v - \tilde{i}_{v,n})^2, \quad (7)$$

where we loop over the Q intensities observable by the scan at $\mathbf{T}_n$. Just as with the BA problem, (7) is general across any map, scan, and state representation, as well as any scan sampling scheme. Gauss-Newton can again be used to iteratively solve for $\mathbf{T}_n$. When used with radar data, we term this method DRL.

IV. IMPLEMENTATION

The general theory from Section III is applied explicitly in the formulation below; an implementation overview is shown in Figure 2. We perform BA directly on spinning radar images to construct dense, globally consistent maps for downstream autonomous driving tasks such as localization. Specifically, we fuse many 2D radar images, such as those shown at the top of Figure 3, collected along an on-road *sequence* into a single map (Figure 1).

A. Map Representation

Since the radar collects 2D data, the map is 2D with locations $\mathbf{m} \in \mathbb{R}^2$ and intensities $i \in [0, 1]$. For simplicity, we sample it regularly with resolution r_v .




Fig. 3: Different processing steps for leveraging raw radar images for Dr-BA. Data from an ‘informative’ suburban location is shown on the left and data from an ‘uninformative’ skyway location is shown on the right.

B. State Representation

We consider the radar data to be collected at discrete poses, $\mathbf{T} = \{\mathbf{T}_n \in SE(2)\}_{n=1}^N$, each defining the transform from the radar to the world frame at that timestamp. We optimize the poses, $\mathbf{T}$, directly on-manifold and select the moving-frame-invariant perturbation scheme to compute Jacobians and iteratively optimize our state [4].

In order to not over-index on stationary segments during which many repeated scans are collected, a keyframing approach, common in visual BA [25], is taken for selecting poses. A new keyframe, and consequently scan, is only added to the problem if sufficient motion (≥ 5 m or $\geq 30^\circ$) has been experienced since the last keyframe.

As Dr-BA performs no explicit loop-closure or feature-based data association, a good initial trajectory estimate

is required for convergence. Here, we initialize using Dr-PoGO [29].

C. Scan Representation

We assume the scans to be Cartesian images with a resolution r_r per (u, v) pixel. Following the standard pixel coordinate convention with the origin in the top-left corner, the positive u axis toward the right, and the positive v axis toward the bottom, a map point $\mathbf{m}_v$ in the world frame is converted to pixel coordinates $\mathbf{p}_v^n$ of a radar scan captured from $\mathbf{T}_n$ as

$$\mathbf{p}_v^n = \begin{bmatrix} 0 & \frac{1}{r_r} & 0 \\ -\frac{1}{r_r} & 0 & 0 \end{bmatrix} \mathbf{T}_n^{-1} \begin{bmatrix} \mathbf{m}_v \\ 1 \end{bmatrix} + \frac{1}{2} \begin{bmatrix} w-1 \\ h-1 \end{bmatrix}, \quad (8)$$

with w and h the width and height of the image. As the pixel position from (8) is unlikely to be a whole number, we use bilinear interpolation to obtain $\tilde{i}_{v,n}(\mathbf{p}_v^n)$. The uncertainty of $\tilde{i}$, used to compute the corresponding $w_{v,n}$ in (2) and (7), is formed from a constant per-pixel uncertainty ($\mathcal{N}(\mathbf{0}, 0.1^2)$) and a range-dependent uncertainty ($\mathcal{N}(\mathbf{0}, (0.005 \|\mathbf{m}_n\|)^2)$, where $\mathbf{m}_n$ is the position of $\mathbf{m}_v$ in radar frame $\mathbf{T}_n$).

The proposed BA formulation optimizes the poses of the sensor for a discrete set of keyframes. However, spinning radars do not take instantaneous snapshots of the environment; they sweep their surroundings with one or more dishes over a period typically between 100 and 500 ms. Accordingly, a radar scan collected during a 360° sweep is subject to motion distortion. Additionally, when the sensor is moving relatively to its environment, measurements will also be subject to the Doppler effect. To obtain undistorted radar scans, both in terms of motion and Doppler effect, we again leverage DRO [27]. DRO outputs undistorted scans and smooth local maps as Cartesian images. Local maps are used as inputs to Dr-BA due to their inherent noise rejection and the larger convergence basin induced by slight smoothing.

