Article

The International Journal of Robotics Research
2018, Vol. 37(1) 13–20

DOI: 10.1177/0278364917751842
journals.sagepub.com/home/ijr

The Visual–Inertial Canoe Dataset

Martin Miller[1], Soon-Jo Chung[2] and Seth Hutchinson[1]

Abstract

We present a dataset collected from a canoe along the Sangamon River in Illinois. The canoe was equipped with a stereo camera, an inertial measurement unit (IMU), and a global positioning system (GPS) device, which provide visual data suitable for stereo or monocular applications, inertial measurements, and position data for ground truth. We recorded a canoe trip up and down the river for 44 minutes covering a 2.7 km round trip. The dataset adds to those previously recorded in unstructured environments and is unique in that it is recorded on a river, which provides its own set of challenges and constraints that are described in this paper. The dataset is stored on the Illinois Data Bank and can be accessed at: `https://doi.org/10.13012/B2IDB-9342111_V1.`

Keywords

SLAM, dataset, stereo, inertial, IMU, monocular, vision, localization, mapping, robotics

1. Introduction

In recent years some researchers working on simultaneous localization and mapping (SLAM) have turned their attention to SLAM in natural river environments. The typical river environment poses unique problems that are not found in a ground-based urban setting. The most dramatic difference is that in place of a roadway, there is water. The water of a river behaves differently according the conditions of the river, wind, and lighting. When the water is still, the scenery forms a near-perfect reflection in the water. On the other hand, on a windy day or near obstacles, the water forms ripples that distort or eliminate the mirror image of the scene. Further, the mirrorlike surface of water can produce specular reflections that challenge many imaging systems. Scherer et al. (2012) developed a learning based method to segment the river from the rest of the scene in order to avoid the complications of the water. Yang et al. (2017) used correspondences between objects and their reflections as a pseudo-stereo solution to improve depth estimation in a monocular system. However, this system faced difficulties in tracking features in a natural environment with repeated textures. These problems were compounded by inadequate exposure control for the camera, which caused underexposure of landmarks.

There are many publicly available datasets for SLAM research and robotic navigation. Many, such as Smith et al. (2009) and Blanco et al. (2014), are ground-based urban datasets. More complex ground-based datasets have been released such as the KITTI dataset (Geiger et al., 2013), which combines high-quality data, a broad range of environments, and evaluation tools. More recently, Maddern et al. (2017) have released a dataset collected on 100 days over the course of a year to address the challenges of long-term autonomous driving. More similar to the natural environment of a river are the datasets of Furgale et al. (2012b), Tong et al. (2013), and Peynot et al. (2010), collected in open environments containing few, if any, man-made objects. Leung et al. (2017) published a dataset inside of a copper mine, which is a closed irregularly shaped environment.

In this paper we describe a new dataset, similar to previous work in open or irregularly shaped environments, but unique in that it was collected on a river. The Visual–Inertial Canoe Dataset augments the current datasets by making available a high-quality dataset collected on a river in a natural setting. We collected data using a high-quality global positioning system (GPS) and inertial navigation system (INS) sensor and a stereo camera on the Sangamon River near Mahomet, Illinois. Inertial data is included as an aid in estimating the pose of the canoe, particularly when vision

[1]University of Illinois at Urbana-Champaign, Urbana, IL, USA
[2]Graduate Aerospace Laboratories (GALCIT) and Jet Propulsion Laboratory, Pasadena, CA, 91125, USA

Corresponding author:
Seth Hutchinson, Department of Electrical and Computer Engineering, University of Illinois, Urbana, IL, 61801, USA.
Email: seth@illinois.edu

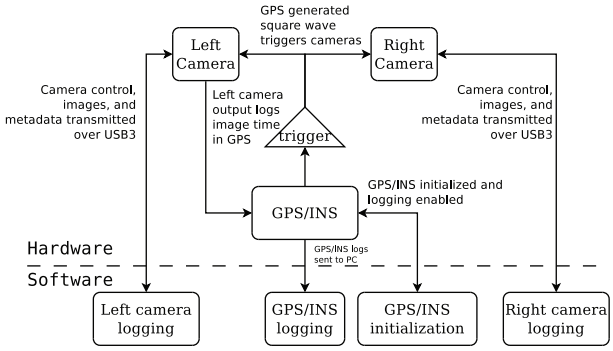

Fig. 1. Block diagram of the sensor hardware and software.

