

Motion-Bias-Free Feature-Based SLAM

Alejandro Fontan¹
fontanvillacampa@gmail.com

Javier Civera²
jcivera@unizar.es

Michael Milford¹
michael.milford@qut.edu.au

¹ QUT Centre for Robotics
Queensland University of Technology
Brisbane, Australia

² School of Engineering
University of Zaragoza
Zaragoza, Spain

Abstract

For SLAM to be safely deployed in unstructured real world environments, it must possess several key properties that are not encompassed by conventional benchmarks. In this paper we show that SLAM *commutativity*, that is, consistency in trajectory estimates on forward and reverse traverses of the same route, is a significant issue for the state of the art. Current pipelines show a significant bias between forward and reverse directions of travel, that is in addition inconsistent regarding which direction of travel exhibits better performance. In this paper we propose several contributions to feature-based SLAM pipelines that remedies the motion bias problem. In a comprehensive evaluation across four datasets, we show that our contributions implemented in ORB-SLAM2 substantially reduce the bias between forward and backward motion and additionally improve the aggregated trajectory error. Removing the SLAM motion bias has significant relevance for the wide range of robotics and computer vision applications where performance consistency is important.

1 Introduction

Simultaneous Localization and Mapping (SLAM) refers to the concurrent estimation of a mobile robot state and a model of the environment, known as the *map*, from its onboard sensor readings. Visual SLAM utilizes video as the main input to build a 3D map of an unknown environment, which is then used in another processing thread to determine the camera pose within the environment in real-time [1].

The progress in this field has been remarkable over the past three decades, leading to the extensive adoption of SLAM in real-world applications across various industries (e.g., robotics, autonomous driving or augmented and virtual reality). Starting from the introduction of probabilistic formulations for Simultaneous Localization and Mapping (SLAM) [2, 3, 4, 5], researchers have investigated the fundamental properties of SLAM, such as observability, convergence, and consistency [6, 7, 8, 9, 10, 11, 12, 13]. Current pipelines

© 2023. The copyright of this document resides with its authors. It may be distributed unchanged freely in print or electronic forms.

This research was partially supported by funding from ARC Laureate Fellowship FL210100156 to MM and the QUT Centre for Robotics, the Spanish Government under Grant PID2021-127685NB-I00 and TED2021-131150B-I00 and the Aragon Government under Grant DGA_FSET45_20R.

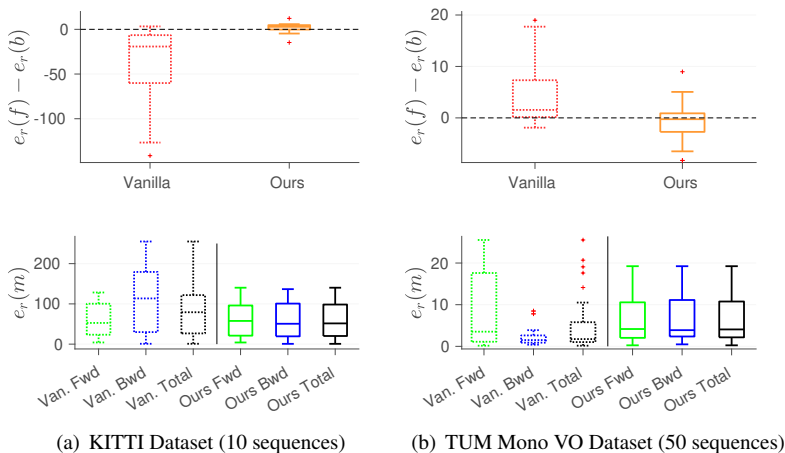


Figure 1: Motion bias and errors in KITTI [20] (a) and TUM Mono VO [15] (b) datasets. **Top row displays the motion bias** as the arithmetic error difference between forward (f) and backward (b) passes. Bigger \pm values represent higher motion errors. Note in the box-and-whiskers plots the substantial biases in the Vanilla ORB-SLAM2 version, and how our implementation is able to remove them. **Bottom row shows forward (green), backward (blue) and total (black) errors** for Vanilla ORB-SLAM2 and our version. Note that our errors are similar for forward and backward passes, in contrast to substantial differences in Vanilla ORB-SLAM2. Moreover, observe the opposite bias sign in KITTI and TUM Mono VO, indicating data dependency.

[14, 16, 19, 27, 28, 32, 35] show an impressive performance in terms of accuracy and robustness. However, although successful implementations in certain domains might rise the question “*is fundamental research in SLAM still necessary?*”, challenges still remain, particularly in complex and dynamic environments [3, 9, 36]. Achieving truly **robust performance** is critical for a wide range of applications. To ensure that SLAM can provide reliable perception and navigation for long-lived autonomous robots, fundamental aspects such as *self-tuning*, *resource awareness*, and *task-driven perception* are yet to be addressed [6].