However, even with the smoothing offered by local maps, radar measurements collected in very sparse environments (see Figure 3) can contain almost no gradient information for pixel-level optimization. To address this, we adopt an adaptive Gaussian blurring approach. For each scan, we compute the amount of ‘valuable information’ within the scan by a simple check of the percent of pixels that contain intensities above a threshold (0.5). If the percent is lower than some bound (0.3%), we apply a progressively stronger Gaussian blur to the image until the bound is reached. This dynamically changes the amount of blurring applied to scans throughout a sequence, keeping informative scans almost untouched and heavily blurring scans that need a wider convergence basin around the few available features. Dynamic blurring is the key to getting good alignment in sparse environments such as the one visualized on the right of Figure 4.

Another practical problem with direct radar BA is that objects in the environment can prevent the radar

signal from measuring to its full range, despite the radar image always returning the same number of azimuth-range-intensity measurements. This means that some zero-intensity measurements correspond to a lack of material in the environment, while others correspond to regions that the radar signal never reached due to occlusions. In the latter case, incorporating these ‘fake’ zero measurements yields an incorrect down-weighting of the corresponding map locations. A related inverse problem is radar saturation, often caused by ‘ringing’, where the signal repeatedly reflects between parallel surfaces. This ringing, visually seen in the right radar scan in Figure 3, returns high intensities all along a given azimuth regardless of what is physically there. These returns incorrectly up-weight those regions. To handle both of these cases, we construct a ‘cumulative’ radar scan based on the input to Dr-BA. For each azimuth, we sum the intensities outward along each ray and store the cumulative sum in the corresponding pixel of the cumulative image. We then ignore all zero measurements that have a cumulative score above a threshold (0.2) and ignore all measurements with a score above another higher threshold (0.9).

The dynamic smoothing and cumulative image processing are handled as part of the Filtering block in Figure 2. Visual examples of the different types of scans that take part in the Dr-BA pipeline are shown in Figure 3. Finally, the impact on performance from these different parts of the pre-processing is studied in Section V-E.

V. EXPERIMENTS

A. Dataset

To demonstrate the performance of Dr-BA and DRL, we use a subset of the Boreas Road Trip (Boreas-RT) dataset [32], which contains data from a Navtech RAS6 radar, a Silicon Sensing DMU41 6-DoF IMU, and an Applanix RTK-GNSS/INS solution for ground-truthing. The radar produces scans with a resolution of $r_r = 0.0438$ m and a range of up to 300 m, although only the first 100 m have useful returns for BA in practice. The subset consists of 20 sequences in 5 different environments (4 repeats per route) for a total of 206.4 km: **Suburbs** (7.9 km), **Industrial** (5.4 km), **Skyway** (11.1 km), **Forest** (16.4 km), and **Farm** (10.8 km). The **Suburbs** route goes along suburban streets before returning to the starting location along the same set of streets in the opposite direction. The **Industrial** route loops through an industrial area with commercial buildings before coming back to the starting location, with some sequences collected during a snowstorm. **Skyway** sequences were collected driving up and down a busy high-speed skyway with very limited structural features but a lot of measurement overlap, making both the problem of odometry and loop-closure detection/registration extremely challenging. The **Forest** and **Farm** routes form large loops in very challenging environments with very few buildings along the vehicle’s path and not a lot of measurement overlap. The **Suburbs**




Fig. 4: Qualitative comparison of trajectory estimates from TBV-SLAM, our pose-graph initialization (Dr-PoGO), and the proposed bundle adjustment method Dr-BA. Each image overlays two scans from the same location collected in opposite driving directions (blue/orange). While Dr-PoGO achieves low ATE, it does not guarantee local consistency (see longitudinal misalignment). Dr-BA uniquely enforces local consistency, producing well-overlapping scans critical for downstream tasks such as localization.

and **Industrial** sequences represent ‘easy’, structured on-road conditions that are typically tested by autonomous vehicle SLAM methods. The **Skyway**, **Forest**, and **Farm** sequences represent environments in which current state-of-the-art (SOTA) algorithms are expected to perform poorly, and in some cases fail.

