

MIT Open Access Articles

*The Blackbird Dataset: A Large-Scale Dataset
for UAV Perception in Aggressive Flight*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Antonini, Amado, Guerra, Winter, Murali, Varun, Sayre-McCord, Thomas and Karaman, Sertac. 2020. "The Blackbird Dataset: A Large-Scale Dataset for UAV Perception in Aggressive Flight." PROCEEDINGS OF THE 2018 INTERNATIONAL SYMPOSIUM ON EXPERIMENTAL ROBOTICS, 11.

As Published: 10.1007/978-3-030-33950-0_12

Publisher: Springer Science and Business Media LLC

Persistent URL: <https://hdl.handle.net/1721.1/145377>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



The Blackbird Dataset: A large-scale dataset for UAV perception in aggressive flight

Amado Antonini¹, Winter Guerra¹, Varun Murali¹, Thomas Sayre-McCord¹,
and Sertac Karaman¹

Laboratory for Information and Decision Systems,
Massachusetts Institute of Technology *

Abstract. The Blackbird unmanned aerial vehicle (UAV) dataset is a large-scale, aggressive indoor flight dataset collected using a custom-built quadrotor platform for use in evaluation of agile perception. Inspired by the potential of future high-speed fully-autonomous drone racing, the Blackbird dataset contains over 10 hours of flight data from 168 flights over 17 flight trajectories and 5 environments at velocities up to 7.0 ms^{-1} . Each flight includes sensor data from 120 Hz stereo and downward-facing photorealistic virtual cameras, 100 Hz IMU, ~ 190 Hz motor speed sensors, and 360 Hz millimeter-accurate motion capture ground truth. Camera images for each flight were photorealistically rendered using FlightGoggles [1] across a variety of environments to facilitate easy experimentation of high performance perception algorithms. The dataset is available for download at <http://blackbird-dataset.mit.edu/>.

1 Introduction

Aggressive Unmanned Aerial Vehicle (UAV) flight using visual inertial simultaneous localization and mapping (VI-SLAM) has received increasing attention over recent years [1, 2]. With the availability of better hardware, aggressive indoor flight maneuvers that were previously only possible using motion capture systems are now becoming achievable using on-board visual inertial state estimation algorithms. In the near future, it is conceivable that complex high-speed tasks, such as fully-autonomous drone racing, will be possible in realtime. To aid in the development of these high-performance algorithms, we provide a large scale, high rate, and high accuracy dataset for the improvement and evaluation of VI-SLAM for agile indoor flight.

Related Work. As can be seen from Table 1, existing UAV datasets focus on either slow speed indoor flight [3] where high accuracy ground truth is achievable, or outdoor high speed flight [4–6] with lower quality ground truth using GPS systems. Burri et al [3] present the EuRoC MAV datasets, a collection of 11 trajectories with an average speed of 1 ms^{-1} and highly accurate

*Emails: {amadoa, winterg, mvarun, rtism, sertac}@mit.edu

Table 1: UAV Visual Inertial Datasets Comparison

	EuRoC MAV[3]	UPenn fast flight [4]	Zurich Urban MAV [5]	Ours
Environments	2	1	3	5^a
Sequences	11	4	1	186
Camera	20 Hz	40 Hz	20 Hz	120 Hz
IMU	200 Hz	200 Hz	10 Hz	100 Hz
Motor Encoders	n/a	n/a	n/a	~190 Hz
Max Distance	130.9 m	700 m	2 km	860.8 m
Top Speed	2.3 m s^{-1}	17.5 m s^{-1}	$3.9 \text{ m s}^{-1 b}$	7.0 m s^{-1}
mm Ground Truth	(100 Hz) ^c	n/a	n/a	360 Hz

