Poster: Do Privacy-Preserving Obfuscation Techniques Degrade the Accuracy of Odometry?

Nikolaos Ntokos, Nahin Kumar Dey, Jiayi Meng, Faysal Hossain Shezan
The University of Texas at Arlington
{nikolaos.ntokos,nahin.dey,jiayi.meng,faysal.shezan}@uta.edu

Abstract

On-device sensors in mobile systems, *e.g.,* autonomous vehicles and AR/VR, use odometry for real-time positioning, but they risk capturing sensitive data of non-consenting bystanders. Prior works have investigated various privacy-preserving techniques to protect those sensitive data. However, it is still unclear about the impact of such approaches on the accuracy of odometry. In this work, we investigate the impact of various privacy-preserving obfuscation techniques on the accuracy of monocular visual odometry. We focus on three widely used obfuscation methods: Gaussian Blur, Gaussian Noise, and Laplacian Noise, applied to protect bystander privacy. Our investigation reveals that some obfuscation techniques can increase the odometry errors by up to 56.9%, while others surprisingly reduce the errors by up to 66.8%, compared to raw data. Our key findings indicate that data obfuscation primarily affects the duration of tracking loss in ORB-SLAM3, which is the main source of the errors, and successful relocalization immediately following tracking loss plays a crucial role in reducing the overall errors.

ACM Reference Format:
Nikolaos Ntokos, Nahin Kumar Dey, Jiayi Meng, Faysal Hossain Shezan. 2024. Poster: Do Privacy-Preserving Obfuscation Techniques Degrade the Accuracy of Odometry?. In *The 30th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '24), November 18–22, 2024, Washington D.C., DC, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3636534.3697459

1 Introduction

Odometry plays a crucial role in many mobile systems, *e.g.,* autonomous vehicles and augmented/virtual reality (AR/VR). It relies on data from on-device sensors (*e.g.,* cameras, gyroscopes, and accelerometers) to estimate the device's pose (location and orientation) in real time. However, the sensor data being collected may include sensitive information of individuals who share the same physical space with the device. These individuals, whom we refer to as bystanders, may not consent to having their sensitive information exposed, resulting in severe privacy leakage issues for odometry.

To protect bystanders' privacy, various innovative approaches have been proposed. Researchers have investigated the use of adversarial perturbations to render bystanders unrecognizable to machine learning models while maintaining the visual integrity. However, these perturbations often lack robustness against sophisticated attacks such as denoising or gradient-based methods. Another avenue of research has focused on real-time computer vision techniques to detect and then blur or block human figures or faces in video streams, effectively anonymizing bystanders. For example, GB has been extensively utilized to selectively blur regions in images that contain bystander data (*e.g.,* [2]); and differential privacy has emerged as another one of the most widely adopted solutions, which blocks sensitive data by adding noises, *e.g.,* Gaussian Noise and Laplacian Noise, to protect such data from leakage.

However, prior works (*e.g.,* [2]) have not thoroughly investigated the effect of obfuscating data on the accuracy of fundamental tasks in mobile systems, particularly odometry. Considering odometry heavily relies on precise sensor data to estimate poses, changes in input data *e.g.,* by data obfuscation, can affect the accuracy of odometry. However, it is still unclear how odometry algorithms can be affected by different obfuscation techniques. Understanding this aspect is essential for designing privacy-preserving odometry algorithms that maximize the accuracy while still protecting sensitive information of bystanders.

In this work, we seek to answer the following important question: *do privacy-preserving obfuscation techniques degrade the accuracy of odometry?* To that end, we investigate three widely used obfuscation techniques: Gaussian Blur (GB), Gaussian Noise (GN), and Laplacian Noise (LN). We explore how these techniques and their corresponding parameters can affect the accuracy of monocular visual odometry, *i.e.,* ORB-SLAM3 [1], which estimates the pose of a device

ACM MobiCom '24, November 18–22, 2024, Washington D.C., DC, USA

ACM ISBN 979-8-4007-0489-5/24/11
https://doi.org/10.1145/3636534.3697459

using images captured by a single camera. Surprisingly, we observe that, obfuscating bystanders does not always degrade the accuracy of odometry. Specifically, we find that (1) various parameters of obfuscation techniques affect the accuracy differently for different videos; and (2) GB with the handpicked parameters achieves 5.3%∼66.8% lower errors than the Default case (*i.e.,* without any obfuscation) and outperforms the other techniques, which have up to 56.9% higher errors compared to Default.