B. Metrics

The ATE evaluates the overall accuracy of Dr-BA and the baseline methods for trajectory estimation. After aligning the estimated and ground-truth trajectories via SVD-based $SE(2)$ registration, the ATE is computed as $\sqrt{\frac{1}{N} \sum_{n=1}^N \|\hat{\mathbf{r}}_n - \mathbf{r}_n^{\text{GT}}\|^2}$, where $\hat{\mathbf{r}}_n$ is the translation component of the pose estimate $\hat{\mathbf{T}}_n$ and $\mathbf{r}_n^{\text{GT}}$ is the ground-truth translation. Since all sequences finish close to the starting location, we use the end-pose-error (EPE) to

TABLE I: Average ATE and EPE (all [m]) per sequence type for TBV-SLAM, DRO, Dr-PoGO, and Dr-BA trajectories.

Method	Suburbs	Industrial	Skyway	Forest	Farm
DRO [27]	9.53 / 38.9	6.32 / 17.2	18.6 / 60.7	25.8 / 80.5	38.8 / 170.4
TBV-SLAM [2]	6.11 / 0.55	4.72 / 0.82 ¹	- / - ⁴	- / - ⁴	- / - ⁴
Dr-PoGO [29]	0.75 / 0.42	1.58 / 0.77	3.16 / 0.23	4.31 / 0.75	4.19 / 0.56
Dr-BA (Ours)	0.54 / 0.25	1.31 / 0.31	2.46 / 0.17	4.36 / 0.78	4.24 / 0.42

Results reported as XX / YY, with XX the ATE and YY the EPE.

¹ The superscript indicates the number of failed sequences (ATE above 1% of the sequence distance).

TABLE II: Average translation ([m]) & rotation ([°]) self-consistency RMSE per sequence type for Dr-PoGO and Dr-BA.

Sequence	Dr-PoGO [29]	Dr-BA (ours)
Suburbs	0.34 / 0.11	0.14 / 0.09
Industrial	0.36 / 0.18	0.10 / 0.10
Skyway	3.49 / 0.10	1.29 / 0.24
Forest	0.57 / 0.12	0.52 / 0.12
Farm	0.28 / 0.08	0.16 / 0.11

XX / YY, with XX the translation and YY the rotation RMSE.

assess the effectiveness of loop-closure constraints in each method. The EPE is defined as the norm of the translation component of $(\mathbf{T}_N^{\text{GT}})^{-1}\mathbf{T}_1^{\text{GT}}(\hat{\mathbf{T}}_n)^{-1}\hat{\mathbf{T}}_N$. We also introduce a ‘self-consistency’ metric to evaluate the quality of the produced maps. For each mapping pose, we identify the nearest pose that is at least 300m away by distance travelled, but within 25m Euclidean distance, then compute the relative error between the associated poses in the estimated trajectory and the ground truth. This metric can be interpreted as a localization error within a trajectory when revisiting previously explored areas. The more accurately these poses align, the more locally consistent the resulting map is. To assess the accuracy of the proposed localization pipeline, which further validates the mapping accuracy, we analyze the relative pose error between the trajectory used for mapping and the localization run following the Boreas dataset localization benchmark [7].

C. Direct radar bundle adjustment

Dr-BA is compared to three baselines. First, we compare to a SOTA odometry method, DRO [27], with gyroscope measurements and explicit Doppler compensation, with no loop closures. This approach, unsurprisingly, performs the worst. Second, we compare to an off-the-shelf SOTA radar SLAM method TBV-SLAM [2], which employs a robust loop closure candidate verification approach. Finally, we evaluate performance using Dr-PoGO [29]. Table I shows the average ATE and EPE metrics for all considered methods. Figure 1 shows a full Dr-BA-created map of an **Industrial** sequence. Table II shows the translation and rotation self-consistency RMSE from Dr-PoGO and Dr-BA for all sequences in all routes. We omit DRO and TBV-SLAM results for brevity and because they were significantly worse (> 10 m).