^aAdditional environments may be rendered using FlightGoggles^bInstantaneous velocity from GPS^cHigh accuracy only available for half the sequences

ground truth for one half of the sequences. However, the scenes are not representative of typical indoor environments and are captured using comparatively lower rate cameras. Sun et al [4] present a fast outdoor flight dataset with the same trajectory at 4 different speeds and GPS ground truth. Although this does allow the evaluation of online algorithms in outdoor settings, it does not provide high quality ground truth or different environments. The Zurich Urban MAV Dataset is presented by Majdik, Till, and Scaramuzza. It contains 2 km of visual and inertial data recorded from a tethered UAV flying in an urban setting, but it lacks high-precision ground truth pose. Wang et al [6] present TorontoCity, a very large UAV dataset with data from multiple perspectives of the city of Toronto captured from different cameras and a LiDAR. TorontoCity focuses on tasks such as segmentation and classification of the environment. It, however does not contain inertial information and cannot be used in the context of visual inertial navigation.

2 Data Collection Setup

UAV Platform. Data was collected using a custom built quadrotor UAV designed for agile autonomous flight, which we call Blackbird (Figure 1). The UAV carries an Xsens MTi-3 IMU, custom made optical motor encoders for accurate motor speed measurements, a DJI Snail propulsion system, and a NVIDIA Jetson TX2. The body of the vehicle is constructed from 3D printed MarkForged Onyx continuous carbon fiber composite. Rubber dampeners are used to mechanically isolate vibrations from the propulsion system from flight sensors. The physical properties of the quadrotor as well as sensor statistics are shown in Table 2.

Experimental Setup. Flights were performed in an $11 \text{ m} \times 11 \text{ m} \times 5.5 \text{ m}$ motion capture room, with 24 OptiTrack Prime 17W cameras providing the 6D pose of the drone at 360 Hz. Each flight in the dataset is between 3-4 minutes long as the drone traces out a pre-defined periodic trajectory using a non-linear




Fig. 1: Coordinate frames in use for this dataset. Note that **Camera_D** and **Body Frame** are coincident, but are translated in the figure for visualization.

Table 2: Quadrotor characteristics

Property	Value	Description
Mass	0.915 kg	Mass with battery
I_{xx}	$4.9 \times 10^{-2} \text{ kg m}^{-2}$	X moment of inertia
I_{yy}	$4.9 \times 10^{-2} \text{ kg m}^{-2}$	Y moment of inertia
I_{zz}	$6.9 \times 10^{-2} \text{ kg m}^{-2}$	Z moment of inertia
Arm Length	0.13 m	Center to end of arm
f_x, f_y	665.108 mm	Cameras' focal length
FOV	60.0°	Cameras' vertical FOV
σ_{gyro}	$1.2 \times 10^{-4} \text{ rad s}^{-1} \sqrt{\text{Hz}}$	Gyroscope noise density
σ_b^{gyro}	$4.7 \times 10^{-6} \text{ rad s}^{-2} \sqrt{\text{Hz}}$	Gyroscope random walk bias
σ_{accel}	$2.0 \times 10^{-3} \text{ m s}^{-2} \sqrt{\text{Hz}}$	Accelerometer noise density
σ_b^{accel}	$4.4 \times 10^{-5} \text{ m s}^{-3} \sqrt{\text{Hz}}$	Accelerometer random walk bias
C_T	$2.27 \times 10^{-8} \text{ N/rpm}^2$	RPM to thrust coefficient
Image Size	1024 px \times 768 px	Image width and height

dynamic inversion controller [7]. The drone is controlled and data is recorded by a custom software framework [1] using the Lightweight Communications and Marshalling (LCM) protocol [8].

Visual Data Generation. Visual data was generated in post process using the FlightGoggles photo-realistic image generation system [1]. FlightGoggles uses the ground truth 6D pose of the drone from motion capture to generate images from the viewpoint of each camera on the drone in a virtual environment. The system allows for complete control over the visual appearance of the environment, the rate of camera images (up to the 360 Hz motion capture rate), the number of cameras, and each camera's location and intrinsic and extrin-

sic properties. The visual data generated by FlightGoggles has been previously validated for use in visual inertial state estimation in [1].

