



# ADVIO: An Authentic Dataset for Visual-Inertial Odometry

Santiago Cortés<sup>1</sup>, Arno Solin<sup>1</sup>, Esa Rahtu<sup>2</sup>, and Juho Kannala<sup>1</sup>

<sup>1</sup> Department of Computer Science, Aalto University, Espoo, Finland  
{santiago.cortesreina, arno.solin, juho.kannala}@aalto.fi

<sup>2</sup> Tampere University of Technology, Tampere, Finland  
esa.rahtu@tut.fi

**Abstract.** The lack of realistic and open benchmarking datasets for pedestrian visual-inertial odometry has made it hard to pinpoint differences in published methods. Existing datasets either lack a full six degree-of-freedom ground-truth or are limited to small spaces with optical tracking systems. We take advantage of advances in pure inertial navigation, and develop a set of versatile and challenging real-world computer vision benchmark sets for visual-inertial odometry. For this purpose, we have built a test rig equipped with an iPhone, a Google Pixel Android phone, and a Google Tango device. We provide a wide range of raw sensor data that is accessible on almost any modern-day smartphone together with a high-quality ground-truth track. We also compare resulting visual-inertial tracks from Google Tango, ARCore, and Apple ARKit with two recent methods published in academic forums. The data sets cover both indoor and outdoor cases, with stairs, escalators, elevators, office environments, a shopping mall, and metro station.

**Keywords:** Visual-inertial odometry · Navigation · Benchmarking


## 1 Introduction

Various systems and approaches have recently emerged for tracking the motion of hand-held or wearable mobile devices based on video cameras and inertial measurement units (IMUs). There exist both open published methods (*e.g.* [2, 12, 14, 16, 21]) and closed proprietary systems. Recent examples of the latter are ARCore by Google and ARKit by Apple which run on the respective manufacturers' flagship smartphone models. Other examples of mobile devices with built-in visual-inertial odometry are the Google Tango tablet device and Microsoft HoloLens augmented reality glasses. The main motivation for developing odometry methods for smart mobile devices is to enable augmented reality

---

**Access data and documentation at:** <https://github.com/AaltoVision/ADVIO>.

**Electronic supplementary material** The online version of this chapter ([https://doi.org/10.1007/978-3-030-01249-6\\_26](https://doi.org/10.1007/978-3-030-01249-6_26)) contains supplementary material, which is available to authorized users.

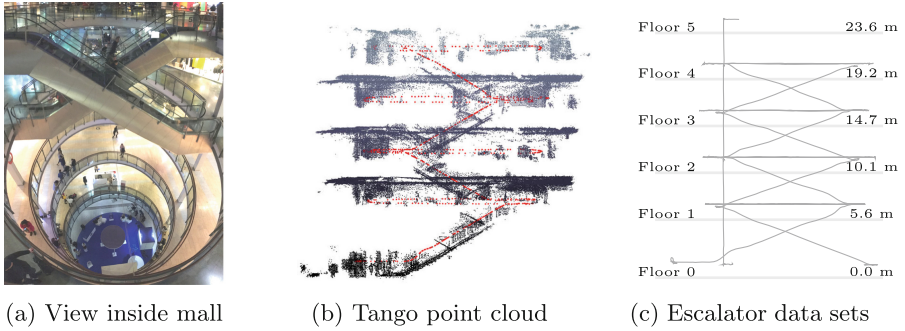


**Fig. 1.** The custom-built capture rig with a Google Pixel smartphone on the left, a Google Tango device in the middle, and an Apple iPhone 6s on the right.

applications which require precise real-time tracking of ego-motion. Such applications could have significant value in many areas, like architecture and design, games and entertainment, telepresence, and education and training.

Despite the notable scientific and commercial interest towards visual-inertial odometry, the progress of the field is constrained by the lack of public datasets and benchmarks which would allow fair comparison of proposed solutions and facilitate further developments to push the current boundaries of the state-of-the-art systems. For example, since the performance of each system depends on both the algorithms and sensors used, it is hard to compare methodological advances and algorithmic contributions fairly as the contributing factors from hardware and software may be mixed. In addition, as many existing datasets are either captured in small spaces or utilise significantly better sensor hardware than feasible for low-cost consumer devices, it is difficult to evaluate how the current solutions would scale to medium or long-range odometry, or large-scale simultaneous localization and mapping (SLAM), on smartphones.

Further, the availability of realistic sensor data, captured with smartphone sensors, together with sufficiently accurate ground-truth would be beneficial in order to speed up progress in academic research and also lower the threshold for new researchers entering the field. The importance of public datasets and benchmarks as a driving force for rapid progress has been clearly demonstrated in many computer vision problems, like image classification [9,19], object detection [13], stereo reconstruction [10] and semantic segmentation [6,13], to name a few. However, regarding visual-inertial odometry, there are no publicly available datasets or benchmarks that would allow evaluating recent methods in a typical smartphone context. Moreover, since the open-source software culture is not as common in this research area as, for example, it is in image classification and object detection, the research environment is not optimal for facilitating rapid progress. Further, due to the aforementioned reasons, there is a danger that the field could become accessible only for big research groups funded by large corporations, and that would slow down progress and decay open academic research.