Fig. 2. The canoe and sensors are in the water of the Sangamon River just prior to data collection.

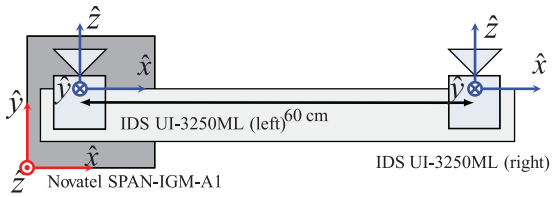

Fig. 3. Diagram of sensor locations and axes.

measurements are sparse or unavailable. The visual data is calibrated and suitable for monocular or stereo applications. We expect that this data will be valuable not only to researchers working on river navigation, but for those seeking data in a forested environment. The dataset is being released to the public for the benefit of researchers and, to the best of the authors' knowledge, is the first publicly released riverine dataset.

2. Sensors

This section describes the hardware and software used for data collection. This includes the sensors, computer, mounting hardware, and the vehicle itself. We use the term platform to describe the collection of these items and their relative locations. A block diagram of the sensor platform is shown in Figure 1. A complete listing of the hardware is given in Table 1.

2.1. Platform

The platform consists of a 16 ft long canoe with the sensors mounted at the nose, and a laptop in the middle. The canoe seats one paddler in the front behind the sensors and one in the rear behind the laptop. The sensors are mounted to an aluminum bar, which is precisely machined so that the cameras and GPS/INS are rigidly attached. There are several drillings in the bar that the left and right cameras can be mounted to for recording with different stereo baselines. The cameras are mounted to the top and the GPS/INS device is mounted below, under the left-hand camera. The bar is attached to a wooden board using clamps and this board, in turn, is clamped to the canoe. The location of the bar with respect to the canoe is not precise. A photo of the platform is shown in Figure 2, and a diagram showing the sensor locations and axes is provided in Figure 3.

2.2. GPS/INS

We used a Novatel SPAN-IGM-A1 GPS/INS[1] device for ground truth, inertial measurements, and time synchronization. A GPS/INS device refines GPS position estimates with inertial measurements from an inertial measurement unit (IMU). The refined position, velocity, and attitude, as well as the raw acceleration, and angular velocity measurements are recorded for the dataset. The Novatel unit can log external events marked by a rising or falling edge on one of its input lines, which allows us to use the internal clock of the device as a common clock for camera events as well.

The SPAN-IGM-A1 has 1.2 m RMS horizontal position accuracy. We log position, velocity, and attitude at 10 Hz and their corresponding covariance matrices at 1 Hz. The inertial measurements are made with a microelectromechanical system (MEMS) IMU, an Analog Devices ADIS-16488 located inside the Novatel device, and recorded at 200 Hz. Position is also logged in Universal Transverse Mercatur (UTM) coordinates, a system of representing position on the globe in units of meters on a Cartesian grid, at 0.2 Hz. See Snyder (1987) for a description of the UTM system.

The Novatel unit logs all data using GPS time. That is, the number of seconds since midnight January 6, 1980. The device outputs a square wave that is used to trigger the stereo camera. When the camera exposure begins, the device receives a rising edge from the left camera and records the time of this event. The logs are transmitted in binary to the computer over a serial connection.

2.3. Stereo Camera

The stereo camera setup consists of two IDS UI-3250ML[2] cameras fitted with a pair of Kowa LM3NC1M[3] 3.5 mm $f/2.4$ lenses and mounted with an approximate baseline of

Table 1. Description of each component of the platform.