We note that Dr-PoGO already achieves notably better performance than TBV-SLAM, likely due to the high-quality odometry provided by DRO. Dr-PoGO improves ATE performance by three to ten times on the structured **Suburbs** and **Industrial** sequences, and prevents failure on the challenging sequences. In environments with extensive trajectory overlap (**Suburbs**, **Skyway**) or well-defined geometric features (**Industrial**), Dr-BA further refines the Dr-PoGO trajectory, achieving the best average ATE and EPE results. In the unstructured environments with very little trajectory overlap (**Forest**, **Farm**), Dr-BA performs on par with Dr-PoGO. However, the most significant improvements can be found in the local consistency of the produced maps. Figure 4 highlights regions from a **Suburbs** and **Skyway** sequence where Dr-PoGO yields a good trajectory, but lacks local map consistency. In contrast, Dr-BA generates a locally consistent map. The self-consistency metric further backs up this finding. Dr-BA significantly improves translational self-consistency on all routes with notable pose overlap (meaning all except **Forest**). Rotational self-consistency stays roughly similar, except for the **Skyway**, where the adaptive blurring results in a loss of rotational self-consistency in exchange for a significant boost in translational performance.

In practice, we find that $r_v = 1.0$ m is more than sufficient for high quality BA. This choice also helps maintain a multi-threaded runtime of half an hour to a few hours per map on a standard laptop, depending on the trajectory length and parameters, with maps being represented by up to four million points. Further runtime improvements, such as GPU-acceleration, are left for future work.

D. Direct radar localization

While trajectory and visual evaluations are useful for comparing BA algorithms, their ultimate goal is to produce high-quality maps for downstream tasks. To assess the practical impact of Dr-BA, we evaluate its maps within our novel direct radar localization framework, DRL. We note that these maps are also compatible with cross-correlation-based localization or point-cloud methods by representing non-zero map regions as points.

For a full comparison, we construct maps using (3) with poses from Dr-BA (DRL-Dr-BA), ground truth (DRL-GT), and Dr-PoGO (DRL-Dr-PoGO). DRL localizes a single scan to the map, using DRO to propagate the estimate to the next frame, without explicit priors or costs

TABLE III: RMSE localization errors for RTR and DRL. DRL is run on maps created using Dr-BA, ground truth (GT), and Dr-PoGO poses. **Forest** results are greyed out as they are at the limit of groundtruth noise.

Seq. type & Method	Long. [m]	Lat. [m]	Yaw [°]
Suburbs			
RTR [6]	0.092	0.052	0.107
DRL-Dr-BA (ours)	0.076	0.049	0.061
DRL-GT	0.071	0.061	0.074
DRL-Dr-PoGO	0.146	0.105	0.078
Industrial			
RTR [6]	0.080	0.054	0.116
DRL-Dr-BA (ours)	0.081	0.062	0.055
DRL-GT	0.076	0.082	0.073
DRL-Dr-PoGO	0.109	0.076	0.085
Skyway			
RTR [6]	-	-	₃
DRL-Dr-BA (ours)	0.301	0.200	0.145
DRL-GT	0.220	0.126	0.090
DRL-Dr-PoGO	1.517	0.419	0.157
Forest			
RTR [6]	-	-	₃
DRL-Dr-BA (ours)	0.429	0.371	0.122
DRL-GT	0.415	0.369	0.138
DRL-Dr-PoGO	0.425	0.371	0.141
Farm			
RTR [6]	-	-	₃
DRL-Dr-BA (ours)	0.404	0.595	0.285
DRL-GT	0.155	0.159	0.107 ¹
DRL-Dr-PoGO	0.191	0.251	0.148 ²

¹ The superscript indicates the number of failed sequences.

beyond the direct intensity objective in (7). We again find that $r_v = 1.0$ m is sufficient for high-quality localization.

We compare DRL with an open-source point cloud-based radar localization framework, radar teach & repeat (RTR) [6]. RTR creates locally consistent, but globally inconsistent, point maps of the environment, and then does topometric localization relative to them on repeated traversals of the same environment.