We further investigate the reasons for such improvement with . We highlight two key findings: (1) data obfuscation affects the duration of tracking loss in ORB-SLAM3, which accounts for the majority of errors. (2) effective data obfuscation can expedite successful relocalization, which can mitigate the accuracy drop caused by tracking loss.

Finally, we discuss future research directions.

2 Existing Obfuscation Techniques

Gaussian Blur (GB), Gaussian Noise (GN), and Laplacian Noise (LN) are effective obfuscation techniques for odometry purposes due to their ability to distort or mask sensitive information while preserving overall data utility (*e.g.,* [4]). GB works by convolving the image with a Gaussian function, effectively smoothing out details and reducing high-frequency components. The standard deviation (σ) of the Gaussian kernel typically ranges from 0 to 100, with higher values resulting in more aggressive blurring. GN adds random variations to pixel values following a normal distribution, while LN adds noise based on the Laplace distribution. Both techniques are controlled by the parameter ϵ. It usually ranges from 0 to 1.0. Lower ϵ values (*i.e.,* below 0.1) provide stronger privacy guarantees. These techniques effectively obfuscate data by introducing controlled levels of uncertainty or distortion, making it difficult to extract precise information about individual data while maintaining the overall structure and statistical properties of the dataset, thus protecting privacy of bystanders in odometry applications.

3 Experimental Setup

Dataset. We choose ADVIO [3], a dataset for pedestrian odometry, to estimate a person's pose (position and orientation) while walking through varying real-world environments, both indoors and outdoors. We select four video sequences, *i.e.,* A01 (indoor, mall A), A10 (indoor, mall B), A14 (indoor, office) and A20 (outdoor), representing different scenarios with varying numbers of bystanders. Among the four videos, A01 has the largest number of bystanders.

Methodology. To protect the privacy of bystanders, we first run the pre-trained YOLOv10 model to detect bystanders in the videos, considering the entire human body as the object of interest We then apply the three obfuscation techniques to those regions using various parameters, *i.e.,* $\{\sigma_{10}, \sigma_{30}, \sigma_{50}, \sigma_{70}\}$ for GB and $\{\epsilon_{0.1}, \epsilon_{0.01}, \epsilon_{0.001}\}$ for GN and LN. Using the obfuscated images as input, we run ORB-SLAM3 in monocular mode to generate estimated trajectories.

To evaluate the accuracy, for each estimated trajectory, we first align it with the corresponding ground truth trajectory [1]. We then calculate the Absolute Pose Error (APE), one of the most widely used metrics for odometry, which essentially measures the euclidean distance between an estimated pose and its corresponding ground truth pose. We finally report the mean APE for each estimated trajectory.

Baseline (Default). Considering only raw images as input.

4 Results

Gaussian Blur. Figure 1 shows that, surprisingly, GB does not always degrade the accuracy of odometry. Compared with Default, GB reduces errors by up to 5.3%∼66.8% for the four sampled videos. Specifically, GB achieves the lowest mean APE of 2.33 m (σ_{30}), 1.28 m (σ_{70}), 1.79 m (σ_{50}), and 6.33 m (σ_{70}), compared to 2.46 m, 1.89 m, 2.35 m, and 19.05 m for Default, for A01, A10, A14, and A20, respectively.