In this paper we address *motion bias*, a particular problem that mainly affects feature-based visual SLAM. Several works [15, 16, 37] and our own results show that variations in the direction of the camera motion (i.e., playing the same image sequence forwards and backwards) lead to significant differences in the performance of SLAM methods. Figure 1 contains experimental results showing this bias in KITTI [20] and TUM Mono VO [15].

The existence of this motion bias in SLAM pipelines is relevant and should be addressed. On the one hand, motion bias should not appear from an optimization perspective, as the geometric configuration and scene appearance do not change when playing an image sequence forwards or backwards. This reveals flaws in current feature-based pipelines, and lack of explainability of their inner working modes. It may also indicate that they are tuned to overfit benchmarks and reported results do not generalize in more general setups.

On the other hand, bias-free forward and backward estimates are desirable for SLAM to generalize across a wide range of use cases. We can cite several application domains in which agents perform varied motions, such as robots equipped with a front-facing and a rear-facing camera [1, 25, 26], highly maneuverable robots such as vacuum cleaners, drones

or robotic arms, and systems operating in environments with strong motion constraints, such as mine exploration or endoscopic surgery [60].

It is true that there might be applications with dominant motion patterns, such as autonomous driving. However, it should be remarked that the motion bias in current SLAM pipelines is neither consistent nor predictable, impacting in their reliability and trustworthiness in critical applications. Figure 1 shows how ORB-SLAM2 presents opposite biases in KITTI and TUM Mono VO: forward motion estimates are more accurate in KITTI, whereas backward ones are more accurate in TUM Mono VO.

In this paper we make the following contributions: (1) we extend previous evaluations [15, 16, 37] that characterized the bias present in feature-based SLAM for forward and backward camera motion, (2) we identify potential sources that may cause such bias and provide theoretical insights and empirical analysis for a better design of new methods, (3) we implement several novel contributions for motion bias removal on the state-of-the-art ORB-SLAM2 [27], specifically in a deterministic version¹ to facilitate ablation studies and analyses, and (4) we evaluate our proposals in 4 different and complementary datasets. Our results show that we are able to reduce drastically the motion bias in feature-based SLAM, while keeping the general accuracy.

2 Related Work

Motion Bias was unveiled by Engel et al. [16], that showed the difference in accuracy of feature-based monocular ORB-SLAM2 [27] between forward and backward runs in the TUM Mono VO Dataset [15]. They arrived at three significant outcomes: 1) They showed that ORB-SLAM2 performed significantly better for backwards motion, 2) that, contrary to ORB-SLAM2, their direct sparse odometry (DSO) was largely unaffected by motion bias, and 3) that image resolution was a relevant factor for ORB-SLAM2, while DSO was only marginally affected due to subpixel accuracy. From these outcomes they also concluded that benchmarking on large datasets, covering a diverse range of environments and motion patterns, was of high importance. While Engel et al. raised these issues, no analysis, conclusion, or possible remedies were given in their work.

Yang et al. [37] performed systematic and quantitative evaluations of motion bias on the three most popular SLAM formulations, namely direct, feature-based and semi-direct methods. Their initial findings agreed with those of Engel et al. [16]. They claim that the performance degradation of feature-based SLAM come from implementation details, and proposed that future work should take into consideration the following issues: *depth representation*, *point sampling strategy*, *point management* and *discretization artifacts*. They implemented a version of ORB-SLAM2 with sub-pixel matching accuracy that reduced the errors of the forward passes but maintained a substantial gap between forward and backward motions. Motion bias remained hence unsolved, and Yang et al.'s final conclusion highlighted the importance of the problem, stating that more efforts should be taken to address such problem for the adoption of feature-based SLAM in safety-critical applications like autonomous driving. In this paper we continue their work by formulating and evaluating six novel contributions that realistically minimize motion bias with minimal impact in SLAM general accuracy.

¹Our deterministic implementation of ORB-SLAM2 [27] described in Section 4 can be found in https://github.com/alejandrofontan/ORB_SLAM2_Deterministic.

3 Motion-Bias-Free Feature-Based SLAM

In this section, we formulate the contribution of visual covariances to motion bias and introduce an alternative robust cost function for outliers. We also address the sources of the motion bias by providing theoretical insights and proposing motion-bias-free models for five components within feature-based SLAM, specifically related to *point representation*, *optimization* and *data association*.

3.1 Point Representation

Feature-based SLAM estimates a sparse map of 3D points $\{\mathbf{p} \in \mathcal{P} \mid \mathbf{p} \in \mathbb{R}^3\}$ and keyframes $k \in \mathcal{K}$. Keyframes are composed by a grayscale image \mathbf{I}_k and its 6DOF camera pose, represented by a rotation matrix $\mathbf{R}_k \in \text{SO}(3)$ and translation vector $\mathbf{t}_k \in \mathbb{R}^3$. Each point \mathbf{p} is associated with a set of 2D keypoints $\mathcal{U} \equiv \{\mathbf{u}_{k1}, \mathbf{u}_{k2}, \dots, \mathbf{u}_{kn}\}$, $\mathbf{u}_{ki} \in \Omega_{ki}$ with corresponding feature descriptors $\mathcal{B} \equiv \{\mathbf{b}_{k1}, \mathbf{b}_{k2}, \dots, \mathbf{b}_{kn}\}$, $\mathbf{b}_{ki} \in \mathbb{R}^d$. The keypoints correspond to the projections of \mathbf{p} in the keyframes \mathcal{K} in which it is visible, and Ω_{ki} are their image domains.