Table III presents the longitudinal, lateral, and yaw localization error RMSE for all sequence types and methods. On the structured **Suburbs** and **Industrial** routes, DRL-Dr-BA and DRL-GT achieve SOTA performance, keeping localization errors within 10 cm and 0.1°. While translational performance is comparable to RTR, DRL-Dr-BA substantially outperforms RTR, and slightly outperforms DRL-GT, in rotation. This is likely due to higher local map quality, as even small rotational inconsistencies at the radar frame cause large intensity offsets at long ranges. The value of Dr-BA is further shown by the fact that DRL-Dr-BA outperforms DRL-Dr-PoGO on all metrics on the **Suburbs**, **Industrial**, and **Skyway** sequences.

Performance for all methods drops on the more challenging sequences, with RTR failing to localize until the end of all **Skyway**, **Forest**, and **Farm** sequences. In contrast, DRL-based methods only experience failures on **Farm** sequences, with DRL-Dr-BA not having any unrecoverable failures

TABLE IV: Ablation study on the impact of different parameters on BA and consequently localization performance. Average ATE/EPE/localization translation RMSE per sequence type are shown (all [m]). The ‘Dr-BA’ method uses cumulative images (cumul. img.), adaptive blurring (AdBlur), and local maps (LocMaps.)

Method	Suburbs	Skyway
Dr-BA	0.51/0.15/0.09	1.33/0.10/0.36
AdBlur: off	0.53/0.15/ 0.09	1.60/0.10/1.19
No cumul. img.	0.54/0.20/0.10	1.33/0.09/0.39
No LocMaps	0.54/ 0.13 /0.10	1.73/ 0.08 /1.37 ²

¹ The superscript indicates the number of failed sequences.




Fig. 5: Average ATE and convergence for the **Industrial** route with increasingly degraded Dr-PoGO initialization.

across the entire dataset, something not achieved even by DRL-GT. Despite preventing failures, the Dr-BA map yields worse average localization performance in the successful **Farm** sequences than DRL-GT or DRL-Dr-PoGO. Possible reasons for this are outlined in Section V-F. All DRL methods produce similar results on the **Forest** route, as performance is limited by groundtruth noise due to the large distance to the nearest RTK base station and tree canopy [32]. Nearly identical errors are reported by baselines in [32] and in the SOTA lidar method in [28]. These results are included for completeness but not compared between implementations.

DRL takes on average 200 ms per frame running on a single thread of a Lenovo P1 Laptop with Intel(R) Core(TM) i7-12800H CPU @ 4.8 GHz and 32 GB of RAM, making it ‘real-time’ for the 4 Hz radar used in these tests.

E. Ablation studies

1) *Processing impact*: Table IV presents an ablation study on the impact of different processing parameters on BA and ultimately localization performance. Ablations are done on only a **Suburbs** and **Skyway** sequence as they contain a large number of overlapping regions in their traversals. DRL-Dr-BA localization results on the remaining three sequences are also included, using each map from the ablation study, for a complete evaluation. All options generally perform well on the **Suburbs** route, though combining all options yields a slight improvement




Fig. 6: Total solve time, Hessian memory use, and ATE comparison between separable and combined Dr-BA.

in ATE. This is due to the **Suburbs** having consistently well-defined geometric structure in all scans. However, the use of adaptive blurring and local maps is required to get good performance on the **Skyway** route, where the radar has very little to observe, as shown in Figure 7 (a).

2) *Initialization sensitivity*: We add increasingly aggressive random (uniform) initialization error to Dr-BA’s PGO initialization and plot ATE and iterations to convergence in Figure 5. Results show average performance from all **Industrial** sequences (chosen for smaller size and faster testing). Perturbations up to 3.0m and 1.50° are recovered without significant ATE degradation, although convergence naturally takes longer. This demonstrates that while general global consistency is required to initialize Dr-BA, the local consistency can be quite poor.

3) *VarPro impact*: To demonstrate the benefits of VarPro over the naive solution, we construct a test problem with 11 poses (first fixed) at a stationary location and progressively increase voxel resolution to scale the problem size. We evaluate ATE, solve time, and Hessian non-zeros (as a memory proxy). Figure 6 compares our separable Dr-BA implementation against a combined implementation that jointly solves (2) for poses and voxel intensities. Both yield near-identical ATE, but the combined formulation requires substantially more compute and memory, becoming prohibitively slow for full sequences.

