GPS DENIED - THE NEXT LEVEL IN AUTONOMOUS NAVIGATION TOWARDS FLYING WITHOUT GPS

Georg Puhlfürst[1], Jürgen Hess[1], Stefan Milz[1]

This paper gives an overview of current challenges in GPS-denied drone navigation, highlighting a crucial advancement towards achieving full autonomy in flights devoid of GPS. It emphasizes the essentiality of GPS-denied solutions for ensuring uninterrupted and safe flight operations in environments where GPS signals are unreliable or absent. The exploration is divided into two principal methodologies: managing sudden GPS loss during flight, and initiating flight without GPS reliance. For scenarios involving GPS disruption mid-flight, the study proposes the employment of visual odometry, utilizing both traditional and deep-learning approaches, to sustain navigation accuracy. Alternatively, for flights planned without GPS from the beginning, we rely on a method of aligning live camera images with satellite imagery. This approach not only enhances navigation reliability in GPS-compromised situations but also sets a new standard for autonomous flight technology.

INTRODUCTION

Traditionally, UAVs and drones rely on GPS guidance for navigating their flight paths. This system is popular and serves as the basis of modern drone navigation as it is highly precise and reliable. However, there exist scenarios where utilizing GPS for navigation proves to be challenging if not impossible. This includes situations such as dense urban environments where high buildings create "urban canyons" that disrupt GPS signals, or scenarios where GPS signals are intentionally jammed or spoofed. Under such conditions, GPS can be lost for varied durations, lost permanently, or altogether absent. Without reliable alternatives, drones risk becoming disoriented or ultimately lost, compromising mission success and operational safety. Given these challenges, our objective is to develop robust GPS-denied navigation solutions, ensuring drones can maintain efficacy across a diverse array of operational hurdles.

In this paper, we address the two pivotal use cases of GPS-Denied navigation: experiencing GPS loss mid-flight and operating entirely void of GPS. For the first use case we employ visual odometry techniques, both classical and deep-learning-based. These methods enable drones to estimate their movement continuously, from one time step to the next by analyzing subsequent camera images. For the second use case, navigating entirely without GPS, we combine this visual odometry with an image based global localization. This allows drones to estimate their position within a larger area by matching onboard camera images with pre-existing maps or satellite imagery. Together, these strategies can significantly increase drone autonomy and effectiveness across diverse operational landscapes. Figure 3a shows an example flight.

[1] Spleenlab GmbH

https://doi.org/10.52202/075106-0017

APPROACHES FOR GPS-DENIED NAVIGATION

In this section, we detail our solution to the two GPS-Denied use cases, a sudden loss of GPS during flight and navigating entirely without GPS

Use Case 1: Losing GPS During Flight

For drones that lose the GPS signal mid-flight, we advocate the implementation of visual odometry techniques. These methods, which can be either classical or based on deep learning, allow drones to estimate their movement over time by analyzing the visual differences between sequential images captured by onboard cameras.

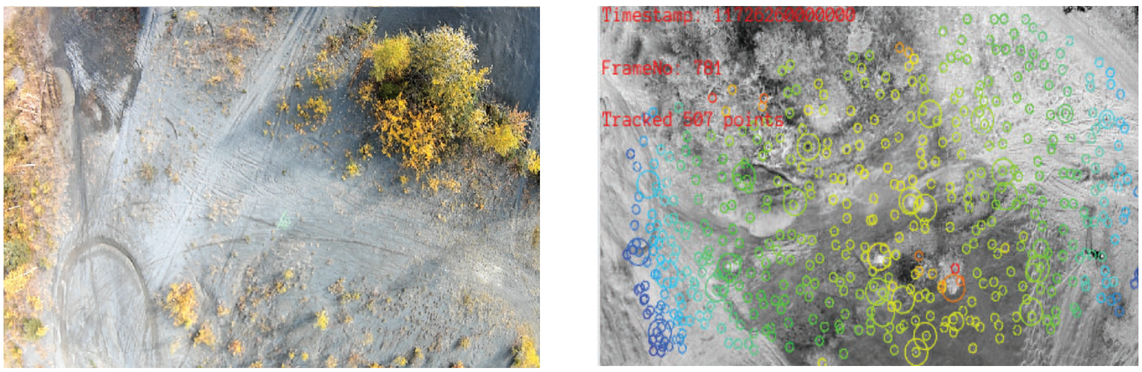

Figure 1: An example image recorded during the flight tracked in Figure 3 (left). An example of image features tracked during flight (right). The color symbolizes the distance to the drone. Dark blue represents features furthest away from the drone camera while dark red mark features closest to the camera.