As part of the rendering process, a number of transforms are used to transform NED ground truth data from world frame into FlightGoggles’ environment frame. To ensure that all recorded flights in each trajectory takeoff from a common altitude and overlap in the XY plane, T_{mocap}^{norm} is introduced as a per-flight translational offset applied to the ground truth data to correct for offsets introduced during dataset collection. $T_{norm}^{FG_{env}}$ is a per-trajectory common transform that positions flights into the FlightGoggles environment in a collision-free manner. The full transform chain from ground-truth coordinates to render coordinates is shown in equation 1.

$$T^{FG_{env}} = T_{norm}^{FG_{env}} * T_{mocap}^{norm} * T^{mocap} \quad (1)$$

Where $T^{FG_{env}}$ is the render pose in FlightGoggles’ virtual environment.

Sensor Calibration and Temporal Synchronization. The Kalibr package [9] was used to find the noise characteristics of the IMU and the IMU-to-camera transform. A 3 second period at rest is included in every flight to allow for initialization of the time varying IMU bias. Force and torque coefficients of the drone were found experimentally through measurements in a wind tunnel to obtain the relationship between motor speeds and vehicle dynamics. Clock synchronization between motion capture data and on-board sensors was performed using a combination of clock estimation over gigabit ethernet and chrony [10] over the wireless network, with an upper bounded offset of $\pm 5\text{ms}$.

3 Dataset Format

Each flight within the dataset contains timestamped values for the following: ground truth 6D pose of the UAV at 360 Hz, IMU measurements at 100 Hz, RPM measurements for each motor at ~ 190 Hz, and three camera streams (forward facing stereo pair and downward facing) at 120 Hz. The data is provided as grayscale images, LCM logs, and ROS bags for easy use in typical pipelines. Scripts and binaries necessary to re-render images using FlightGoggles at other rates (up to 360 Hz) or camera parameters and configurations are available at <http://blackbird-dataset.mit.edu/>.

In addition to the raw data streams, the full calibration information of the UAV system is included in the dataset i.e, IMU noise characteristics, IMU-camera transform, camera intrinsic and extrinsic parameters (as currently rendered), and torque and thrust curves. The dataset’s file structure is specified in figure 2.

4 Data Validation


Validation of ground truth data and inertial measurements in both quality and temporal synchronization was performed by comparing raw inertial measurements with derivatives of the ground truth pose using a Savitzky-Golay filter




Fig. 2: Dataset file hierarchy.

[11]. Figure 3 shows a comparison of the IMU angular rate and acceleration with respect to ground truth. The accuracy of the drone’s motor speed sensors were verified through the use of an external tachometer.

Trajectory Tracking. A feature of the provided dataset is the ability to repeatedly run perception algorithms on a nominal trajectory pattern while flying at different speeds with new inertial, dynamical, and visual data. A comparison of the ground truth pose of the same trajectory (Sphinx, see Fig. 6) flown four times at speeds between 1 m s^{-1} and 4 m s^{-1} is shown in Figure 4, with sufficient tracking accuracy for a user to isolate the speed of flight from other parameters that may affect VI-SLAM algorithms.

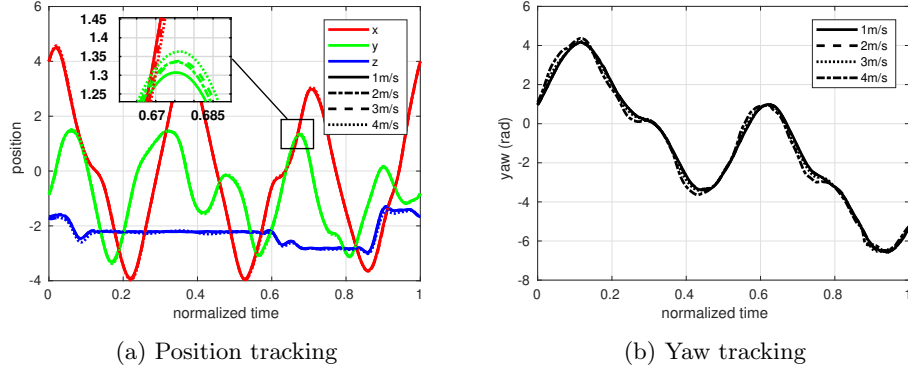


(a) IMU accelerometer vs ground truth.