Reference Descriptor Selection. To enhance the matching process and acquire more observations with new keypoints, a common technique is to choose a distinctive reference descriptor \mathbf{b}_r and match subsequent keypoints against it [24, 27, 28]. Traditionally, the reference descriptor is selected based on image appearance. For instance, one approach is to choose the descriptor with the least median Hamming distance d_H to the remaining descriptors: $\mathbf{b}_r = \text{argmin}_{\mathbf{b}_{ki}} \text{median}(d_H(\mathbf{b}_{ki}, \mathcal{B}))$ [27]. In contrast, we propose to leverage 3D geometry by selecting the descriptor with the smallest Euclidean distance d_E between the translation vectors of its holding keyframe \mathbf{t}_k and the keyframe \mathbf{t}_i containing the keypoint being matched: $\mathbf{b}_r = \text{argmin}_{\mathbf{b}_{ki}} d_E(\mathbf{t}_k, \mathbf{t}_i)$. With this policy, we aim to enhance matching quality by utilizing the descriptor that exhibits closest geometric proximity.

Point Appearance Invariance. In order to reduce the number of incorrect observations resulting from the matching process, we apply a geometric constraint that is robust to motion bias to filter the query keypoints u_q . The image keypoints \mathcal{U} are detected at different octaves $l \in (0, l_{max} \in \mathbb{Z})$ corresponding to scaled images, where $w(l) = s^l \cdot w_0$ and $h(l) = s^l \cdot h_0$. Here, s represents the scaling factor, and w_0 and h_0 are the widths and heights at the highest resolution. To establish the appearance invariance interval for the depth z_q of the query keypoints u_q , we set a maximum change in the octave of the keypoints denoted by Δl :

$$z_q \in [\max_k(z_k \cdot s^{-\Delta l - 0.5}), \min_k(z_k \cdot s^{\Delta l + 0.5})], \quad (1)$$

where z_k is the depth of the point \mathbf{p} with respect to the keyframe k (see Figure 2(a)).

3.2 Optimization

Visual Covariances contribution to motion bias. The camera pose optimization in feature-based SLAM is based on minimizing a weighted residual $\mathbf{r}\sigma_r^{-2}\mathbf{r}^\top$, in which \mathbf{r} and σ_r^2 are respectively the reprojection error and its covariance. In this section we formulate the contribution to motion bias of the standard approximation of visual covariances and propose a new cost function that is robust to motion bias.

We identify an *observation* as the projection of a 3D point \mathbf{p} in two different keyframes $\{i, j\} \in \mathcal{K}$. We label with subscript j the reference frame, and with i any other frame from which the point is visible. The image coordinates of the projection of \mathbf{p} in reference frame j and its depth are denoted as $\mathbf{u} \in \Omega$ and $z \in \mathbb{R}$ respectively (we drop subindex j for simplicity).

The function $\varphi(\mathbf{u})$ projects a point \mathbf{p} from its camera coordinates \mathbf{u} in the reference frame j into the frame i ,

$$\varphi(\mathbf{u}) = \Pi(\mathbf{R}\Pi^{-1}(\mathbf{u}, z) + \mathbf{t}), \quad (2)$$

where $\Pi(\mathbf{p})$ (determined by the intrinsic camera parameters) projects \mathbf{p} in the image; and $\Pi^{-1}(\mathbf{u}, z)$ back-projects the image point with coordinates \mathbf{u} at depth z . $\mathbf{R} \in \text{SO}(3)$ and $\mathbf{t} \in \mathbb{R}^3$ stand here for the relative rotation and translation between keyframes j and i .

Feature-based errors are usually defined as variations of the following expression:

$$\mathbf{r} = \mathbf{u}_i - \varphi(\mathbf{u}), \quad \sigma_r^2 = \sigma_{u_i}^2 + \varphi(\sigma_{u_i}^2), \quad (3)$$

where $\mathbf{u}_i \{ \mathbf{u}_i \in \mathcal{U} \mid \mathbf{u}_i \in \mathbb{R}^2 \}$ stands for the feature point in image i and $\varphi(\mathbf{u})$ for the corresponding point in the frame j reprojected in the frame i . Hence the residual covariance is expressed as the sum of the projection covariance and the feature subpixel noise $\sigma_{u_i}^2$.

We approximate the residual function (3) by its **first-order Taylor approximation**

$$\varphi(\mathbf{u} + d\mathbf{u}) \approx \varphi(\mathbf{u}) + \nabla_{\mathbf{u}}\varphi d\mathbf{u}, \quad \sigma_r^2 \approx \sigma_{u_i}^2 + \nabla_{\mathbf{u}}\varphi \sigma_{u_i}^2 \nabla_{\mathbf{u}}\varphi^T, \quad (4)$$

with the **perspective deformation gradient tensor** $\nabla_{\mathbf{u}}\varphi$ [14] being the Jacobian of the projection function (2).