Model	Description	Coordinates
Novatel SPAN-IGM-A1	A GPS/INS device. Logs position, velocity, attitude and inertial measurement unit (IMU) data from onboard sensors. Triggers camera shutter and logs time of image capture. Sensors are logged using the following Novatel messages defined in reference manuals Novatel (2013) and Novatel (2015): • inspvas: Position (latitude–longitude), velocity, and attitude at 10 Hz • inscovs: Covariance matrices of position, velocity, and attitude at 1 Hz • bestutm: Best available position estimate in UTM coordinates at 0.2 Hz • rawimus: Raw acceleration and angular rate at 200 Hz • mark2time: Event log of left camera trigger at 20 Hz	right, forward, up
IDS UI-3250ML	A camera with 1600×1200 native resolution. Data is collected using 800×600 resolution at 20 Hz. The shutter is triggered externally by a signal from the SPAN-IGM-A1. A pair of these cameras is mounted on an aluminum bar with a $\sim$60 cm baseline.	right, down, forward
Kowa LM3NC1M	A 3.5 mm $f/2.4$ machine vision lens. The lens provides a roughly 90° field of view (FOV). Used on both cameras.	N/A
Trigger unit	A custom circuit that conditions the signals between the SPAN-IGM-A1 and the pair of UI-3250MLs	N/A
Dell Latitude E5450	A laptop running Fedora Linux and equipped with a solid state drive (SSD). It runs the software that initializes the SPAN-IGM-A1 and the UI-3250MLs and it stores the data logs and video feed.	N/A
Old Town Penobscot 16 canoe	A 16 ft canoe that can seat two paddlers. The sensors are mounted at the nose of the canoe.	N/A

60 cm. The cameras were configured in order to maintain a constant frame rate with a low probability of dropped frames, to automatically adjust exposure settings under dynamic lighting conditions, and to provide precise timing information to the rest of the system.

In order to reduce the strain on bandwidth over the USB3 bus and the hard disk, images are stored in the raw Bayer format (Bayer, 1976) and the resolution is reduced using on chip pixel binning (Howell, 2006).

As the canoe winds along the river, the lighting conditions can change very quickly due to the position of the sun with respect to the camera and the amount of tree canopy over the river. Automatic exposure control is used to adjust the cameras to these dynamic lighting conditions. River scenes can be dominated by the sky and its reflection causing a typical exposure controller to compensate for the brightness of the sky by reducing exposure time. This causes the trees and the shore to appear dark and without detail. To reduce this effect, we used spot metering, which is the application of automatic exposure control to an area of interest (AOI). An AOI is a preselected region that is used for the averaging and the area outside of the AOI does not contribute to the average. We set the AOI to a region in the middle of the frame, where the trees and shoreline are typically found. Automatic exposure control has some downsides. It adds processing overhead that has to be compensated for by binning or reduced frame rate. In addition, the exact exposure time used is not known and it is possible that each camera uses a different value for the exposure time, because the control operates in each camera independently. This was mitigated by setting an upper limit of 16 ms on the exposure time. In practice, we found that in daylight the cameras typically operate with much shorter exposure times than the maximum.

The timing requirements were satisfied by using an externally supplied square wave to trigger the right and left cameras simultaneously. Upon receiving a falling edge from the hardware trigger, the right and left cameras begin their exposure. The left camera is configured with an active high output that results in a rising edge when an exposure begins. The edge is transmitted to the Novatel SPAN-IGM-A1 GPS/INS unit through a custom hardware trigger circuit, which logs the GPS time of the event.

The above settings were coupled with a 20 Hz camera trigger. The combination of these settings results in very few or no dropped frames in typical use and excellent exposure under most lighting conditions. It should be noted that binning and exposure control were disabled during the calibration procedures where the additional precision in timing and resolution are necessary.

2.4. Computer

A Dell Latitude E5450 laptop running Fedora Linux is used to control the sensors and record the data. The computer has a 500 GB SSD hard drive. USB3 is used for the cameras and USB2 is used for serial communication with the Novatel unit. This computer ran all of the software used for operating the sensors.