**Fig. 2.** Multi-floor environments such as (a) were considered. The point cloud (b) and escalator/elevator paths captured in the mall. The Tango track (red) in (b) has similar shape as the ground-truth in (c). Periodic locomotion can be seen in (c) if zoomed in. (Color figure online)

In this work, we present a dataset that aims to facilitate the development of visual-inertial odometry and SLAM methods for smartphones and other mobile devices with low-cost sensors (*i.e.* rolling-shutter cameras and MEMS based inertial sensors). Our sensor data is collected using a standard iPhone 6s device and contains the ground-truth pose trajectory and the raw synchronized data streams from the following sensors: RGB video camera, accelerometer, gyroscope, magnetometer, platform-provided geographic coordinates, and barometer. In total, the collected sequences contain about 4.5 km of unconstrained hand-held movement in various environments both indoors and outdoors. One example sequence is illustrated in Fig. 2. The data sets are collected in public spaces, conforming the local legislation regarding filming and publishing. The ground-truth is computed by combining a recent pure inertial navigation system (INS) [24] with frequent manually determined position fixes based on a precise floor plan. The quality of our ground-truth is verified and its accuracy estimated.

Besides the benchmark dataset, we present a comparison of visual-inertial odometry methods, including three recent proprietary platforms: ARCore on a Google Pixel device, Apple ARKit on the iPhone, and Tango odometry on a Google Tango tablet device, and two recently published methods, namely ROVIO [1,2] and PIVO [25]. The data for the comparison was collected with a capture rig with the three devices and is illustrated in Fig. 1. Custom applications for data capture were implemented for each device.

The main contributions of our work are summarized in the following:

- A public dataset of iPhone sensor data with 6 degree-of-freedom pose ground-truth for benchmarking monocular visual-inertial odometry in real-life use cases involving motion in varying environments, and also including stairs, elevators and escalators.
- Comparing state-of-the-art visual-inertial odometry platforms and methods.
- A method for collecting ground-truth for smartphone odometry in realistic use cases by combining pure inertial navigation with manual position fixes.

**Table 1.** An overview of related datasets.

	Rawseeds [5]	KITTI [10]	NCLT [4]	EuRoC [3]	PennCOSYVIO [18]	Proposed
Year	2006	2012	2015	2016	2017	(this paper)
Carrier	Wheeled robot	Car	Segway	MAV	Hand-held device	Hand-held device
Environment	Indoors/Outdoors	Outdoors	Indoors/Outdoors	Indoors	Indoors/Outdoors	Indoors/Outdoors
Scene setup	Campus-scale	City-scale	Campus-scale	2 Rooms	150 m path on campus (walked 4 times)	Multiple levels in 3 buildings + outdoor scenes
Hardware setup	Custom	Custom	Custom	Custom	Custom	Standard smartphone
Distance (total)	~10 km	~39 km	~147 km	~800 m	~600 m	~4.5 km
Long-range use-case	✓	✓	✓	—	✓	✓
3D point cloud	—	✓	✓	✓	—	✓
Ground-truth	GPS/Visual tags	GPS/IMU	GPS/IMU/Laser	MoCap/Laser	Visual tags	IMU + Position fixes
Accuracy	~m	~dm	~dm	~mm	~dm	dm-m

## 2 Related Work


Despite visual-inertial odometry (VIO) being one of the most promising approaches for real-time tracking of hand-held and wearable devices, there is a lack of good public datasets for benchmarking different methods. A relevant benchmark should include both video and inertial sensor recordings with synchronized time stamps preferably captured with consumer-grade smartphone sensors. In addition, the dataset should be authentic and illustrate realistic use cases. That is, it should contain challenging environments with scarce visual features, both indoors and outdoors, and varying motions, also including rapid rotations without translation as they are problematic for monocular visual-only odometry. Our work is the first one addressing this need.

Regarding pure visual odometry or SLAM, there are several datasets and benchmarks available [6, 8, 23, 26] but they lack the inertial sensor data. Further, many of these datasets are limited because they (a) are recorded using ground vehicles and hence do not have rapid rotations [6, 23], (b) do not contain low-textured indoor scenes [6, 23], (c) are captured with custom hardware (e.g. fisheye lens or global shutter camera) [8], (d) lack full 6-degree of freedom ground-truth [8], or (e) are constrained to small environments and hence are ideal for SLAM systems but not suitable for benchmarking odometry for medium and long-range navigation [26].

Nevertheless, besides pure vision datasets, there are some public datasets with inertial sensor data included, for example, [3–5, 10, 18]. Most of these datasets are recorded with sensors rigidly attached to a wheeled ground vehicle. For example, the widely used KITTI dataset [10] contains LIDAR scans and videos from multiple cameras recorded from a moving car. The ground-truth is obtained using a very accurate GPS/IMU localization unit with RTK correction signals. However, the IMU data is captured only with a frequency of 10 Hz, which would not be sufficient for tracking rapidly moving hand-held devices. Further, even if high-frequency IMU data would be available, also KITTI has the constraints (a), (b), and (c) mentioned above and this limits its usefulness for smartphone odometry.