Our proposed solution employs a monocular camera system for UAV navigation. This approach, however, introduces the challenge of accurately estimating the scale of movement, a task that is inherently complex due to the reliance on a single viewpoint for depth perception. While stereo camera systems might seem like an alternative, offering direct depth information through the use of two cameras, they fall short in aerial applications where the UAV's flight altitude significantly exceeds the baseline, i.e. the distance between the two cameras. In such cases, the stereo system's effectiveness in depth estimation diminishes, rendering it inadequate for large-scale depth perception.

To calibrate the scale for the monocular camera, we utilize GPS data when it is available, e.g., at the start of the flight. While the GPS data is available, we already run and analyze the current scaling of the visual odometry output. This is done by matching the poses that are computed by the visual odometry system to the drone positions as received from GPS. As a result, once the GPS is failing, we can leverage the previously computed scale information together with the continuous visual odometry for calculating an estimated GPS location. An overview of how our approach integrates into the entire navigation system is shown in Figure 2.

To implement monocular visual odometry, we leverage two approaches depending on the use environment the drone is operating in and the computational resources available.

The first system utilizes classical features to compute optical flow between images, such as ORB (Oriented FAST and Rotated BRIEF)[1], which are renowned for their efficiency and low computational cost. ORB excels in identifying and matching distinctive points between

https://doi.org/10.52202/075106-0017

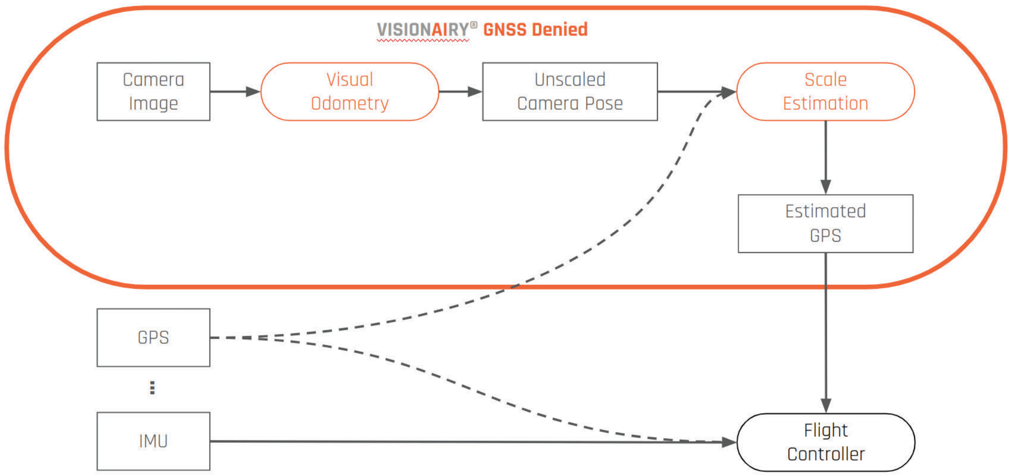

Figure 2: System overview for the use case of losing GPS during the flight. At times where GPS is available, this information can be used to scale the monocular visual odometry. Once the GPS fails, an estimated GPS location can be sent to the flight controller that can also integrate additional information such as IMU or wind sensor data.

consecutive images, enabling the drone to estimate its movement by comparing these points over time. We enhance this with additional steps that help refine the movement and camera position estimation. After obtaining feature correspondences between images, we apply a graph-based optimization together with a keyframe selection scheme to estimate the movement and pose of the camera. This method is particularly effective in environments where the visual features remain consistent and well-defined, allowing for precise movement estimation without the need for intensive computational resources.

Figure 3 shows an example result of applying feature-based visual odometry inside the GPS-Denied system. The blue dot marks the start of the flight. In the beginning, GPS information is still available. We simulate a GPS cutoff at the position of the orange dot. The resulting trajectories are shown in green (true location as received from the receiver) and red (estimated GPS location after the simulated cutoff). As is visible from the figure, the algorithm is highly accurate. While experiencing some expected drift, the estimated trajectory stays close to the measured location.