Assuming pure forward/backward movement, $\mathbf{R} = \mathbf{I}_3$, $\mathbf{t} = (0, 0, t)^T$, the feature subpixel noise, the gradient tensor and the residual covariances can be modeled as isotropic Gaussians,

$$\sigma_r^2 = \sigma_r^2 \mathbf{I}_2 = (\sigma_{u_i}^2 + \varepsilon_{ij}^2 \sigma_{u_j}^2) \mathbf{I}_2, \quad \text{with} \quad \sigma_{u_i}^2 = \sigma_{u_i}^2 \mathbf{I}_2, \quad \sigma_{u_j}^2 = \sigma_{u_j}^2 \mathbf{I}_2, \quad \text{and} \quad \nabla_{\mathbf{u}}\varphi = \varepsilon_{ij} \mathbf{I}_2. \quad (5)$$

The **standard approach for visual covariance** is to approximate the residual covariance of the keypoint as $\sigma_r^2|_{appr} = 2\sigma_{u_i}^2$ [15]. Compared against the value from the linearization

$$\alpha = \frac{\sigma_r^2|_{appr}}{\sigma_r^2} = \frac{2\sigma_{u_i}^2}{(\sigma_{u_i}^2 + \varepsilon_{ij}^2 \sigma_{u_j}^2)} = \frac{2}{(1 + \varepsilon_{ij}^2)} \begin{cases} > 1 & \text{if } \varepsilon_{ij}^2 < 1, \text{ over-estimation} \\ < 1 & \text{if } \varepsilon_{ij}^2 > 1, \text{ under-estimation.} \end{cases} \quad (6)$$

this leads to over- and under-estimation effects that result in different outlier rejection criteria at forward and backward motions. To address this issue, we compute the reprojection error in both the projection and reference keyframes (assuming constant depth) and normalize them with the covariance of the keypoints in their respective views

$$\mathbf{r}_i = (\mathbf{u}_i - \varphi(\mathbf{u}))^2, \quad \mathbf{r}_j = (\mathbf{u} - \varphi(\mathbf{u}_i))^2, \quad \sigma_{r_i}^2 = 2\sigma_{u_i}^2, \quad \sigma_{r_j}^2 = 2\sigma_{u_j}^2. \quad (7)$$

Figure 2(b) compares the robustness of the standard approach and our reprojection function against motion bias

$$\alpha = \frac{\sigma_r^2|_{appr}}{\sigma_r^2} = \frac{2\sigma_{u_i}^2 + 2\sigma_{u_j}^2}{(1 + \varepsilon_{ij}^2)\sigma_{u_j}^2 + (1 + \varepsilon_{ji}^2)\sigma_{u_i}^2} \stackrel{\sigma_{u_i}^2 = \sigma_{u_j}^2}{=} \frac{4}{2 + \varepsilon_{ij}^2 + \varepsilon_{ji}^2}. \quad (8)$$

Outlier Rejection. ORB-SLAM2 [16] and DSO [17], representative pipelines for feature-based and photometric SLAM, implement two strategies to reject outliers. Firstly, they minimize a Huber cost function [18] of the geometric and photometric errors respectively in order to downweight their influence. And secondly, both approaches remove point observations that have a Mahalanobis residual above a certain threshold after optimization. Although




Figure 2: **Left:** We show the **depth invariance interval** (1) for every octave of a point found at $z = 2(m)$. **Right:** Relative error in the estimation of visual covariances (6), (8).

this second strategy may have benefits, it may have a negative effect if removal is done too early. For instance, for a frame with a large number of noisy observations (due to dynamic objects, motion blur or lighting changes), accurate matches can suffer from early removal due to inaccurate initial estimates.

Early removal of potential outliers may not have a significant influence for photometric methods, as observations are abundant and very efficiently computed by simply retrieving the intensity values of the image points [18]. However, it can degrade the performance of feature-based methods. As features have a high extraction footprint, their re-detection after removal comes with certain delay. In order to avoid this effect, we do not remove potential outliers and keep all observations in the optimization.

3.3 Data Association

Symmetric Data Association. The implementation of data association varies depending on its specific function within the SLAM pipeline. The left part of Figure 3 provides pseudo-code for the tracking and mapping threads of ORB-SLAM2 [27]. Design choices differ based on whether the association aims to create new map points (i.e., *Search For Triangulation*), search for new observations (i.e., *Search By Projection*), or fuse map points (i.e., *Fuse*).