2.5. Software

Custom software communicates with the Novatel device in order to turn on logging, set the logging rates, and start camera triggering. Incoming log messages arrive at the PC as binary, are parsed into ASCII, and saved to the hard drive. All messages sent by the Novatel unit are followed by a CRC-32 hash, which can be used to check for transmission errors. The CRC-32 hash of the message is computed on the laptop in accordance with Novatel (2016) and compared with the received hash. In the case that the hashes do not match, the entire message is discarded. In practice, these events occur rarely, if ever.

Another custom program is used to control the cameras and save the incoming images. Each camera is controlled by a separate camera process, which allows us to take advantage of the parallel processing ability of a multicore CPU. This software initializes the camera settings, displays a live view of the video, saves the incoming images, and logs the image metadata. The image metadata includes a software defined frame number, a camera defined frame number, the computer's clock time when the image was received, and the camera's internal clock time when the image was captured. By comparing this data to the image times logged by the Novatel device, we can determine the presence of glitches or dropped frames.

3. Calibration

In this section we describe the calibration of our system. We will discuss the timing accuracy of the SPAN-IGM-A1, the stereo camera calibration process, the camera to IMU calibration, and the measurements made of the canoe itself. We will also describe how to use the transforms given by the calibration results.

3.1. Time synchronization

The Novatel SPAN-IGM-A1 logs events using GPS time and has 20 ns time accuracy. Position, velocity, attitude, and their respective covariance matrices are logged synchronously, that is the logs are written at regular intervals. This is not the case for the IMU measurements, which are asynchronous. The IMU measurements are written to the log as soon as new readings are received from the sensor. While these measurements arrive nominally at 200 Hz, the period is not precisely fixed and they do not necessarily arrive in phase with other messages.

3.2. Stereo cameras

Stereo calibration images were collected the day after the dataset was recorded. During calibration automatic exposure control was disabled and the exposure period was fixed at 2 ms. Pixel binning was disabled and images were recorded at full 1600×1200 resolution. The camera intrinsic parameters, lens distortion and camera extrinsics were determined simultaneously using the *kalibr* package for Robot Operating System (ROS) (Maye et al., 2013).

With the camera intrinsic and extrinsic parameters recovered, the raw images can be rectified into row-aligned format for stereo matching and the camera to IMU calibration procedure can be performed.

3.3. Camera to IMU

Camera to IMU calibration was performed using the *kalibr_calibrate_imu_camera* tool (Furgale et al., 2013, 2012a) from the *kalibr* package for *ROS*. This tool computes the rotation and translation of the cameras with respect to the IMU axes. The results of the calibration procedure are provided in the file `camchain-imucam.yaml`, which conforms to the output in the *kalibr* documentation.

3.4. Canoe measurements

A measurement from the origin of the IMU to the waterline was made using a tape measure during data collection, while the canoe was moving with two paddlers. During canoe motion, the shifting weight of the paddlers and general rocking motion of the canoe can cause the actual height above water to vary significantly. Nevertheless, the measurement is provided here as a rough guideline. The measured height of the IMU of above the water is 46 cm.

3.5. Transformations

There are separate coordinate frames for the IMU, the cameras and the rectified cameras. The image transformations are of the same form as Geiger et al. (2013) and we use the same notation as much as possible for clarity. The superscripts 0 and 1 are used to represent the left and right cameras, respectively, and the subscript *rect* is used for the rectified frame. The absence of a subscript refers to the unrectified frame. For example, $\mathbf{R}^{(0)}_{\text{rect}}$ represents the rectifying rotation for the left camera.