The second system makes use of neural networks and is based on Deep Pose Visual Odometry (DPVO)[2] to infer the drone's pose and movement from raw image data, offering highly accurate results in complex environments. Unlike classical methods, which rely on explicit feature detection and matching, DPVO learns to recognize patterns and correlations in the visual data, enabling it to adapt to a wider range of visual scenarios. This adaptability makes DPVO particularly suitable for environments where traditional feature-based methods struggle, such as areas with limited distinct visual markers or in varied lighting conditions. However, it does require more computational resources, i.e., the usage of an onboard GPU which also leads to a higher power consumption. In addition, DPVO utilizes dense depth data for training which is hard to obtain, especially for UAV applications. In another work, we have shown how to train DPVO with minimal supervision and thus mitigate this challenge.[5]

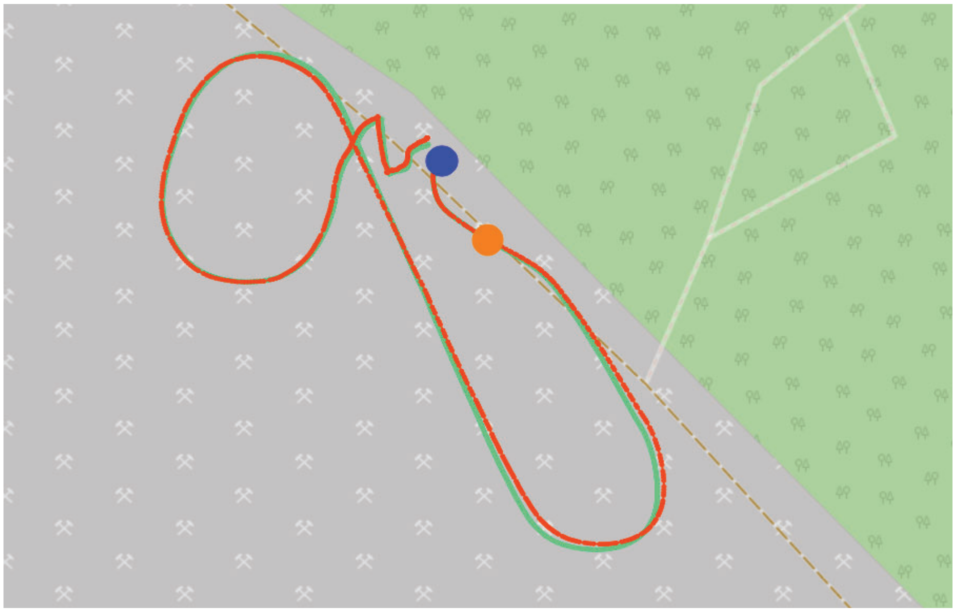

a) Example trajectory flown over a quarry. Plotted are the true GPS trajectory (green) and the estimated trajectory (red). The flight time of the drone was approximately 4.5 minutes. The start is marked in blue, the simulated GPS cutoff is marked in orange. The estimated trajectory stays close to the true trajectory even through the sharp turn at the end.

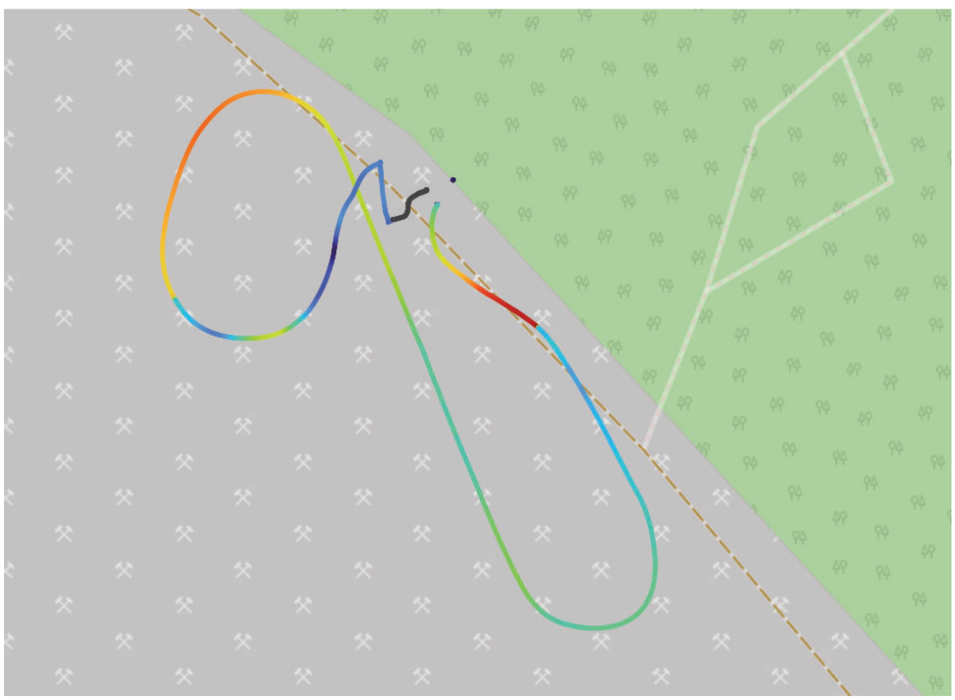

b) Visualization of the absolute trajectory error (ATE) over the entire trajectory. Dark red marks the largest error of about 27m. The mean ATE over the part of the entire flight is ~15m.