Another analogue to KITTI is that we also use pure inertial navigation with external location fixes for determining the ground-truth. In our case, the GPS fixes are replaced with manual location fixes since GPS is not available or accurate indoors. Further, in contrast to KITTI, by utilizing recent advances in inertial navigation [24] we are able to use the inertial sensors of the iPhone for the ground-truth calculation and are therefore not dependent on a high-grade IMU, which would be difficult to attach to the hand-held rig. In our case the manual location fixes are determined from a reference video (Fig. 3a), which views the recorder, by visually identifying landmarks that can be accurately localized from precise building floor plans or aerial images. The benefit of not using optical methods for establishing the ground-truth is that we can easily record long sequences and the camera of the recording device can be temporarily occluded. This makes our benchmark suitable also for evaluating occlusion robustness of VIO methods [25]. Like KITTI, the Rawseeds [5] and NCLT [4] datasets are recorded with a wheeled ground vehicle. Both of them use custom sensors (*e.g.* omnidirectional camera or industrial-grade IMU). These datasets are for evaluating odometry and self-localization of slowly moving vehicles and not suitable for benchmarking VIO methods for hand-held devices and augmented reality.

The datasets that are most related to ours are EuRoC [3] and PennCOSYVIO [18]. EuRoC provides visual and inertial data captured with a global shutter stereo camera and a tactical-grade IMU onboard a micro aerial vehicle (MAV) [17]. The sequences are recorded in two different rooms that are equipped with motion capture system or laser tracker for obtaining accurate ground-truth motion. In PennCOSYVIO, the data acquisition is performed using a hand-held rig containing two Google Tango tablets, three GoPro Hero 4 cameras, and a similar visual-inertial sensor unit as used in EuRoC. The data is collected by walking a 150 m path several times at UPenn campus, and the ground-truth is obtained via optical markers. Due to the need of optic localization for determining ground-truth, both EuRoC and PennCOSYVIO contain data only from a few environments that are all relatively small-scale. Moreover, both datasets use the same high-quality custom sensor with wide field-of-view stereo cameras [17]. In contrast, our dataset contains around 4.5 km of sequences recorded with regular smartphone sensors in multiple floors in several different buildings and different outdoor environments. In addition, our dataset contains motion in stairs, elevators and escalators, as illustrated in Fig. 2, and also temporary occlusions and lack of visual features. We are not aware of any similar public dataset. The properties of different datasets are summarized in Table 1. The enabling factor for our flexible data collection procedure is to utilize recent advances in pure inertial navigation together with manual location fixes [24]. In fact, the methodology for determining the ground-truth is one of the contributions of our work. In addition, as a third contribution, we present a comparison of recent VIO methods and proprietary state-of-the-art platforms based on our challenging dataset.



**Fig. 3.** Example of simultaneously captured frames from three synchronized cameras. The external reference camera (a) is used for manual position fixes for determining the ground-truth trajectory in a separate post-processing stage.

### 3 Materials

The data was recorded with the three devices (iPhone 6s, Pixel, Tango) rigidly attached to an aluminium rig (Fig. 1). In addition, we captured the collection process with an external video camera that was viewing the recorder (Fig. 3). The manual position fixes with respect to a 2D map (*i.e.* a structural floor plan image or an aerial image/map) were determined afterwards from the view of the external camera. Since the device was hand-held, in most fix locations the height was given as a constant distance above the floor level (with a reasonable uncertainty estimate), so that the optimization could fit a trajectory that optimally balances the information from fix positions and IMU signals (details in Sect. 4).

The data streams from all the four devices are synchronized using network provided time. That is, the device clock is synchronized over a network time protocol (NTP) request at the beginning of a capture session. All devices were connected to 4G network during recording. Further, in order to enable analysis of the data in the same coordinate frame, we calibrated the internal and external parameters of all cameras by capturing multiple views of a checkerboard. This was performed before each session to account for small movements during transport and storage. The recorded data streams are listed in Table 2.

#### 3.1 Raw iPhone Sensor Capture

An iOS data collection app was developed in Swift 4. It saves inertial and visual data synchronized to the Apple ARKit pose estimation. All individual data points are time stamped internally and then synchronized to global time. The global time is fetched using the Kronos Swift NTP client<sup>1</sup>. The data was cap-

<sup>1</sup> <https://github.com/lyft/Kronos>.

**Table 2.** Data captured by the devices.

Device	Data	Format	Units	Capture rate
Ground-truth	Pose	Position/orientation	Metric position	100 Hz
iPhone	ARKit pose	Position/orientation	Metric position	60 Hz
	Video	RGB video	Resolution 1280×720	60 Hz
	GNSS	Latitude/Longitude	World coordinates (incl. meta)	~1 Hz
	Barometer	Pressure	kPa	~10 Hz
	Gyroscope	Angular rate	rad/s	100 Hz
	Accelerometer	Specific force	g	100 Hz
	Magnetometer	Magnetic field	$\mu$ T	100 Hz
Pixel	ARCore pose	Position/orientation	Metric position	30 Hz
Tango	Raw pose	Position/orientation	Metric position	60 Hz
	Area learning	Position/orientation	Metric position	60 Hz
	Fisheye video	Grayscale video	Resolution: 640×480	60 Hz
	Point cloud	Array of 3D points	Point colour	~5 Hz

tured using an iPhone 6s running iOS 11.0.3. The same software and an identical iPhone was used for collecting the reference video. This model was chosen, because the iPhone 6s (published 2015) is hardware-wise closer to an average smartphone than most recent flagship iPhones and also matches well with the Google Pixel hardware.