F. Limitations

Despite SOTA performance on structured sequences and a decreased failure rate in challenging unstructured environments, there is a notable decrease in performance




Fig. 7: Challenging segments of the **Skyway**, **Farm**, and **Forest** routes, characterized by sparse features, diffuse natural structure, or poor scan co-visibility. The length-scale is shared by all plots.

of Dr-BA and DRL on the challenging routes. Particularly challenging segments of these routes are shown in Figure 7, where all three difficult sequences contain extensive regions with very few features visible from a single radar scan. The **Forest** sequences also contain tight vegetation corridors with poor co-visibility from scans collected short distances apart. These are fundamental sensor limitations in a given environment, and additional interoceptive data must be included in the estimation to mitigate them. This is likely why DRL-Dr-PoGO performs better in successful localization runs than DRL on the **Farm** sequences, as Dr-PoGO poses, and consequently the map produced using them, are constrained by relative motion information provided by Doppler velocity measurements and a gyroscope. Additionally, this work ignores the dependency of radar intensity returns on the viewpoint from which the sensor is observing a feature. Explicitly incorporating the viewpoint significantly increases complexity, but would likely yield a more accurate BA result.

VI. CONCLUSION

This paper introduces Dr-BA, a first-of-its-kind direct radar bundle adjustment framework. Dr-BA is based on a general separable formulation that enables independent optimization of the trajectory and the map, resulting in an efficient algorithm with linear scaling in the number of estimated map states. Using 2D spinning radar data, we demonstrate that Dr-BA can tractably construct dense maps spanning tens of kilometres while maintaining both local and global consistency. Building on this formulation, we also propose a novel direct radar localization (DRL) method. Experimental results on over 200km of real-world data show SOTA performance in both BA and localization using Dr-BA-generated maps. Tests are run on five diverse trajectories, ranging from structured suburban environments to highly unstructured farm roads.