Figure 3: Example of a GPS-Denied trajectory flight over a quarry plotted onto an OpenStreetMap.

https://doi.org/10.52202/075106-0017

Use Case 2: Completely Navigating Without GPS

In scenarios where drones are required to operate entirely without GPS, combining visual odometry with global localization techniques offers a comprehensive solution. This dual approach enables drones not only to track their immediate movement through visual odometry but also to locate their position within a larger mapped area. By matching real-time imagery with pre-existing GPS indexed maps or satellite images, drones can achieve a level of precision navigation that would otherwise be unattainable without GPS. Figure 4 shows an example of matching an image obtained from sensory input during the flight to a Google maps aerial image.

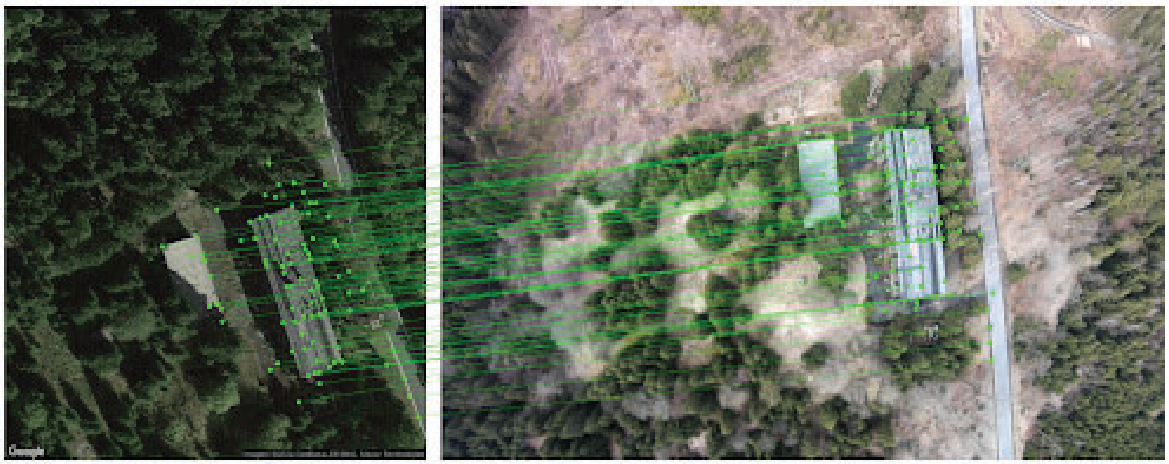

Figure 4: Example of matching visual features of a camera view (right) to Google satellite imaginary (left). The matching proves robust to different viewpoints and even larger vegetation changes.

To implement this matching process, we employ cutting-edge deep-learning based algorithms such as SuperGlue[3] and LightGlue[4], which are exceptional in their ability to identify and match features between the drone's immediate visual input and the corresponding sections of pre-stored maps or satellite images. We have tuned and highly optimized these algorithms for the specific hardware onboard the drone, allowing a processing speed beyond the current state-of-the-art. We then leverage the found feature matches to compute the current GPS position of the drone and thus realize global localization. While the matching algorithm is highly optimized, it yields a lower frame rate than the visual odometry. While visual odometry and global localization can run independently, we propose a hybrid system. This system runs the visual odometry and estimates the relative movement of the drone in addition to the global localization runs to provide a continuous estimate to the flight controller. Figure 5 shows the integration of this approach into the system control view.

This approach, however, does come with its set of challenges. First, its reliance on a GPU for processing means that the energy consumption are factors that must be carefully managed to ensure prolonged operational periods. Furthermore, for the matching process to be effective, the drone needs to maintain a sufficient altitude, ensuring it captures a broad enough view of the surrounding environment to match with the aerial data. This requirement on a high altitude poses limitations on the flight path in densely populated areas, environments with significant vertical obstructions, or other areas where high altitude flight is highly regulated. Additionally, aerial image data, comparable in altitude to that at which the drone is flying, must be available, processed beforehand, and stored onboard. This introduces constraints regarding data storage and