During the capture the camera is controlled by the ARKit service. It is performing the usual auto exposure and white balance but the focal length is kept fixed (the camera matrix returned by ARKit is stored during capture). The resolution is also controlled by ARKit and it is 1280×720. The frames are packed into an H.264/MPEG-4 video file. The GNSS/network location data is collected through the CoreLocation API. Locations are requested with the desired accuracy of ‘kCLLocationAccuracyBest’. The location service provides latitude and longitude, horizontal accuracy, altitude, vertical accuracy, and speed. The accelerometer, gyroscope, magnetometer, and barometer data are collected through the CoreMotion API and recorded at the maximum rate. The approximate capture rates of the multiple data streams are shown in Table 2. The magnetometer values are uncalibrated. The barometer samples contain both the barometric pressure and associated relative altitude readings.

### 3.2 Apple ARKit Data

The same application that captures the raw data is running the ARKit framework. It provides a pose estimate associated with every video frame. The pose is saved as a translation vector and a rotation expressed in Euler angles. Each pose is relative to a global coordinate frame created by the phone.

### 3.3 Google ARCore Data

We wrote an app based on Google’s ARCore example<sup>2</sup> for capturing the ARCore tracking result. Like ARKit, the pose data contains a translation to the first

<sup>2</sup> <https://github.com/google-ar/arcore-android-sdk>.

frame of the capture and a rotation to a global coordinate frame. Unlike ARKit, the orientation is stored as a unit quaternion. Note that the capture rate is slower than with ARKit. We do not save the video frames nor the sensor data on the Pixel. The capture was done on a Google Pixel device running Android 8.0.0 Oreo and using the Tango Core AR developer preview (Tango core version 1.57:2017.08.28-release-ar-sdk-preview-release-0-g0ce07954:250018377:stable).

### 3.4 Google Tango Data

A data collection app developed and published by [11], based on the Paraview project<sup>3</sup>, was modified in order to collect the relevant data. The capture includes the position of the device relative to the first frame, the orientation in global coordinates, the fisheye grayscale image, and the point cloud created by the depth sensor. The Tango service was run on a Project Tango tablet running Android 4.4.2 and using Tango Core Argentine (Tango Core version 1.47:2016.11-22-argentine.tango-release-0-gce1d28c8:190012533:stable). The Tango service produces two sets of poses, referred to as *raw odometry* and *area learning*<sup>4</sup>. The raw odometry is built frame to frame without long term memory whereas the area learning uses ongoing map building to close loops and reduce drift. Both tracks are captured and saved.

### 3.5 Reference Video and Locations

One important contribution of this paper is the flexible data collection framework that enables us to capture realistic use cases in large environments. In such conditions, it is not feasible to use visual markers, motion capture, or laser scanners for ground-truth. Instead, our work takes advantage of pure inertial navigation together with manual location fixes as described in Sect. 4.1.

In order to obtain the location fixes, we record an additional reference video, which is captured by an assisting person who walks within a short distance from the actual collector. Figure 3a illustrates an example frame of such video. The reference video allows us to determine the location of the data collection device with respect to the environment and to obtain the manual location fixes (subject to measurement noise) for the pure inertial navigation approach [24].

In practice, the location fixes are produced as a post-processing step using a location marking tool developed for this paper. In this tool, one can browse the videos, and mark manual location fixes on the corresponding floor plan image. The location fixes are inserted on occasions where it is easy to determine the device position with respect to the floor plan image (*e.g.* in the beginning and the end of escalators, entering and exiting elevator, passing through a door, or walking past a building corner). In all our recordings it was relatively easy to find enough such instances needed to build an accurate ground-truth. Note that it is enough to determine the device location manually, not orientation.

<sup>3</sup> <https://github.com/Kitware/ParaViewTangoRecorder>.

<sup>4</sup> <https://developers.google.com/tango/overview/area-learning>.



The initial location fixes have to be further transformed from pixel coordinates of floor plan images into metric world coordinates. This is done by first converting pixels to meters by using manually measured reference distances (*e.g.* distance between pillars). Then the floor plan images are registered with respect to each other using manually determined landmark points (*e.g.* pillars or stairs) and floor height measurements.

## 4 Methods

### 4.1 Ground-Truth

The ground-truth is an implementation of the purely inertial odometry algorithm presented in [24], with the addition of manual fixation points recorded using the external reference video (see Sec. 3.5). The IMU data used in the inertial navigation system for the ground-truth originated from the iPhone, and is the same data that is shared as part of the dataset. Furthermore, additional calibration data was acquired for the iPhone IMUs accounting for additive gyroscope bias, additive accelerometer bias, and multiplicative accelerometer scale bias.