REFERENCES

- [1] Nader J. Abu-Alrub and Nathir A. Rawashdeh. Radar Odometry for Autonomous Ground Vehicles: A Survey of Methods and Datasets. *IEEE Trans. on Intelligent Vehicles (T-IV)*, 9(3):4275–4291, 2024. doi: 10.1109/TIV.2023.3340513.
- [2] Daniel Adolfsson, Mattias Karlsson, Vladimír Kubelka, Martin Magnusson, and Henrik Andreasson. TBV Radar SLAM – Trust but Verify Loop Candidates. *IEEE Robotics and Automation Letters (RA-L)*, 8(6):3613–3620, 2023. doi: 10.1109/LRA.2023.3268040.
- [3] Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. Lidar-Level Localization With Radar? The CFEAR Approach to Accurate, Fast, and Robust Large-Scale Radar Odometry in Diverse Environments. *IEEE Trans. on Robotics (TRO)*, 39(2):1476–1495, 2023. doi: 10.1109/TRO.2022.3221302.
- [4] Timothy D Barfoot. *State Estimation for Robotics*. Cambridge University Press, 2024.
- [5] Dan Barnes, Rob Weston, and Ingmar Posner. Masking by Moving: Learning Distraction-Free Radar Odometry from Pose Information. In *Conf. Robot Learn.*, 2019. URL <https://arxiv.org/pdf/1909.03752>.
- [6] Keenan Burnett, Yuchen Wu, David J Yoon, Angela P Schoellig, and Timothy D Barfoot. Are We Ready for Radar to Replace Lidar in All-weather Mapping and Localization? *IEEE Robotics and Automation Letters (RA-L)*, 7(4):10328–10335, 2022.
- [7] Keenan Burnett, David J Yoon, Yuchen Wu, Andrew Z Li, Haowei Zhang, Shichen Lu, Jingxing Qian, Wei-Kang Tseng, Andrew Lambert, Keith YK Leung, Angela P Schoellig, and Timothy D Barfoot. Boreas: A Multi-season Autonomous Driving Dataset. *Int. J. Robot. Res.*, 42(1-2):33–42, 2023. doi: 10.1177/02783649231160195.
- [8] Keenan Burnett, Angela P. Schoellig, and Timothy D. Barfoot. Continuous-Time Radar-Inertial and Lidar-Inertial Odometry Using a Gaussian Process Motion Prior. *IEEE Trans. on Robotics (TRO)*, 41: 1059–1076, January 2025. ISSN 1552-3098. doi: 10.1109/TRO.2024.3521856.
- [9] Jonas Callmer, David Törnqvist, Fredrik Gustafsson, Henrik Svensson, and Pelle Carlbom. Radar SLAM Using Visual Features. *EURASIP Journal on Advances in Signal Processing*, 2011, 12 2011. doi: 10.1186/1687-6180-2011-71.
- [10] Sarah Cen and Paul Newman. Precise Ego-Motion Estimation with Millimeter-Wave Radar Under Diverse and Challenging Conditions. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 1–8, 05 2018. doi: 10.1109/ICRA.2018.8460687.
- [11] Paul Checchin, Franck Gérossier, Christophe Blanc, Roland Chapuis, and Laurent Trassoudaine. Radar Scan Matching SLAM Using the Fourier-Mellin Transform. In *Field and Service Robotics: Results of the 7th International Conference*, pages 151–161. Springer, 2010.
- [12] M.W.M.G. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, and M. Csorba. A Solution to the Simultaneous Localization and Map Building (SLAM) Problem. *IEEE Trans. on Robotics and Automation*, 17(3):229–241, 2001. doi: 10.1109/70.938381.
- [13] Paul Furgale and Timothy D Barfoot. Visual Teach and Repeat for Long-range Rover Autonomy. *Journal of Field Robotics (JFR)*, 27(5):534–560, 2010.
- [14] Gene Golub and Victor Pereyra. Separable Nonlinear Least Squares: the Variable Projection Method and its Applications. *Inverse Problems*, 19:R1–R26(1), 01 2003.
- [15] G.H. Golub and Victor Pereyra. The Differentiation of Pseudo-Inverses and Nonlinear Least Squares Problems Whose Variables Separate. *SIAM Journal on Numerical Analysis*, 10:413–432, 04 1973. doi: 10.1137/0710036.
- [16] Kyle Harlow, Hyesu Jang, Timothy D. Barfoot, Ayoun Kim, and Christoffer Heckman. A New Wave in Robotics: Survey on Recent MmWave Radar Applications in Robotics. *IEEE Trans. on Robotics (TRO)*, 40:4544–4560, 2024. doi: 10.1109/TRO.2024.3463504.
- [17] Li He, Xiaolong Wang, and Hong Zhang. M2DP: A Novel 3D Point Cloud Descriptor and its Application in Loop Closure Detection. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 231–237, 2016. doi: 10.1109/IROS.2016.7759060.
- [18] Marian Himstedt, Jan Frost, Sven Hellbach, Hans-Joachim Böhme, and Erik Maehle. Large scale place recognition in 2D LIDAR scans using Geometrical Landmark Relations. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 5030–5035, 2014. doi: 10.1109/IROS.2014.6943277.
- [19] Martin Holder, Sven Hellwig, and Hermann Winner. Real-Time Pose Graph SLAM based on Radar. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1145–1151, 2019. doi: 10.1109/IVS.2019.8813841.
- [20] Connor Holmes and Timothy D. Barfoot. An Efficient Global Optimality Certificate for Landmark-Based SLAM. *IEEE Robotics and Automation Letters*, 8(3): 1539–1546, 2023. doi: 10.1109/LRA.2023.3238173.
- [21] Ziyang Hong, Yvan Petillot, Andrew Wallace, and Sen Wang. RadarSLAM: A Robust Simultaneous Localization and Mapping System for All Weather Conditions. *Intl. Journal of Robotics Research (IJRR)*, 41 (5):519–542, 2022. doi: 10.1177/02783649221080483.
- [22] Hyesu Jang, Minwoo Jung, and Ayoun Kim. RaPlace: Place Recognition for Imaging Radar using Radon Transform and Mutable Threshold. In *Proc. of*