We take the unrectified left camera to be the reference frame. A point in the three-dimensional world represented in the unrectified left camera frame as $\mathbf{X} = \begin{bmatrix} x & y & z & 1 \end{bmatrix}^\top$ is projected onto the rectified camera plane by its respective camera matrix, rotation, and translation, denoted by $\mathbf{K}^{(\mathrm{i})} \in \mathbb{R}^{3\times3}$, $\mathbf{R}^{(\mathrm{i})} \in \mathbb{R}^{3\times3}$, and $\mathbf{t}^{(\mathrm{i})} \in \mathbb{R}^{3\times1}$, respectively. In the case of the left camera, $\mathbf{R}^{(0)} = \mathbf{I}_{3\times3}$ and $\mathbf{t}^{(0)} = \mathbf{0}_{3\times1}$

$$\mathbf{x}^{(\mathrm{i})} = \mathbf{K}^{(\mathrm{i})} \begin{bmatrix} \mathbf{R}^{(\mathrm{i})} & \mathbf{t}^{(\mathrm{i})} \end{bmatrix} \mathbf{X}$$

We use the projection matrix $\mathbf{P}^{(\mathrm{i})}_{\text{rect}} \in \mathbb{R}^{4\times4}$ to project the same point into the rectified frame. For dimensional consistency, a fourth row and column is added to $\mathbf{R}^{(0)}_{\text{rect}}$ with the element in the fourth row and column set to 1 and the other values of the fourth row and column set to zero. The following equation describes the rectification:

$$\mathbf{x}^{(\mathrm{i})}_{\text{rect}} = \mathbf{P}^{(\mathrm{i})}_{\text{rect}} \mathbf{R}^{(0)}_{\text{rect}} \mathbf{X}$$

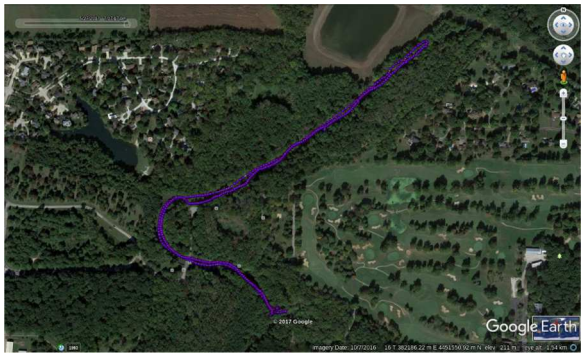

Fig. 4. The GPS trajectory overlaid onto a satellite image. Data was recorded during both the outbound and return journeys.

Note that while the *i*th $\mathbf{P}$ matrix is used, $\mathbf{R}_{\text{rect}}^{(0)}$ is used for both cameras, because this matrix restores the rectified plane (shared by both cameras) to the reference plane of the left camera.

The result of the camera to IMU calibration described in Section 3.3 provides a transformation matrix $\mathbf{T}_{\text{imu}}^{\text{cam}} \in \mathbb{R}^{4\times4}$ that transforms points from the IMU coordinate system into the left unrectified camera frame. A point $\mathbf{y}^{\text{imu}} \in \mathbb{R}^{4\times1}$ in the IMU coordinate system is projected into the left unrectified camera frame by

$$\mathbf{x}^{(0)} = \mathbf{T}_{\text{imu}}^{\text{cam}}\mathbf{y}^{\text{imu}}$$

and it is projected into the rectified camera frame by

$$\mathbf{x}_{\text{rect}}^{(\text{i})} = \mathbf{P}_{\text{rect}}^{(\text{i})}\mathbf{R}_{\text{rect}}^{(0)}\mathbf{T}_{\text{imu}}^{\text{cam}}\mathbf{y}^{\text{imu}}$$

4. Dataset

The data being released was collected on May 27, 2017 over a period of 44 minutes. The route traverses upstream on the Sangamon River and then back downstream, returning to the starting point, traveling a total distance of approximately 2.7 km. The recorded GPS track is shown in Figure 4. There are 53,275 images per camera and zero dropped frames. There are 532,744 IMU readings. The majority of the trajectory is in a forested area lined with trees on both sides. At times the tree canopy encloses much of the channel. One exception to the forested scenery occurs when the canoe passes under a covered bridge. Representative examples of the scenery are shown in Figures 6–9.

We divided the data into subsets. Each upstream subset has a complementary downstream subset that covers the same portion of the river in the opposite direction, which we believe can be useful for loop closure. The subsets are non-overlapping and the concatenation of all datasets is equal to the full dataset with some truncation at the start and end. A description of the subsets is provided in Table 2.