The inference of the iPhone pose track (position and orientation) was implemented as described in [24] with the addition of fusing the state estimation with both the additional calibration data and the manual fix points. The pose track corresponds to the INS estimates conditional to the fix points and external calibrations,

$$p(\mathbf{p}(t_k), \mathbf{q}(t_k) \mid \text{IMU, calibrations, } \{(t_i, \mathbf{p}_i)\}_{i=1}^N), \quad (1)$$

where  $\mathbf{p}(t_k) \in \mathbb{R}^3$  is the phone position and  $\mathbf{q}(t_k)$  is the orientation unit quaternion at time instant  $t_k$ . The set of fixpoints consists of time–position pairs  $(t_i, \mathbf{p}_i)$ , where the manual fixpoint  $\mathbf{p}_i \in \mathbb{R}^3$  assigned to a time instant  $t_i$ . The ‘IMU’ refers to all accelerometer and gyroscope data over the entire track.

Accounting for uncertainty and inaccuracy in the fixation point locations is taken into account by not enforcing the phone track to match the points, but including a Gaussian measurement noise term with a standard deviation of 25 cm in the position fixes (in all directions). This allows the estimate track to disagree with the fix. Position fixes are given either as 3D locations or 2D points with unknown altitude while moving between floors.

The inference problem was finally solved with an extended Kalman filter (forward pass) and extended Rauch–Tung–Striebel smoother (backward pass, see [24] for technical details). As real-time computation is not required here, we could have also used batch optimization but that would not have caused noticeable change in the results. Calculated tracks were inspected manually frame by frame and the pose track was refined by additional fixation points until the track matched the movement seen in all three cameras and the floor plan images. Figure 2c shows examples of the estimated ground-truth track. The vertical line is an elevator ride (stopping in each floor). Walking-induced periodic movement can be seen if zoomed in. The obtained accuracy can be checked also from the example video in the supplementary material.

## 4.2 Evaluation Metrics

For odometry results captured on the fly while collecting the data, we propose the following evaluation metrics. All data was first temporally aligned to the same global clock (acquired by NTP requests while capturing the data), which seemed to give temporal alignments accurate to about 1–2 s. The temporal alignment was further improved by determining a constant time offset by minimizing the median error between the device yaw and roll tracks. This alignment accounts for both temporal registration errors between devices and internal delays in the odometry methods.

After the temporal alignment the tracks provided by the three devices are chopped to the same lengths covering the same time-span as there may be few seconds differences in the starting and stopping times of the recordings with different devices. The vertical direction is already aligned to gravity. To account for the relative poses between the devices, method estimates, and ground-truth, we estimate a planar rigid transform (2D rotation and translation) between estimate tracks and ground-truth based on the first 60 s of estimates in each method (using the entire path would not have had a clear effect on the results, though). The reason for not using the calibrated relative poses is that especially ARCore (and occasionally ARKit) showed wild jumps at the beginning of the tracks, which would have had considerable effects and ruined those datasets for the method.

The aligned tracks all start from origin, and we measure the absolute error to the ground-truth for every output given by each method. The empirical cumulative distribution function for the absolute position error is defined as

$$\hat{F}_n(d) = \frac{\text{number of position errors } \leq d}{n} = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{e_i \leq d}, \quad (2)$$

where  $\mathbf{1}_E$  is an indicator function for the event  $E$ ,  $\mathbf{e} \in \mathbb{R}^n$  is a vector of absolute position errors compared to ground-truth, and  $n$  is the number of positions. The function tells the proportion of position estimates being less than  $d$  meters from ground-truth.

## 5 Data and Results

The dataset contains 23 separate recordings captured in six different locations. The total length of all sequences is 4.47 km and the total duration is 1 h 8 min. There are 19 indoor and 4 outdoor sequences. In the indoor sequences there is a manual fix point on average every 3.7 m (or 3.8 s), and outdoors every 14.7 m (or 10 s). The ground-truth 3D trajectories for all the sequences are illustrated in the supplementary material, where also additional details are given. In addition, one of the recordings and its ground-truth are illustrated in the supplementary video. The main characteristics of the dataset sequences and environments are briefly described below.