(b) IMU gyroscope vs ground truth.

Fig. 3: Derivative of position and rotational ground truth data compared with accelerometer and gyroscope data for a flight at 4 m s^{-1} .

Fig. 4: Tracking precision while flying the same trajectory at speeds of 1 to 4 m s^{-1}

Collision Checking. To verify that the trajectories being rendered do not collide with any obstacles in the virtual environment, bullet3 [12] based simulation was used to verify that the trajectories were collision free prior to rendering the camera streams.

Table 3: Dataset Flights

Trajectory	Constant Yaw								Forward-Facing							
Top speed (ms ⁻¹)	0.5	1.0	2.0	3.0	4.0	5.0	6.0	7.0	0.5	1.0	2.0	3.0	4.0	5.0	6.0	7.0
3D Figure 8	✓	✓	✓	✓	✓	✓	-	-	-	-	-	-	-	-	-	-
Amperсанд	-	M	M	H	-	-	-	-	-	H	H	-	-	-	-	-
Bent Dice	-	E	E	M	M	-	-	-	E	E	E	E	-	-	-	-
Clover	-	H	H	H	H	H	H	-	H	H	H	H	H	H	-	-
Dice	-	-	E	E	M	-	-	-	-	E	E	E	-	-	-	-
Flat Figure 8	✓	✓	✓	✓	-	✓	-	-	-	-	-	-	-	-	-	-
Half-Moon	-	E	E	E	M	-	-	-	-	M	M	M	M	-	-	-
Mouse	-	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M
Oval	-	-	M	M	H	-	-	-	-	M	M	H	H	-	-	-
Patrick	-	E	E	E	E	M	-	-	E	E	E	E	-	-	-	-
Picasso	M	M	M	M	M	M	M	-	M	M	-	H	H	H	-	-
Sid	-	E	E	E	E	E	E	E	M	M	M	M	M	M	-	-
Sphinx	-	H	H	H	H	-	-	-	-	M	M	M	M	-	-	-
Star	-	M	M	M	H	H	-	-	M	M	M	M	M	H	-	-
Thrice	-	E	E	E	E	E	M	M	E	E	E	E	E	E	M	-
Tilted Thrice	-	E	E	E	E	E	M	M	E	E	E	E	E	E	E	-
Winter	-	M	M	M	M	M	-	-	M	-	H	H	H	-	-	-

Click flight for grayscale video preview of flight in all rendered environments.

5 Datasets

The trajectories were designed to allow for independent variation in the following flight characteristics: speed, yaw, trajectory complexity, and period. The

included trajectories are classified by difficulty ([E]asy, [M]edium, [H]ard) according to mean features tracked using the visual inertial state estimation pipeline outlined in [1] and are shown in Table 3. These trajectories range in complexity from an oval with constant yaw and altitude to trajectories with varying speed, altitude, and yaw as they weave through visual obstacles (e.g. the Sphinx) as shown in Fig. 6. For flights that are rendered in multiple environments (see Fig. 5), some environments are harder for state estimation than others.



Fig. 5: Five rendering environments for visual data: (a) Butterfly Apartment, (b) Hazelwood Loft, (c) Museum Pillars, (d) Museum Sphinx, and (e) NYC Subway

To generate smoothly trackable trajectories, minimum snap optimization was performed over a set of requested waypoint positions using the non-linear optimization technique from [13], with boundary conditions to make the trajectory periodic. Table 3 shows all the sequences included. There are 163 unique flights of approximately three minutes each for a total of over 10 h and 60 km of ground truth, inertial, and dynamical sensor data, as well as rendered imagery in multiple environments. Trajectories were designed for specific flight environments (e.g. Sphinx for Museum Sphinx) and are therefore particularly well suited to those environments, however, where it does not result in virtual collisions with objects, the same trajectory can be re-rendered in multiple environments. Flights marked in Table 3 by ✓ do not have visual data associated with them and are intended for calibration use.