4.1. Organization and file formats

The dataset is available in subsets that can be downloaded individually. In each subset, the data organization follows the same format. The directory structure for the data is shown in Figure 5.

```
name/
├── name.csv
├── name.cam0
├── name.cam1
├── camchain-imucam.yaml
├── params.yaml
├── README
├── cam0/
│   └── <imagenumber>.bmp
├── cam1/
│   └── <imagenumber>.bmp
└── rectified/
    ├── cam0/
    │   └── <imagenumber>.png
    └── cam1/
        └── <imagenumber>.png
```

Fig. 5. Directory structure for a data subset. `name` refers to the subset's name. `imagenumber` refers to the sequential frame number of an image.

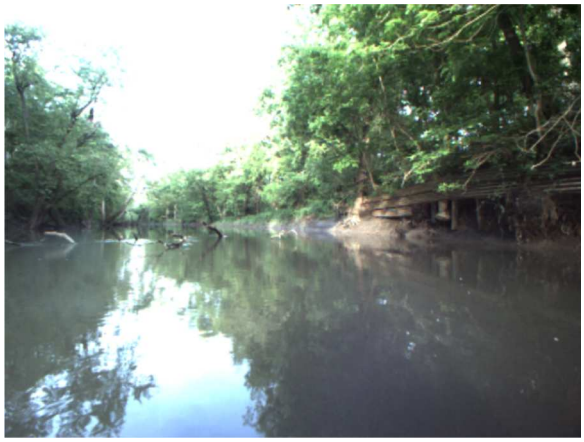

Fig. 6. Spot metering prevents the bright sky from dominating the exposure control and making the trees and shore appear too dark.

Fig. 7. A typical image from the dataset. The shoreline is well exposed, many reflections are visible, and the canopy blocks the sun

Table 2. Each row describes an upstream subset and the corresponding downstream dataset over the same part of the river.

Name	Length (mm:ss/m)	Notes	Name	Length (mm:ss/m)	Notes
0	5:05/65	We drag the canoe over land and launch into the water.	9	3:25/40	We return to the starting point.
1	4:10/222	There is some lens flare present.	8	2:20/216	Well exposed throughout.
2	5:30/335	We cross under a covered bridge and there is minor lens flare until the turn.	7	3:45/329	We cross back under the bridge. Well exposed throughout.
3	8:00/479	Well exposed.	6	5:10/492	Well exposed.
4	3:50/258	We approach the turnaround	5	2:50/257	Perform U-turn and head downstream. There is some lens flare during turnaround.

Fig. 8. The water with some ripples.

Fig. 9. The canoe passes under a covered bridge.

The sensor data is provided in `name.csv`, where *name* is the name of the dataset. One sensor message is stored per line. Each line contains a timestamp, the message name, and the message data. The message names and field order are consistent with Novatel documentation. In the case of the `rawimus` message, the accelerometer and angular velocity measurements have been converted from their native fixed point machine representation to units of meters per second squared and radians per second, respectively. Images are printed with the `IMG` message name and followed by the left and right image filenames.

Images are available in the raw Bayer format and in stereo rectified format. The raw Bayer images have had no pre-processing applied to them. The stereo rectified images have been converted into RGB color, transformed to eliminate lens distortion and row-aligned so that they are ready for stereo matching. The rectification was performed using custom software based on OpenCV's sample stereo rectification code (Bradski and Kaehler, 2008).

The camera and IMU calibration parameters are provided in `camchain-imucam.yaml` and `params.yaml`. The file format of `camchain-imucam.yaml` matches the output of the *kalibr* calibration suite, which was used for both the stereo camera and camera to IMU calibration. The camera intrinsic parameters are provided for full resolution 1600×1200 images. In order to use these parameters with the 800×600 data, the values of the camera focal length and center pixel should be scaled by one half and the distortion coefficients remain the same. The format of `params.yaml` is suitable for use with the stereo rectified images and follows the format used in Geiger et al. (2013). This file provides the rotation and projection matrices for working with the undistorted and rectified stereo images provided in the `rectified` directory. The notation follows that of Section 3.5 and the parameters are corrected for the scaled resolution.