https://doi.org/10.52202/075106-0017

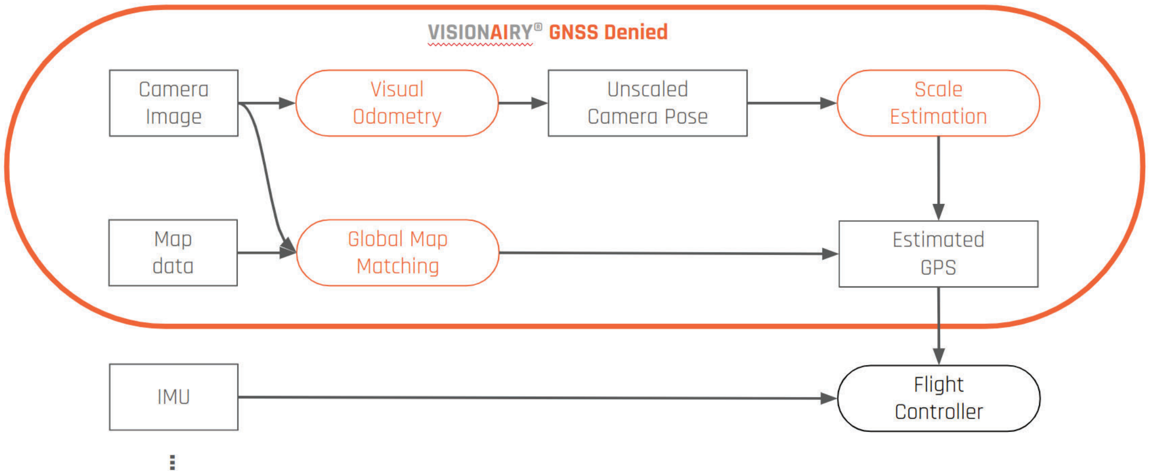

Figure 5: System overview for the use case of navigating completely without GPS. The global map matching typically runs at a lower frequency, providing global GPS locations. The visual odometry estimates the local movement between the global matching runs. The flight controller integrates the estimated GPS locations together with additional sensor information such as IMU data.

management. Despite these challenges, the advantages offered by this method, such as the ability to navigate with high precision in GPS-denied environments, make it a promising avenue for enhancing drone autonomy and operational capabilities.

GPS-DENIED NAVIGATION COMPETITION

To benefit the community with valuable datasets as well as advance our own research, we are launching a competition that challenges participants to navigate drones without GPS by solely relying on visual data from three onboard cameras. Hosted as part of the CVPR OmniCV2024 Workshop, the aim is to develop an algorithm capable of accurately calculating the drone's flight path in GPS-denied scenarios. The competition, named "Omni-BALLOON" in homage to the historic East Germany balloon escape, underscores the transformative potential of camera-assisted navigation. This endeavor not only fosters innovation but also provides a platform for showcasing novel solutions in the realm of autonomous flight. CVPR OmniCV2024 Workshop. We encourage everyone to participate in showcasing and comparing your solution through Kaggle[2].

CONCLUSION

GPS-Denied drone navigation in real-time is extremely challenging. We have presented approaches for two use cases, losing GPS mid-flight and flying the entire mission without GPS. In addition, through our "Omni-BALLOON" competition, part of the CVPR OmniCV2024 Workshop, we offer the community valuable datasets for fostering innovation in autonomous flight technology. There are still a number of challenges ahead, including flying in low-light or night time conditions and flying robustly without GPS for several hours at a time. However, through these efforts, we lay the groundwork for future advancements in drone navigation, moving closer to reliable operations in environments where GPS is compromised or unavailable.

[2] https://www.kaggle.com/competitions/omni-balloon

https://doi.org/10.52202/075106-0017

REFERENCES

[1] E. Rub2lee, V. Rabaud, K. Konolige and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," *2011 International Conference on Computer Vision*, Barcelona, Spain, 2011, pp. 2564-2571

[2] Teed, Zachary, Lahav Lipson, and Jia Deng. "Deep patch visual odometry." *Advances in Neural Information Processing Systems* 36 (2024).

[3] Sarlin, Paul-Edouard, et al. "Superglue: Learning feature matching with graph neural networks." *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020.

[4] Lindenberger, P., Sarlin, P. E., & Pollefeys, M. (2023). Lightglue: Local feature matching at light speed. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 17627-17638).

[5] Abouee, Amin, et al. "Weakly Supervised End2End Deep Visual Odometry", *In Proceedings of the 7th International Workshop on Visual Odometry and Computer Vision Applications Based on Location Clues at the CVPR 2024* (accepted, to be published)

https://doi.org/10.52202/075106-0017