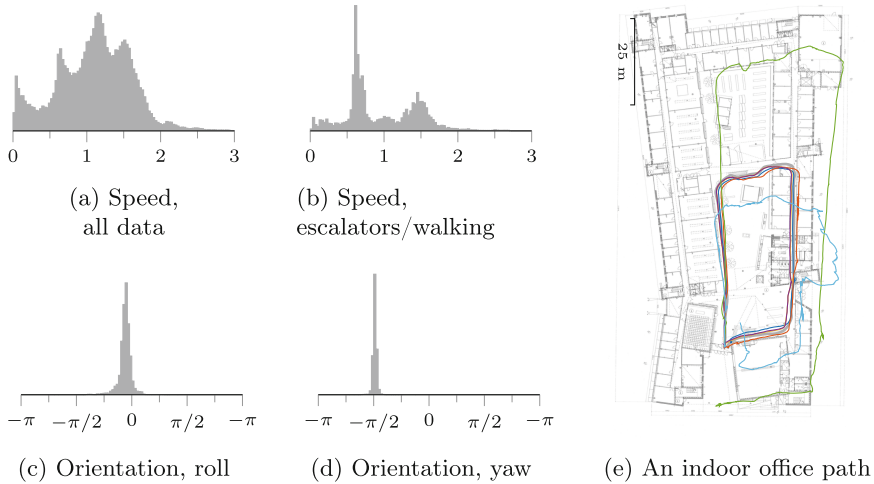


**Fig. 4.** Example frames from datasets. There are 7 sequences from two separate office buildings, 12 sequences from urban indoor scenes (malls and metro station), two from urban outdoor scenes, and two from suburban (campus) outdoor scenes.

Our dataset is primarily designed for benchmarking medium and long-range odometry. The most obvious use case is indoor navigation in large spaces, but we have also included outdoor paths for completeness. The indoor sequences were acquired in a 7-storey shopping mall ( $\sim 135,000 \text{ m}^2$ ), in a metro station, and in two different office buildings. The shopping mall and station are in the same building complex. The metro and bus station is located in the bottom floors, and there are plenty of moving people and occasional large vehicles visible in the collected videos, which makes pure visual odometry challenging. Also the lower floors of the mall contain a large number of moving persons. Figure 2 illustrates an overall view of the mall along with ground-truth path examples and a Tango point cloud (Fig. 2b). Figure 4b shows example frames from the mall and station. The use cases were as realistic as possible including motion in stairs, elevators and escalators, and also temporary occlusions and areas lacking visual features. There are ten sequences from the mall and two from the station.

Office building recordings were performed in the lobby and corridors in two office buildings. They contain some people in a static position and a few people moving. The sequences contain stair climbs and elevator rides. There are closed and open (glass) elevator sequences. Example frames are shown in Fig. 4a.

The outdoor sequences were recorded in the city center (urban, two sequences) and university campus (suburban, two sequences). Figures 4c and 4d illustrate example frames from both locations. Urban outdoor captures were performed through city blocks; they contain open spaces, people, and vehicles. Suburban outdoor captures were performed through sparsely populated areas. They contain a few people walking and some vehicle encounters. Most of the spaces are open. The average length of the outdoor sequences is 334.6 m, ranging from 133 to 514 m. The outdoor sequences were acquired in different times of the day illustrating several daylight conditions.



**Fig. 5.** (a) Speed histograms; peaks correspond to escalators, stairs, and walking. (b) the histogram for one data set with escalator rides/walking. (c-d) the histogram for roll and yaw. (e) the paths for ground-truth (—), ARKit (—), ARCore (—), Tango/Raw (—), Tango/Area learning (—), ROVIO (—), and PIVO (—)


Figure 5a shows the histograms of different motion metrics extracted from the ground-truth. Figure 5a shows the speed histogram which has three peaks that reflect the three main motion modes. From slower to faster they are escalator, stairs, and walking. Figure 5b shows the speed histogram for just one sequence that contained both escalator rides and normal walking. The orientation histograms show that the phone was kept generally in the same position relative to the carrier (portrait orientation, slightly pointing downward). The pitch angle which reflects the heading direction has a close to uniform distribution.

## 5.1 Benchmark Results

We evaluated two research level VIO systems using the raw iPhone data and the three proprietary solutions run on the respective devices (ARCore on Pixel, ARKit on iPhone, and Tango on the tablet). The research systems used were ROVIO [1, 2, 20] and PIVO [25]. ROVIO is a fairly recent method, which has been shown to work well on high-quality IMU and large field-of-view camera data. PIVO is a recent method which has shown promising results in comparison with Google Tango [25] using smartphone data. For both methods, implementations (ROVIO as part of maplab<sup>5</sup>) from the original authors were used (in odometry-only mode without map building or loop-closures). We used pre-calibrated camera parameters and rigid transformation from camera to IMU, and pre-estimated the process and measurement noise scale parameters.

<sup>5</sup> <https://github.com/ethz-asl/maplab>.

For testing purposes, we also ran two visual-only odometry methods on the raw data (DSO [7] and ORB-SLAM2 [15]). Both were able to track subsets of the paths, but the small field-of-view, rapid motion with rotations, and challenging environments caused them not to succeed for any of the entire paths.