- the *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 11194–11201, 2023. doi: 10.1109/IROS55552.2023.10341883.
- [23] Hyesu Jang, Wooseong Yang, Ayoung Kim, Dongje Lee, and Hanguen Kim. XPRESS: X-Band Radar Place Recognition via Elliptical Scan Shaping. *IEEE Robotics and Automation Letters*, 10(12): 13121–13128, 2025.
- [24] Giseop Kim, Sunwook Choi, and Ayoung Kim. Scan Context++: Structural Place Recognition Robust to Rotation and Lateral Variations in Urban Environments. *IEEE Trans. on Robotics (TRO)*, 38(3): 1856–1874, 2022. doi: 10.1109/TRO.2021.3116424.
- [25] Georg Klein and David Murray. Parallel Tracking and Mapping for Small AR Workspaces. In *Proc. Sixth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR’07)*, Nara, Japan, November 2007.
- [26] Pou-Chun Kung, C. Wang, and Wen-Chieh Lin. A Normal Distribution Transform-Based Radar Odometry Designed For Scanning and Automotive Radars. *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 14417–14423, 2021.
- [27] Cedric Le Gentil, Leonardo Brizi, Daniil Lisus, Xinyuan Qiao, Giorgio Grisetti, and Timothy D. Barfoot. DRO: Doppler-Aware Direct Radar Odometry. In *Proc. of Robotics: Science and Systems (RSS)*, 2025.
- [28] Cedric Le Gentil, Raphael Falque, Daniil Lisus, and Timothy D. Barfoot. 2FAST-2LAMAA: Large-Scale Lidar-Inertial Localization and Mapping with Continuous Distance Fields, 2025. URL <https://arxiv.org/abs/2410.05433>.
- [29] Cedric Le Gentil, Weican Li, Leonardo Brizi, and Timothy D. Barfoot. Dr-PoGO: Direct Radar Pose-Graph Optimization. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2026.
- [30] Hyungtae Lim, Kawon Han, Gunhee Shin, Giseop Kim, Songcheol Hong, and Hyun Myung. ORORA: Outlier-robust Radar Odometry. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 2046–2053, 2023.
- [31] Daniil Lisus, Keenan Burnett, David J. Yoon, Richard Poulton, John Marshall, and Timothy D. Barfoot. Are Doppler Velocity Measurements Useful for Spinning Radar Odometry? *IEEE Robotics and Automation Letters (RA-L)*, 10(1):224–231, 2025. doi: 10.1109/LRA.2024.3505821.
- [32] Daniil Lisus, Katya M. Papais, Cedric Le Gentil, Elliot Preston-Krebs, Andrew Lambert, Keith Y. K. Leung, and Timothy D. Barfoot. Boreas Road Trip: A Multi-Sensor Autonomous Driving Dataset on Challenging Roads, 2026. URL <https://arxiv.org/abs/2602.16870>.
- [33] David G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *Intl. Journal of Computer Vision (IJCV)*, 60(2):91–110, November 2004.
- [34] Zhehua Mao, Liang Zhao, Shoudong Huang, Yiting Fan, and Alex P. W. Lee. DSR: Direct Simultaneous Registration for Multiple 3D Images. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, pages 98–107, 2022.
- [35] Alan Papalia, Nikolas Sanderson, Haoyu Han, Heng Yang, Hanumant Singh, and Michael Everett. Sparse Variable Projection in Robotic Perception: Exploiting Separable Structure for Efficient Nonlinear Optimization, 2025. URL <https://arxiv.org/abs/2512.07969>.
- [36] Yeong Sang Park, Young-Sik Shin, and Ayoung Kim. PhaRaO: Direct Radar Odometry using Phase Correlation. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 2617–2623, 2020. doi: 10.1109/ICRA40945.2020.9197231.
- [37] Elliot Preston-Krebs, Daniil Lisus, and Timothy D. Barfoot. The Finer Points: A Systematic Comparison of Point-Cloud Extractors for Radar Odometry. *Proceedings of the Conference on Robots and Vision*, may 27 2025. <https://crv.pubpub.org/pub/urc6alfo>.
- [38] Arthur Venon, Yohan Dupuis, Pascal Vasseur, and Pierre Merriaux. Millimeter Wave FMCW RADARs for Perception, Recognition and Localization in Automotive Applications: A Survey. *IEEE Trans. on Intelligent Vehicles (T-IV)*, 7:1–1, 09 2022. doi: 10.1109/TIV.2022.3167733.