The ray from a point observation $\mathbf{u}_i \in \mathbb{R}^3$ on an image $I \in \{i, j\}$ is $\mathbf{r}_i = R_i^{wc}(K^{-1}\mathbf{u}_i)$, where $K^{-1} \in \mathbb{R}^{3 \times 3}$ contains the camera intrinsics and R_i^{wc} the rotation matrix of the corresponding camera pose. The parallax angle between two rays is $\cos(\alpha) = (\mathbf{r}_i \cdot \mathbf{r}_j) / (|\mathbf{r}_i| |\mathbf{r}_j|)$. Figure 3 shows the angle between point observations. Despite the angle at the end of both forward and backward passes being the same, it varies throughout the trajectory. Consequently, if a threshold for the angle is set during point initialization, the number of past and future observations may differ between the forward and backward passes. This discrepancy becomes more prominent when matching functions utilize distinct parameters or heuristics. These variations can result in a different number of observations during forward and backward motion, ultimately leading to differences in performance.

In our modification of ORB-SLAM2, we have changed the implementation of association functions to ensure the utilization of consistent geometry and appearance constraints.

Robust Data Association. A common approach to expedite data association is to employ greedy strategies, where the matching of keypoints is performed strictly in sequential order. This sequential ordering can lead to biased matching results, as keypoints may be ordered based on their position in the image or their resolution octave. We have modified the matching strategy in the ORB-SLAM2 code to avoid its potential bias. We prioritize matches based on the lowest Hamming distance instead of relying on sequential ordering.

```

1: function TRACKING ( $\mathcal{I}$ )
2:      $\triangleright \mathcal{I} = \text{RGB image}$ 
3:     if Not Initialized then
4:         Try to initialize ( $\mathcal{I}$ )
5:     end if
6:     Track With Motion Model ()
7:         Search By Projection v1 ( $c_1$ )
8:     Track Local Map ()
9:         Search By Projection v2 ( $c_2$ )
10:    if Need New KeyFrame () then
11:         $\mathcal{K}_n \leftarrow$  Create New KeyFrame ()
12:    end if
13: end function
    
```

```

1: function MAPPING ( $\mathcal{K}_n$ )
2:      $\triangleright \mathcal{K}_n = \text{new keyframe}$ 
3:      $\mathcal{P} \leftarrow$  Map Point Culling ( $\mathcal{P}$ )
4:      $\mathcal{P}_n \leftarrow$  Create New Pts ( $\mathcal{K}_n, \mathcal{K}$ )
5:         Search For Triangulation ( $c_3$ )
6:     Search In Neighbors ()
7:         Fuse ( $\mathcal{P}_n, \mathcal{K}, c_4$ )
8:         Fuse ( $\mathcal{P}, \mathcal{K}_n, c_4$ )
9:      $\mathcal{P} \leftarrow \mathcal{P} \cup \mathcal{P}_n, \mathcal{K} \leftarrow \mathcal{K} \cup \mathcal{K}_n$ 
10:    Local Bundle Adjustment ( $\mathcal{P}, \mathcal{K}$ )
11:    Erase Outliers ()
12:     $\mathcal{K} \leftarrow$  KeyFrame Culling ( $\mathcal{K}$ )
13: end function
    
```

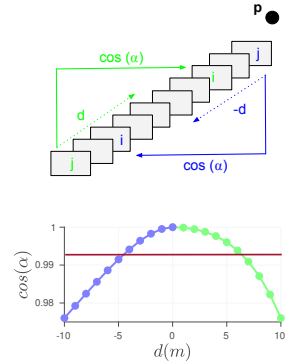


Figure 3: **Left: Tracking and mapping threads** of [27] illustrating the distribution of data association functions throughout a standard pipeline (underlined in the pseudo-code). $c_{1,2,3,4}$ represents variations in the chosen descriptor distance thresholds used for feature matching. **Right: Asymmetry in the angle** between observations of a point under fwd/bwd motion.

4 Ablation Study

In this section, we evaluate quantitatively how the methods proposed in Section 3 contribute to the reduction of motion bias. Specifically, we examine the effects of *descriptor selection*, *point appearance invariance*, *visual covariances*, *outlier rejection*, *symmetric data association*, and *robust data association*. Following [67], we evaluate the accuracy of the trajectories by calculating an **alignment error**, denoted as e_r . This error is estimated as the translational root-mean-squared-error (RMSE) between the estimated trajectory, after it is aligned to the start and end segments of the ground truth trajectory. For further details, please refer to [45]. To quantify **the motion bias**, we define it as the difference between the alignment error of the forward pass and the alignment error of the backward pass: $e_r(f) - e_r(b)$.

4.1 Implementation Details

Deterministic ORB-SLAM2. SLAM systems are susceptible to non-deterministic effects that can significantly impede their performance and hinder our ability to perform ablation studies on specific processes or variables. In this study, we introduce a modified version of ORB-SLAM2 [27] that has been specifically designed to guarantee full determinism, rendering it highly suitable for conducting ablation studies². Specifically, 1) we have altered both the tracking and mapping threads to ensure that they operate sequentially, thereby eliminating online non-deterministic behavior resulting from the real-time implementation. 2) In addition, we have addressed implicit non-deterministic behavior that arose from dynamic memory allocation and random sample consensus (RANSAC).