The file `README` provides basic information about the dataset formatting. Each message is described and the message fields in `name.csv` are described. Researchers should be able to utilize the data referring only to information provided in these directories, however reference to this paper and hardware documentation listed here is still recommended.

4.2. Challenges

Collecting data in a river environment presented difficulties with controlling exposure, lens flare, and maintaining a GPS fix. Precautions were taken to mitigate these issues, and the data subsets were defined in an attempt to isolate periods with issues to a single subset. Overall our approach

was able to limit the periods with low-quality data to a small portion of the total data collected.

To compensate for exposure problems, the auto-exposure approach described in Section 2.3 was used to good effect. In a few cases where the cameras are pointed towards the sun, there is noticeable lens flare. As the canoe approaches the covered bridge while traveling upstream the lens flare is constant, but minor. Later, when the canoe makes a U-turn and returns downstream, there is a large lens flare of short duration during the U-turn.

The GPS system was able to maintain a position estimate throughout the data collection, however the quality of the satellite fix deteriorates for about 5 minutes on the return journey. This is largely unavoidable due to the tree canopy in the river environment. The poor estimates are mostly confined to subset 6. During subset 6, the standard deviation of error in the horizontal position starts at ∼15 m, grows to over 50 m, and ends at 5 m. Upon examination of the GPS trajectory, there are obvious errors between seconds 1,179,964,415 and 1,179,964,455 (GPS time). However, the position estimate recovers, and the typical standard deviation in the horizontal position estimate is between 1.5 and 5.0 m. Users of the dataset can consult the `inscovs` messages to determine the quality of the GPS at each second.

5. Evaluation

We propose a metric for evaluating SLAM results generated with this dataset, which is similar to the metric used in Geiger et al. (2012). For translation error we propose using the mean percentage error. That is, at every timestep k, both the positional error,

$$e_p[k] = |\mathbf{p}[k] - \hat{\mathbf{p}}[k]|$$

and the total distance traveled,

$$D[k] = \sum_{j=1}^{k} |\mathbf{p}[j] - \mathbf{p}[j-1]|$$

are computed, where $\mathbf{p}[k] \in \mathbb{R}^3$ is the position of the GPS receiver at time k and $\hat{\mathbf{p}}[k] \in \mathbb{R}^3$ is its estimate. The percentage error is

$$e_{pct}[k] = e_p[k]/D[k]$$

The average of $e_{pct}[k]$ over all k is the translation error.

For attitude error, the mean in degrees per meter is calculated. That is, at every timestep the attitude error is computed in degrees

$$e_q[k] = \frac{360}{\pi} \arccos(\mathbf{q}[k] \cdot \hat{\mathbf{q}}[k])$$

where $\mathbf{q}[k] \in \mathbb{H}$ is the attitude of the GPS receiver represented as a quaternion and $\hat{\mathbf{q}}[k] \in \mathbb{H}$ is its estimate. Degrees per meter is calculated as

$$e^{q}_{pct}[k] = e_q[k]/D[k]$$

and the average of $e^{q}_{pct}[k]$ over all k is the attitude error.

When evaluating error, we propose that errors be computed over entire subsets and that subset 6 not be used for comparisons due to the large uncertainty in the GPS position described in Section 4.2.

6. Summary

We collected stereo images, inertial measurements, and ground-truth data on the Sangamon River in Illinois. This unique environment provides new opportunities for researchers to approach problems in riverine navigation. The abundance of images containing objects and their reflections in the water will provide ample data for those working on SLAM approaches that utilize these features. We expect that this high-quality data will be useful to researchers in SLAM and also computer vision researchers working on vision problems in riverine and forested environments. The data described above is the first publicly available riverine dataset for SLAM research. In the future, we may expand the dataset by adding additional sensors, recording data in other river environments, or using a faster vehicle.