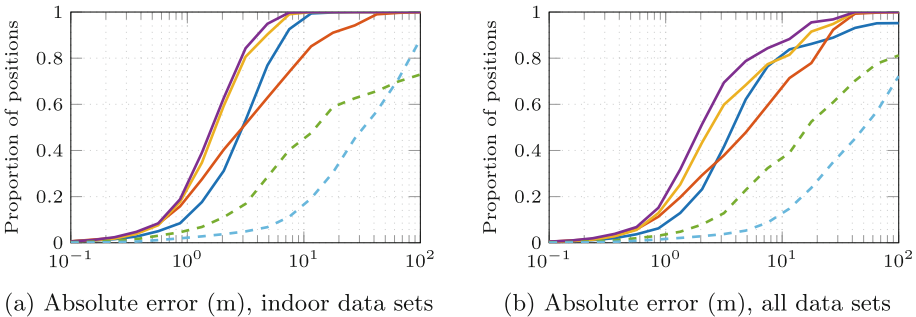
**Fig. 6.** Example paths showing ground-truth (—), ARKit (—), ARCore (—), Tango/Raw (—), and Tango/Area learning (—) that stopped prematurely in (a). Map data © OpenStreetMap. The ground-truth fix points were marked on an architectural drawing. ROVIO and PIVO diverge and are not shown.

In general, the proprietary systems work better than the research methods, as shown in Fig. 7. In indoor sequences, all proprietary systems work well in general (Fig. 7a). Tango has the best performance, ARKit performs well and robustly with only a few clear failure cases (95th percentile  $\sim 10$  m), and ARCore occasionally fails, apparently due to incorrect visual loop-closures. Including the outdoor sequences changes the metrics slightly (Fig. 7b). ARKit had severe problems with drifting in the outdoor sequences. In terms of the orientation error all systems were accurate with less than  $< 2^\circ$  error from the ground-truth on average. This is due to the orientation tracking by integrating the gyroscope performing well if the gyroscope is well calibrated.

As shown in Fig. 7, the research methods have challenges with our iPhone data which has narrow field-of-view and a low-cost IMU. There are many sequences where both methods diverge completely (*e.g.* Fig. 6). On the other hand, there are also sequences where they work reasonably well. This may be partially explained by the fact that both ROVIO and PIVO estimate the calibration parameters of the IMU (*e.g.* accelerometer and gyroscope biases) internally on the fly and neither software directly supports giving pre-calibrated IMU parameters as input. ROVIO only considers additive accelerometer bias, which shows in many sequences as exponential crawl in position. We provide the ground-truth IMU calibration parameters with our data, and it would hence be possible to evaluate their performance also with pre-calibrated values. Alternatively, part of the sequences could be used for self-calibration and others for testing. Proprietary systems may benefit from factory-calibrated parameters. Figures 5e and 6 show examples of the results. In these cases all commercial solutions worked well.

Still, ARCore had some issues at the beginning of the outdoor path. Moreover, in multi-floor cases drifting was typically more severe and there were sequences where also proprietary systems had clear failures.

In general, ROVIO had problems with long-term occlusions and disagreements between visual and inertial data. Also, in Fig. 5e it has clearly inaccurate scale—most likely due to the not modelled scale bias in the accelerations, which is clearly inadequate for consumer-grade sensors that also show multiplicative biases [22]. On the other hand, PIVO uses a model with both additive and multiplicate accelerometer biases. However, with PIVO the main challenge seems to be that without suitable motion the online calibration of various IMU parameters from scratch for each sequence takes considerable time and hence slows convergence onto the right track.



**Fig. 7.** Cumulative distributions of position error: ARKit (—), ARCore (—), Tango/Raw (—), Tango/Area learning (—), ROVIO (---) and PIVO (---).

## 6 Discussion and Conclusion

We have presented the first public benchmark dataset for long-range visual-inertial odometry for hand-held devices using standard smartphone sensors. The dataset contains 23 sequences recorded both outdoors and indoors on multiple floor levels in varying authentic environments. The total length of the sequences is 4.5 km. In addition, we provide quantitative comparison of three proprietary visual-inertial odometry platforms and two recent academic VIO methods, where we use the raw sensor data. To the best of our knowledge, this is the first back-to-back comparison of ARKit, ARCore, and Tango.

Apple’s ARKit performed well in most scenarios. Only in one hard outdoor sequence the ARKit had the classic inertial dead-reckoning failure where the estimated position grew out of control. Google’s ARCore showed more aggressive visual loop-closure use than ARKit, which is seen in false positive ‘jumps’ scattered throughout the tracks (between visually similar areas). The specialized hardware in the Tango gives it an upper hand, which can also be seen in Fig. 7. The area learning was the most robust and accurate system tested. However, all

systems performed relatively well in the open elevator where the glass walls let the camera see the open lobby as the elevator moves. In the case of the closed elevator none of the systems were capable of reconciling the inertial motion with the static visual scene. The need for a dataset of this kind is clear from the ROVIO and PIVO results. The community needs challenging narrow field-of-view and low-grade IMU data for developing and testing new VIO methods that generalize to customer-grade hardware.

The collection procedure scales well to new environments. Hence, in future the dataset can be extended with a reasonably small effort. The purpose of the dataset is to enable fair comparison of visual-inertial odometry methods and to speed up development in this area of research. This is relevant because VIO is currently the most common approach for enabling real-time tracking of mobile devices for augmented reality.

Further details of the dataset and the download links can be found on the web page: <https://github.com/AaltoVision/ADVIO>.