Keyframe management. Selecting a subset of representative keyframes in order to enable real-time global bundle adjustment and relocalization is an open problem in SLAM [18, 23, 24, 51, 53, 68]. State-of-the-art approaches often rely on policies that are designed to maximize performance on specific benchmarks, and may degrade when they are tested on out-of-distribution sequences. To ensure homogeneous executions in our forward/backward

²Our deterministic implementation of ORB-SLAM2 [27] maintained all capabilities of the original baseline. However, similar to [45, 67], we deactivated loop closure to maximize the impact of motion bias in our experiments.

	Kitti Dataset [20]						TUM Mono VO Dataset [15]					
	$e_r(f)$		$e_r(b)$		$e_r(f) - e_r(b)$		$e_r(f)$		$e_r(b)$		$e_r(f) - e_r(b)$	
	rms	mean/std	rms	mean/std	rms	mean/std	rms	mean/std	rms	mean/std	rms	mean/std
Ours (full)	72.7	58.3/45.7	71.4	56.8/45.4	6.9	1.4/ 7.1	6.7	4.8/4.7	6.8	5.3/4.3	3.9	-1.3/3.7
wo. Descriptor Sel.	74.4	59.4/47.2	79.0	63.4/49.6	15.1	-3.9/15.3	13.0	5.9/11.7	5.3	4.3/3.2	5.4	-0.5/5.5
wo. Point App. Inv.	64.7	52.7/39.5	71.9	56.7/46.5	9.4	-4.0/8.9	6.7	5.0/4.5	6.2	4.9/3.7	2.8	-0.1/2.9
wo. Robust Assoc.	72.4	57.0/47.1	69.3	57.6/40.6	17.8	-0.6/18.7	5.8	4.5/3.8	8.3	5.8/6.0	4.3	-0.5/4.3
wo. Sym. Assoc.	90.6	72.4/57.3	71.8	59.1/43.0	24.9	13.3/22.2	7.6	4.4/6.2	4.7	4.1/2.4	4.1	-1.3/3.9
wo. Visual Cov.	74.4	59.4/47.1	72.2	59.4/43.5	13.2	3.0/14.1	8.0	5.0/6.2	7.0	5.3/4.6	6.2	0.3/6.3
wo. Outlier Reject.	59.8	50.7/33.4	88.9	69.7/58.2	31.2	-19.0/26.1	15.2	7.3/13.5	3.3	2.4/ 2.4	6.3	3.3/5.4

Table 1: **Ablation Study in KITTI [20] and TUM Mono VO [15].** We report the root-mean-squared-error (RMSE), mean and standard deviation (std) of the trajectory alignment error (e_r) in meters (**all better if closer to zero**). We show results for the forward (f) and backward (b) passes, and underline in red/green the worst/best result for each metric. Values in the last columns evaluate the motion bias effect (**smaller with errors closer to zero**). Every row evaluates a configuration where all contributions except one are implemented.




Figure 4: **Other motion bias metrics.** We show the differences on the total number [#] of map points (Pt), local keyframes (Key), and observation inliers ($Inl.$) between forward (f) and backward (b) passes. As before, our contributions effectively minimize the differences across the three indicators, providing evidence of reduced motion bias.

experiments across four different datasets, we process every new frame as a keyframe but keep only one out of every five together with the latest five keyframes. These numbers are a reasonable compromise in our experiments between computational efficiency and performance.

4.2 Ablation Results

Table 1 illustrates the impact of each proposed method on accuracy and motion bias. We assess the relative contribution of each method by implementing all of them and then conducting separate evaluations by deactivating each method one at a time.

The main factor contributing to motion bias is the outlier rejection, which creates a significant gap between forward and backward runs by artificially improving one pass while degrading the other. Our contributions effectively balance the forward and reverse passes in both datasets, thereby eliminating motion bias.

An alternative approach to evaluate motion bias without relying on camera trajectory ground truth is by comparing key statistics of the SLAM graph. Figure 4 illustrates how our proposed approaches effectively reduce the differences across three indicators, providing evidence of reduced motion bias.





Dataset			Hz		$e_r(f)$		$e_r(b)$		$e_r(f) - e_r(b)$	
			Fwd	Bwd	Van.	Ours	Van.	Ours	Van.	Ours
TUM Mono VO 	Indoor	Walk	40	Medium	6.10	6.74	1.89	6.86	4.51	3.93
RGB-D SLAM 	Indoor	Walk	30	Soft	0.07	0.06	0.03	0.05	0.02	0.01
EuRoC 	Indoor	Drone	20	Soft	0.07	0.16	0.09	0.07	0.04	0.01
KITTI 	Outdoor	Car	10	Strong	76.42	72.72	132.24	71.40	65.67	6.94

Table 2: **Final Evaluation.** Our approach eliminates motion bias in the four datasets (third column in orange) while keeping global accuracy constant. As expected the reduction of motion bias is stronger in Kitti and TUM Mono VO datasets where the motion occurs mainly in the forward/backward direction of the camera. Values are the medians after outlier removal.