Figure 1 also shows that different values of σ affect the accuracy differently and the best σ value varies across the four videos. For A01, only σ_{30} decreases mean APE from 2.46 m to 2.33 m. The other three σ values increase the mean APE by 12.9%∼45.2%, from 2.46 m to 2.78 m∼3.58 m. As for A10, σ_{10} has slightly larger mean APE of 1.91 m than 1.89 m, while the other three σ values achieve 12.4%∼32.4% smaller mean APE, *i.e.,* 1.28 m∼1.65 m. For A14, all the σ values achieve 7.7%∼23.8% lower errors, *i.e.,* 1.79 m∼2.17 m, than 2.35 m. As for A20, σ_{30}, σ_{50}, and σ_{70} significantly increase mean APE by 17.1%∼37.2%, from 19.05 m (Default) to 22.30 m∼26.13 m whereas σ_{70} drastically decreases mean APE to 6.33 m.

Gaussian & Laplacian Noise. ❶ Similarly, different ϵ values, impact accuracy differently, with the best ϵ varying across the videos. ❷ Considering the best ϵ for each video sequence, Figure 1 shows that GN performs worse than GB by 7.07%∼217.04%, delivering errors at 2.50 m ($\epsilon_{0.1}$), 1.46 m ($\epsilon_{0.01}$), 1.95 m ($\epsilon_{0.001}$), and 20.08 m ($\epsilon_{0.001}$) for the four sequences respectively, but better than Default for A10 and A14 with mean APE of 1.89 m and 2.35 m. ❸ LN again performs worse than GB in A14 and A20 by 12.6%∼19.5% and slightly better in sequences A01 and A10 by only 1.9%, while outperforming the Default in all four videos by 7.1%∼62.6%. Its lowest errors are 2.29 m ($\epsilon_{0.001}$), 1.25 m ($\epsilon_{0.1}$), 2.15 m ($\epsilon_{0.001}$) and 7.13 m ($\epsilon_{0.01}$) for A01, A10, A14, and A20 respectively.

5 Why Does Gaussian Blur Improve Accuracy for ORB-SLAM3?

Given that GB outperforms the other methods, we next focus on GB and investigate why it improves the accuracy of ORB-SLAM3 compared to Default and how different σ values affect the performance of ORB-SLAM3. Due to page

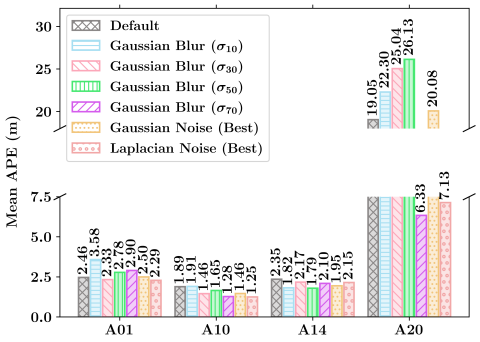

Figure 1: Mean APE comparison.

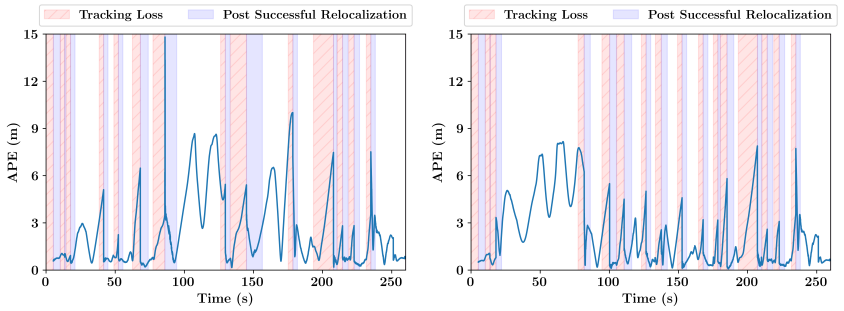

(a) Default, mean APE = 2.46 **(b) Gaussian Blur (σ_{30}), mean APE** = 2.33

Figure 2: APE changes and TRL-PSR periods over time for A01.

Table 1: Comparison of APE spikes, duration in TRL, and mean APE in various time periods for A01.