## References

1. Bloesch, M., Burri, M., Omari, S., Hutter, M., Siegwart, R.: Iterated extended Kalman filter based visual-inertial odometry using direct photometric feedback. *Int. J. Robot. Res.* **36**(10), 1053–1072 (2017)
2. Blösch, M., Omari, S., Hutter, M., Siegwart, R.: Robust visual inertial odometry using a direct EKF-based approach. In: *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, pp. 298–304, Hamburg, Germany (2015)
3. Burri, M., Nikolic, J., Gohl, P., Schneider, T., Rehder, J., Omari, S., Achtelik, M.W., Siegwart, R.: The EuRoC micro aerial vehicle datasets. *Int. J. Robot. Res.* **35**, 1157–1163 (2016)
4. Carlevaris-Bianco, N., Ushani, A.K., Eustice, R.M.: University of Michigan north campus long-term vision and LIDAR dataset. *Int. J. Robot. Res.* **35**, 1023–1035 (2015)
5. Ceriani, S., et al.: Rawseeds ground truth collection systems for indoor self-localization and mapping. *Auton. Robots* **27**(4), 353–371 (2009)
6. Cordts, M., et al.: The Cityscapes dataset for semantic urban scene understanding. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3213–3223, Las Vegas, USA (2016)
7. Engel, J., Koltun, V., Cremers, D.: Direct sparse odometry. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(3), 611–625 (2018)
8. Engel, J., Usenko, V.C., Cremers, D.: A photometrically calibrated benchmark for monocular visual odometry. *arXiv preprint [arXiv:1607.02555](https://arxiv.org/abs/1607.02555)* (2016)
9. Everingham, M., Eslami, A., Van Gool, L., Williams, I., Winn, J., Zisserman, A.: The PASCAL visual object classes challenge: a retrospective. *Int. J. Comput. Vis. (IJCV)* **111**(1), 98–136 (2015)
10. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? The KITTI vision benchmark suite. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3354–3361, Providence, Rhode Island (2012)

11. Laskar, Z., Huttunen, S., Herrera, D., Rahtu, E., Kannala, J.: Robust loop closures for scene reconstruction by combining odometry and visual correspondences. In: Proceedings of the International Conference on Image Processing (ICIP), pp. 2603–2607, Phoenix, AZ, USA (2016)
12. Li, M., Kim, B.H., Mourikis, A.I.: Real-time motion tracking on a cellphone using inertial sensing and a rolling-shutter camera. In: Proceedings of the International Conference on Robotics and Automation (ICRA), pp. 4712–4719 (2013)
13. Lin, T., et al.: Microsoft COCO: common objects in context. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 740–755, Zurich, Switzerland (2014)
14. Mourikis, A.I., Roumeliotis, S.I.: A multi-state constraint Kalman filter for vision-aided inertial navigation. In: Proceedings of the International Conference on Robotics and Automation (ICRA), pp. 3565–3572, Rome, Italy (2007)
15. Mur-Artal, R., Tardós, J.D.: ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras. *IEEE Trans. Robot.* **33**(5), 1255–1262 (2017)
16. Mur-Artal, R., Tardós, J.D.: Visual-inertial monocular SLAM with map reuse. *Robot. Autom. Lett.* **2**(2), 796–803 (2017)
17. Nikolic, J., et al.: A synchronized visual-inertial sensor system with FPGA pre-processing for accurate real-time SLAM. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), pp. 431–437, Hong-Kong, China (2014)
18. Pfrommer, B., Sanket, N., Daniilidis, K., Cleveland, J.: PennCOSYVIO: a challenging visual inertial odometry benchmark. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), pp. 3847–3854, Singapore (2017)
19. Russakovsky, O., et al.: ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis. (IJCV)* **115**(3), 211–252 (2015)
20. Schneider, T., et al.: Maplab: an open framework for research in visual-inertial mapping and localization. *IEEE Robot. Autom. Lett.* **3**(3), 1418–1425 (2018)
21. Schöps, T., Engel, J., Cremers, D.: Semi-dense visual odometry for AR on a smartphone. In: Proceedings of the International Symposium on Mixed and Augmented Reality (ISMAR), pp. 145–150 (2014)
22. Shelley, M.A.: Monocular visual inertial odometry on a mobile device. Master's thesis, Technical University of Munich, Germany (2014)
23. Smith, M., Baldwin, I., Churchill, W., Paul, R., Newman, P.: The new college vision and laser data set. *Int. J. Robot. Res.* **28**(5), 595–599 (2009)
24. Solin, A., Cortes, S., Rahtu, E., Kannala, J.: Inertial odometry on handheld smartphones. In: Proceedings of the International Conference on Information Fusion (FUSION), Cambridge, UK (2018)
25. Solin, A., Cortes, S., Rahtu, E., Kannala, J.: PIVO: probabilistic inertial-visual odometry for occlusion-robust navigation. In: Proceeding of the IEEE Winter Conference on Applications of Computer Vision (WACV), Lake Tahoe, NV, USA (2018)
26. Sturm, J., Engelhard, N., Endres, F., Burgard, W., Cremers, D.: A benchmark for the evaluation of RGB-D SLAM systems. In: Proceedings of the International Conference on Intelligent Robot Systems (IROS), pp. 573–580 (2012)