Sequence	00	02	03	04	05	06	07	08	09	10
$e_r(f)$ Van.	128.3	38.0	9.8	4.1	65.1	93.2	40.1	123.1	100.4	23.4
$e_r(b)$ Van.	254.9	179.3	6.9	0.9	115.3	111.7	53.3	183.2	120.3	29.9
$e_r(f) - e_r(b)$ Van.	-126.6	-141.3	2.9	3.2	-50.1	-18.4	-13.1	-60.0	-19.9	-6.5
$e_r(f)$ Ours	140.1	27.1	4.1	4.8	72.0	96.0	43.6	102.6	71.6	21.0
$e_r(b)$ Ours	136.6	41.8	2.0	0.7	66.2	100.7	38.9	102.9	59.4	19.3
$e_r(f) - e_r(b)$ Ours	3.5	-14.7	2.0	4.1	5.7	-4.6	4.7	-0.2	12.2	1.68

Table 3: **Per-sequence results in KITTI.** Our motion-bias-free implementation removes the bias ($e_r(f) - e_r(b)$) between forward and backward passes in most sequences.

5 Results

We selected four publicly available datasets, encompassing a wide range of potential sources for motion bias, including *indoor/outdoor* environments, *walking/driving* scenarios, *high/low* video frequencies, and *strong/soft* forward/backward motion.

Table 2 provides an overview of these dataset characteristics and shows the errors and motion bias metrics obtained by our deterministic implementation of ORB-SLAM2. Note first how motion bias appear in the four datasets, evidencing that it is a general problem in SLAM. With our contributions, we successfully eliminate motion bias across all datasets, effectively balancing the accuracy between forward and backward passes. Notably, in the KITTI dataset, motion bias is completely eradicated, leading to a significant improvement in the accuracy of the backward pass. For detailed results per sequence, refer to Table 3.

6 Conclusion & Future Work

In this paper, we are motivated by the significant impact of **motion bias** on state-of-the-art feature-based SLAM systems. Our findings reveal substantial inconsistencies in performance, not only between forward and reverse directions of travel but also in terms of which direction exhibits better performance.

Our study reveals that the primary factor contributing to the observed inconsistencies is the **outlier rejection** process. To address this issue, we propose an alternative policy and introduce a **new model for visual covariances** that effectively removes motion bias. By implementing this and other contributions in ORB-SLAM2, we demonstrate a substantial enhancement in the consistency between forward and backward motion, along with improved average trajectory error in both directions, across four diverse datasets. This novel approach yields more trustable and generalizable performance, which is of great significance in critical applications where consistency of performance is crucial for usability and practicality.

Future work should focus on investigating how outlier management techniques can effectively eliminate motion bias while maintaining the highest level of trajectory accuracy (Section 4.2). Additionally, we will further explore the impact of image resolution (Section 2) and keyframe management strategies (Section 4.1) on inducing motion bias.

References

- [1] Siddharth Agarwal, Ankit Vora, Gaurav Pandey, Wayne Williams, Helen Kourous, and James McBride. Ford multi-av seasonal dataset. *The International Journal of Robotics Research*, 39(12):1367–1376, 2020.
- [2] Tim Bailey and Hugh Durrant-Whyte. Simultaneous localization and mapping (slam): Part ii. *IEEE robotics & automation magazine*, 13(3):108–117, 2006.
- [3] Irene Ballester, Alejandro Fontán, Javier Civera, Klaus H Strobl, and Rudolph Triebel. Dot: Dynamic object tracking for visual slam. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11705–11711. IEEE, 2021.
- [4] Berta Bescos, José M Fácil, Javier Civera, and José Neira. Dynaslam: Tracking, mapping, and inpainting in dynamic scenes. *IEEE Robotics and Automation Letters*, 3(4): 4076–4083, 2018.
- [5] Michael Burri, Janosch Nikolic, Pascal Gohl, Thomas Schneider, Joern Rehder, Sammy Omari, Markus W Achtelik, and Roland Siegwart. The euroc micro aerial vehicle datasets. *The International Journal of Robotics Research*, 35(10):1157–1163, 2016.
- [6] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics*, 32(6):1309–1332, 2016.
- [7] Carlos Campos, Richard Elvira, Juan J Gómez Rodríguez, José MM Montiel, and Juan D Tardós. Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam. *IEEE Transactions on Robotics*, 37(6):1874–1890, 2021.
- [8] Margarita Chli and Andrew J Davison. Active matching for visual tracking. *Robotics and Autonomous Systems*, 57(12):1173–1187, 2009.
- [9] Javier Civera, Andrew J Davison, and JM Martinez Montiel. Inverse depth parametrization for monocular slam. *IEEE transactions on robotics*, 24(5):932–945, 2008.
- [10] Alejo Concha and Javier Civera. An evaluation of robust cost functions for rgb direct mapping. In *2015 European Conference on Mobile Robots (ECMR)*, pages 1–8. IEEE, 2015.
- [11] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *IEEE transactions on pattern analysis and machine intelligence*, 29(6):1052–1067, 2007.
- [12] Gamini Dissanayake, Shoudong Huang, Zhan Wang, and Ravindra Ranasinghe. A review of recent developments in simultaneous localization and mapping. In *2011 6th International Conference on Industrial and Information Systems*, pages 477–482. IEEE, 2011.
- [13] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006.