Notes

1. https://www.novatel.com/products/span-gnss-inertial-systems/span-combined-systems/span-igm-a1/
2. See https://en.ids-imaging.com/store/products/cameras/ui-3250ml.html
3. See https://lenses.kowa-usa.com/wide-angle-megapixel-ncm-series/408-lm3nc1m.html

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The work in this article was supported by the Office of Naval Research (grant number N00014-14-1-0265).

ORCID iD

Martin Miller https://orcid.org/0000-0003-3520-4946
Soon-Jo Chung https://orcid.org/0000-0002-6657-3907
Seth Hutchinson https://orcid.org/0000-0002-3949-6061

References

Bayer B (1976) Color imaging array. US Patent 3,971,065.

Blanco JL, Moreno FA and González-Jiménez J (2014) The Málaga Urban Dataset: High-rate stereo and lidars in a realistic urban scenario. *The International Journal of Robotics Research* 33(2): 207–214.

Bradski G and Kaehler A (2008) *Learning OpenCV: Computer Vision with the OpenCV Library*. O'Reilly Media.

Furgale P, Rehder J and Siegwart R (2013) Unified temporal and spatial calibration for multi-sensor systems. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Tokyo, Japan, pp. 1280–1286.

Furgale PT, Barfoot TD and Sibley G (2012a) Continuous-time batch estimation using temporal basis functions. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, St. Paul, MN, pp. 2088–2095.

Furgale PT, Carle P, Enright J and Barfoot TD (2012b) The devon island rover navigation dataset. *The International Journal of Robotics Research* 31(6): 707–713.

Geiger A, Lenz P, Stiller C and Urtasun R (2013) Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* 32(11): 1231–1237.

Geiger A, Lenz P and Urtasun R (2012) Are we ready for autonomous driving? The KITTI vision benchmark suite. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*.

Howell S (2006) *Handbook of CCD astronomy*. Cambridge: Cambridge University Press.

Leung K, Lühr D, Houshiar H, et al. (2017) Chilean underground mine dataset. *The International Journal of Robotics Research* 36(1): 16–23.

Maddern W, Pascoe G, Linegar C and Newman P (2017) 1 year, 1000 km: The Oxford Robotcar Dataset. *The International Journal of Robotics Research* 36(1): 3–15.

Maye J, Furgale PT and Siegwart R (2013) Self-supervised calibration for robotic systems. In: *Intelligent Vehicles Symposium*.

Novatel (2013) Novatel OM-20000141. *SPAN-IGM User Manual*. Novatel. Rev. 2.

Novatel (2015) Novatel OM-20000144. *SPAN on OEM6 Firmware Reference Manual*. Novatel. Rev. 7.

Novatel OM-20000129 (2016) Novatel OM-20000129. *OEM6 Family Firmware Reference Manual*. Novatel. Rev. 9.

Peynot T, Scheding S and Terho S (2010) The Marulan Data Sets: Multi-sensor perception in a natural environment with challenging conditions. *The International Journal of Robotics Research* 29(13): 1602–1607.

Scherer S, Rehder J, Achar S, et al. Cover H, Chambers AD, Nuske ST and Singh S (2012) River mapping from a flying robot: state estimation, river detection, and obstacle mapping. *Autonomous Robots* 32(5): 1–26.

Smith M, Baldwin I, Churchill W, Paul R and Newman P (2009) The new college vision and laser data set. *The International Journal of Robotics Research* 28(5): 595–599.

Snyder JP (1987) Map projections: A working manual. Technical report, US Geological Survey, Washington, DC. http://pubs.er.usgs.gov/publication/pp1395.

Tong CH, Gingras D, Larose K, Barfoot TD and rick Dupuis (2013) The Canadian Planetary Emulation Terrain 3D Mapping Dataset. *The International Journal of Robotics Research* 32(4): 389–395.

Yang J, Dani AP, Chung SJ and Hutchinson S (2017) Vision-based localization and robot-centric mapping in riverine environments. *Journal of Field Robotics* 34: 429–450.