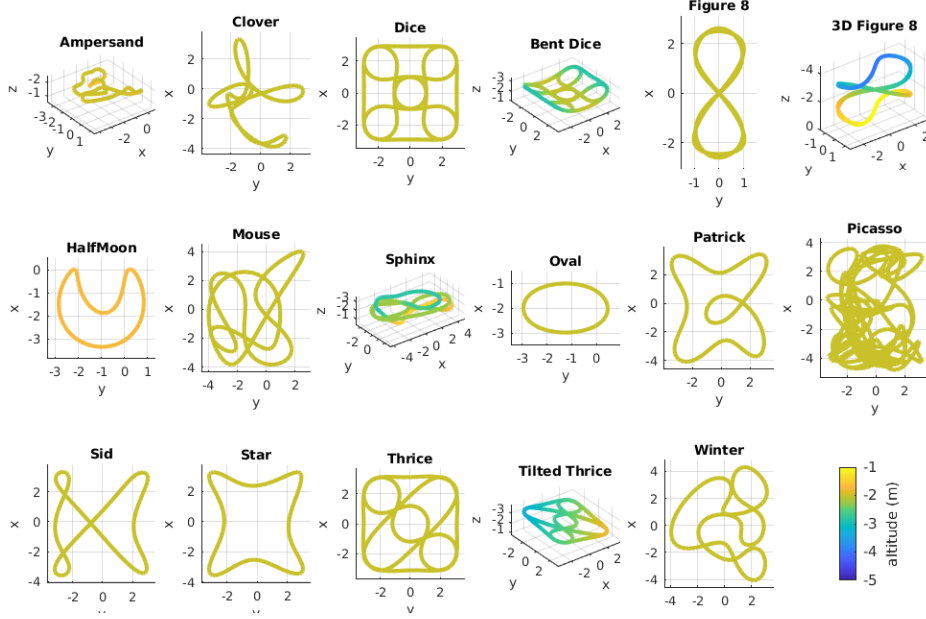


Fig. 6: Top down diagrams of trajectory paths included in this dataset.

6 Known Issues

In this dataset, we synchronized motion capture ground truth data with onboard IMU measurements using camera exposure timestamps provided by OptiTrack, IMU measurement timestamps provided on arrival by our UAV’s TX2, and clock sync and drift correction provided by Chrony [10]. However, due to the complexity and stochastic nature of the systems involved, we are only able to guarantee IMU and ground truth temporal alignment to within $\pm 5\text{ms}$ across all flights in this dataset. This upperbound was verified in post process by cross correlation of IMU and ground truth measurements for each flight.

Over the course of the various data recording sessions required to create this dataset, many recalibrations of the motion capture groundtruth setup were required due to thermal expansion and contraction of the motion capture support beams. Each recalibration’s overall tracking error and groundplane alignment was verified using post process analysis, but are not guaranteed to be exact. Most flights were flown with $\leq 1.5\text{mm}$ of mean tracking error and $\leq 1^\circ$ of ground plane alignment error.

7 Conclusion

The Blackbird dataset allows for the evaluation of the robustness and performance of VI-SLAM algorithms under varying conditions and degrees of agility. It has been used in [14] to evaluate dynamics factors in factor graph VIO. As

an additional example, we show the results of running a tracker-based visual inertial odometry method [1] on the Ampersand trajectory for increasing speed of execution. The performance is summarized in Table 4, where drift rate (the percentage error per meter traveled) serves as a metric of the overall state estimation accuracy. The mean features tracked, the average number of features tracked between every pair of consecutive frames, serves as a metric of success for the visual front-end. The structured nature of the dataset allows for clear evaluations of the strengths and failures of particular algorithms (e.g. tracking vs. matching visual front-ends) depending on the type of trajectory and visual environment being flown in.