- [14] Jakob Engel, Thomas Schöps, and Daniel Cremers. Lsd-slam: Large-scale direct monocular slam. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part II 13*, pages 834–849. Springer, 2014.
- [15] Jakob Engel, Vladyslav Usenko, and Daniel Cremers. A photometrically calibrated benchmark for monocular visual odometry. *arXiv preprint arXiv:1607.02555*, 2016.
- [16] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):611–625, 2017.
- [17] Alejandro Fontan, Laura Oliva, Javier Civera, and Rudolph Triebel. Model for multi-view residual covariances based on perspective deformation. *IEEE Robotics and Automation Letters*, 7(2):1960–1967, 2022.
- [18] Alejandro Fontan, Riccardo Giubilato, Laura Oliva, Javier Civera, and Rudolph Triebel. Sid-slam: Semi-direct information-driven rgb-d slam. *IEEE Robotics and Automation Letters*, 2023.
- [19] Christian Forster, Zichao Zhang, Michael Gassner, Manuel Werlberger, and Davide Scaramuzza. Svo: Semidirect visual odometry for monocular and multicamera systems. *IEEE Transactions on Robotics*, 33(2):249–265, 2016.
- [20] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [21] Christian Kerl, Jürgen Sturm, and Daniel Cremers. Robust odometry estimation for rgb-d cameras. In *2013 IEEE international conference on robotics and automation*, pages 3748–3754. IEEE, 2013.
- [22] Georg Klein and David Murray. Parallel tracking and mapping for small ar workspaces. In *2007 6th IEEE and ACM international symposium on mixed and augmented reality*, pages 225–234. IEEE, 2007.
- [23] Kurt Konolige and Motilal Agrawal. Frameslam: From bundle adjustment to real-time visual mapping. *IEEE Transactions on Robotics*, 24(5):1066–1077, 2008.
- [24] Stefan Leutenegger, Paul Furgale, Vincent Rabaud, Margarita Chli, Kurt Konolige, and Roland Siegwart. Keyframe-based visual-inertial slam using nonlinear optimization. *Proceedings of Robotis Science and Systems (RSS) 2013*, 2013.
- [25] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017.
- [26] Lukas Meyer, Michal Smíšek, Alejandro Fontan Villacampa, Laura Oliva Maza, Daniel Medina, Martin J Schuster, Florian Steidle, Mallikarjuna Vayugundla, Marcus G Müller, Bernhard Rebele, et al. The madmax data set for visual-inertial rover navigation on mars. *Journal of Field Robotics*, 38(6):833–853, 2021.

- [27] Raul Mur-Artal and Juan D Tardós. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE transactions on robotics*, 33(5):1255–1262, 2017.
- [28] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5): 1147–1163, 2015.
- [29] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. Dtam: Dense tracking and mapping in real-time. In *2011 international conference on computer vision*, pages 2320–2327. IEEE, 2011.
- [30] David Recasens, José Lamarca, José M Fácil, JMM Montiel, and Javier Civera. Endo-depth-and-motion: reconstruction and tracking in endoscopic videos using depth networks and photometric constraints. *IEEE Robotics and Automation Letters*, 6(4):7225–7232, 2021.
- [31] Patrik Schmuck and Margarita Chli. On the redundancy detection in keyframe-based slam. In *2019 International Conference on 3D Vision (3DV)*, pages 594–603. IEEE, 2019.
- [32] Thomas Schops, Torsten Sattler, and Marc Pollefeys. BAD SLAM: Bundle Adjusted Direct RGB-D SLAM. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [33] Hauke Strasdat, JMM Montiel, and Andrew J Davison. Real-time monocular slam: Why filter? In *2010 IEEE International Conference on Robotics and Automation*, pages 2657–2664. IEEE, 2010.
- [34] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A Benchmark for the Evaluation of RGB-D SLAM Systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012.
- [35] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34:16558–16569, 2021.
- [36] Binbin Xu, Wenbin Li, Dimos Tzoumanikas, Michael Bloesch, Andrew Davison, and Stefan Leutenegger. Mid-fusion: Octree-based object-level multi-instance dynamic slam. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5231–5237. IEEE, 2019.
- [37] Nan Yang, Rui Wang, Xiang Gao, and Daniel Cremers. Challenges in monocular visual odometry: Photometric calibration, motion bias, and rolling shutter effect. *IEEE Robotics and Automation Letters*, 3(4):2878–2885, 2018.
- [38] Georges Younes, Daniel Asmar, Elie Shammas, and John Zelek. Keyframe-based monocular slam: design, survey, and future directions. *Robotics and Autonomous Systems*, 98:67–88, 2017.