Methods	Avg. Spike (m)	Duration in TRL (s)	mean APE (m)			
			Entire	TRL	PSR	TRL-PSR
Default	5.58	5.08	2.46	2.84	1.31	2.16
σ_{10}	6.05	4.76	3.58	4.22	1.90	3.29
σ_{30}	**4.54**	**4.44**	**2.33**	**2.65**	**1.05**	**1.95**
σ_{50}	4.99	5.14	2.78	3.42	1.15	2.49
σ_{70}	6.02	5.04	2.90	3.50	1.92	2.75

limit, we elaborate on our findings for A01 in the rest of §5. The other videos share similar results.

We first plot the APE variations throughout the entire timeline of A01 for Default and GB (σ_{30}) , which has the lowest mean APE over all the methods, as shown in Figure 2. Due to page limit, we omit the other σ values. We observe two key common trends across different methods.

First, APE starts to increase when ORB-SLAM3 fails to match enough ORB features of the current frame to previous keyframes, *i.e.,* tracking loss (TRL), and hence loses the ability to generate new poses in real time. This is indicated by the red shaded regions in Figure 2, *e.g.,* 134–145 s (Default) and 77–81.5 s (σ_{30}). The longer the TRL persists, the higher the APE could be, especially when the device is moving.

Second, APE drastically drops after TRL *e.g.,* at 145 s (Default) and 81.5 s (σ_{30}). This is due to successful relocalization. During TRL, ORB-SLAM3 attempts relocalization, essentially searching in its keyframe database for the most similar keyframe to the current frame. When the match is found and checked for temporal and geometric consistency, relocalization is considered successful and a new accurate pose is estimated. However, relocalization cannot be performed for every frame due to its high computational cost.

Since APE is related to those APE spikes caused by TRL and relocalization, to understand the differences between GB (σ_{30}) and the other methods, we next calculate the average of the peak values of the spikes and the duration of TRL per method. Table 1 shows that **compared to all the other methods, GB (σ_{30}) has 9.9%~33.3% lower APE spikes and stays in TRL for 7.2%~15.8% less on average.**

To quantify how much relocalization can offset the accuracy degradation caused by TRL for different methods, we define Post Successful Relocalization (PSR) time periods (light blue regions in Figure 2). PSR begins when relocalization succeeds and continues for the same duration as the last TRL. Note that if another TRL occurs before PSR ends, we only consider the time until this next TRL. We calculate the mean APE in the combined time periods of TRL and PSR, denoted as TRL-PSR. We compare it with the mean APE during TRL. We find that **relocalization significantly mitigates the accuracy drop of tracking loss for all the methods.** In particular, GB (σ_{30}) has a mean APE of 1.95 m for TRL-PSR, much lower than 2.65 m for TRL only and 2.33 m for the entire trajectory; and it has the smallest mean APE in all the time periods, compared to the other methods.

Remark: *Effective data obfuscation can facilitate achieving successful relocalization sooner, which results in lower errors.*

6 Future Plans

In this paper, we take the first step towards examining the effect of data obfuscation on odometry. In the future, we plan to continue exploring– (1) how various obfuscation techniques affect the duration of TRL in ORB-SLAM3 and, consequently, odometry accuracy; (2) how to predict the best performing obfuscation parameters in terms of odometry accuracy; (3) how different odometry algorithms perform with different obfuscation techniques; and (4) how our findings adapt to more challenging scenarios, *e.g.,* more bystanders.

References

[1] Carlos Campos et al. 2021. ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual–Inertial, and Multimap SLAM. *IEEE Transactions on Robotics* 37, 6 (2021), 1874–1890.

[2] Matthew Corbett et al. 2023. BystandAR: Protecting Bystander Visual Data in Augmented Reality Systems. In *Proc. of the ACM MobiSys.*

[3] Santiago Cortés et al. 2018. ADVIO: An Authentic Dataset for Visual-Inertial Odometry. arXiv:1807.09828

[4] Zhigang Gao et al. 2022. Protecting location privacy of users based on trajectory obfuscation in mobile crowdsensing. *IEEE Transactions on Industrial Informatics* 18, 9 (2022), 6290–6299.