Table 4: Performance of Visual Inertial Odometry [1] for the same nominal trajectory with varying speed

Speed	Drift Rate	Mean Features Tracked
1.12 m s^{-1}	0.55	118.98
1.96 m s^{-1}	0.95	115.74
3.08 m s^{-1}	1.94	106.83

Due to the large number of aggressive and varied flight patterns available, the Blackbird dataset is well suited to fill a void in the current landscape of available robotics datasets. The inertial, dynamical, and visual sensors provide a complete sensor package for many UAV VI-SLAM algorithms. The photorealistic visual simulation system, FlightGoggles, allows a user to expand beyond the three camera streams provided and re-render images to match the characteristics of their current system, or to evaluate design choices on camera placement, type, and framerate. Variations in visual parameters may all be performed while maintaining the true dynamics and inertial measurements of the UAV, and without sacrificing the high resolution ground truth accuracy provided by an expensive motion capture system.

The ability to repeat nominal trajectories with varying conditions allows for methodical evaluation of perception algorithm performance during high speed flight and enables progressively building capabilities towards more challenging scenarios.

Bibliography

- [1] Sayre-McCord, T., Guerra, W., Antonini, A., Arneberg, J., Brown, A., Cav-alheiro, G., Fang, Y., Gorodetsky, A., McCoy, D., Quilter, S., Riether, F., Tal, E., Terzioglu, Y., Carlone, L., Karaman, S.: Visual-inertial navigation algorithm development using photorealistic camera simulation in the loop. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). (2018)

- [2] Falanga, D., Mueggler, E., Faessler, M., Scaramuzza, D.: Aggressive quadrotor flight through narrow gaps with onboard sensing and computing using active vision. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), IEEE (2017) 5774–5781
- [3] Burri, M., Nikolic, J., Gohl, P., Schneider, T., Rehder, J., Omari, S., Achtelik, M.W., Siegwart, R.: The euroc micro aerial vehicle datasets. *The International Journal of Robotics Research* **35**(10) (2016) 1157–1163
- [4] Sun, K., Mohta, K., Pfrommer, B., Watterson, M., Liu, S., Mulgaonkar, Y., Taylor, C.J., Kumar, V.: Robust stereo visual inertial odometry for fast autonomous flight. *IEEE Robotics and Automation Letters* **3**(2) (2018) 965–972
- [5] Majdik, A.L., Till, C., Scaramuzza, D.: The zurich urban micro aerial vehicle dataset. *The International Journal of Robotics Research* **36**(3) (2017) 269–273
- [6] Wang, S., Bai, M., Mattyus, G., Chu, H., Luo, W., Yang, B., Liang, J., Cheverie, J., Fidler, S., Urtasun, R.: Torontocity: Seeing the world with a million eyes. In: 2017 IEEE International Conference on Computer Vision (ICCV), IEEE (2017) 3028–3036
- [7] Tal, E., Karaman, S.: Accurate tracking of aggressive quadrotor trajectories using incremental nonlinear dynamic inversion and differential flatness. In: 2018. Proceedings. 57th IEEE Conference on Decision and Control, IEEE (2018)
- [8] Huang, A.S., Olson, E., Moore, D.C.: LCM: Lightweight communications and marshalling. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). (2010) 4057–4062
- [9] Furgale, P., Rehder, J., Siegwart, R.: Unified temporal and spatial calibration for multi-sensor systems. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). (Nov 2013) 1280–1286
- [10] Miroslav Lichvar: Chrony. <https://chrony.tuxfamily.org/>
- [11] Savitzky, A., Golay, M.J.: Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry* **36**(8) (1964) 1627–1639
- [12] Coumans, E.: Bullet physics simulation. In: ACM SIGGRAPH 2015 Courses, ACM (2015) 7
- [13] Burri, M., Oleynikova, H., , Achtelik, M.W., Siegwart, R.: Real-time visual-inertial mapping, re-localization and planning onboard mavs in unknown environments. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). (Sept 2015)
- [14] Antonini, A., Leonard, J., Karaman, S.: Pre-Integrated Dynamics Factors and a Dynamical Agile Visual-Inertial Dataset for UAV Perception. Master’s thesis, Massachusetts Institute of Technology (2018)