Brno Urban Dataset - The New Data for Self-Driving Agents and Mapping Tasks

LIGOCKI, A.; ŽALUD, L.; JELÍNEK, A.

Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA) ,
pp. 3284-3290

eISBN: 978-1-7281-7395-5

ISSN: 2577-087X

DOI: <https://doi.org/10.1109/ICRA40945.2020.9197277>

Accepted manuscript

©2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. LIGOCKI, A.; ŽALUD, L.; JELÍNEK, A. "Brno Urban Dataset - The New Data for Self-Driving Agents and Mapping Tasks", IEEE International Conference on Robotics and Automation (ICRA), 2020. DOI: 10.1109/ICRA40945.2020.9197277. Final version is available at <https://ieeexplore.ieee.org/document/9197277>

Brno Urban Dataset - The New Data for Self-Driving Agents and Mapping Tasks

Adam Ligocki¹, Ales Jelinek¹ and Ludek Zalud¹

Abstract—Autonomous driving is a dynamically growing field of research, where quality and amount of experimental data is critical. Although several rich datasets are available these days, the demands of researchers and technical possibilities are evolving. Through this paper, we bring a new dataset recorded in Brno - Czech Republic. It offers data from four WUXGA cameras, two 3D LiDARs, inertial measurement unit, infrared camera and especially differential RTK GNSS receiver with centimetre accuracy which, to the best knowledge of the authors, is not available from any other public dataset so far. In addition, all the data are precisely timestamped with sub-millisecond precision to allow wider range of applications. At the time of publishing of this paper, recordings of more than 350 km of rides in varying environment are shared at: <https://github.com/RoboticsBUT/Brno-Urban-Dataset>.

I. INTRODUCTION

Research in the domain of autonomous mobile vehicles have tremendously expanded in the last few years [1], [2], [3]. From one of many possible applications of general mobile robotics [4] and a geeky interest of technical visionaries it became a large topic for both scientific and commercial sectors. Despite the undoubted motivation of financial bounties and pursuit of the emerging trends, this boom is also fueled with openly available data allowing more people to be part of it. To equip a car with state of the art sensors can easily become too expensive for small subjects such as start-ups or research groups on local universities. Sharing data allows much more researchers to participate in the progress of the field and enrich it with novel ideas, which, in the end, rewards everybody [5]. Second good reason for data sharing is a possibility to bypass the necessity of building and maintaining the sensory apparatus, which otherwise requires extra resources and engineering skills not related to the actual research topic of artificial intelligence. Having the opportunity to build our own data acquisition system and exceeding current state of the art in some of its parameters, we have decided to make the data publicly available.

The research was supported by ECSEL JU under the project H2020 737469 AutoDrive - Advancing fail-aware, fail-safe, and fail-operational electronic components, systems, and architectures for fully automated driving to make future mobility safer, affordable, and end-user acceptable. This research has been financially supported by the Ministry of Education, Youth and Sports of the Czech republic under the project CEITEC 2020 (LQ1601).

¹All the authors are with the Central European Institute of Technology (CEITEC), Cybernetics in Material Science research group, Brno University of Technology, Purkynova 123, Brno-Kralovo Pole, Czechia, adam.ligocki@ceitec.vutbr.cz, ales.jelinek@ceitec.vutbr.cz, ludek.zalud@ceitec.vutbr.cz

Fig. 1. The detail of the Atlas sensory framework, which has been used to record all the data published in the Brno Urban Dataset. On the image there are four RGB cameras, single thermal (IR) camera, two Velodyne HDL-32E LiDARs, the RTK GNSS receiver with pair of differential antennas and the IMU unit in the center of the frame.

Through this paper, we provide an urban dataset recorded in Brno - Czech Republic and its near surroundings. The location offers highly diverse environments - from highways to dirt roads, from densely build-up areas to woods, all recorded in conditions of the real traffic. The sensory system contains all sensors which are considered standard these days [1], [6]: frontal stereo cameras, lateral cameras, two 3D laser scanners (LiDARs), inertial measurement unit (IMU) and global navigation satellite system (GNSS) receiver. Beyond that, there are three key features, which make the system unique. First, to our best knowledge, it is the only measurement system equipped with differential real-time kinematic (RTK) GNSS receiver providing centimetre accurate global positioning and heading at the same time. Second, a thermal camera is used to sense the scene in front of the car, which greatly enhances vulnerable road user awareness [7] and general detection of various objects in bad weather conditions [8]. Third, most of the system is exactly synchronized and timestamped using the GPS signal with sub-millisecond precision.

Usage of such data is wide. From practical point of view, having a reliable reference eliminates many problems in data processing. In general pose estimation, the benefits of differential GNSS are well understood [9], [10] and its application in autonomous road vehicles seems reasonable. Similarly, issues arising from bad synchronization of cameras are subject of dedicated research coming-up with creative ways of dealing with them [11], [12], but our system allows

to avoid such problems entirely, which is probably going to be the case in modern autonomous vehicles as well.

Another possible application of our data consists in development of the previously mentioned methods for dealing with uncertainty, because in some occasions their adoption is necessary. Exact timestamping allows evaluation of such methods, where timing of the data acquisition is part of the stochastic model [13], [14]. With precise global localization, we can benchmark the methods relying on local sensors such as cameras [12], [15], LiDARs [16], [17] and their combinations, frequently enhanced with inertial measurement units [18]. As we will show in the next section, the purpose of most of the other datasets is different and would be of limited usage in such experiments.

II. RELATED WORKS

Publicly available datasets are rapidly evolving according to the fast technological progress of the last two decades. Sensor accuracy and especially resolutions have grown substantially, which led to larger memory requirements for storing and higher bandwidth needs for sharing of the recordings. Today's possibilities fulfill all these demands, which makes older datasets exceeded quite fast. For example, going through the review in the Málaga dataset paper [19], we see that datasets captured five to ten year ago provide mostly 2D laser scans, camera resolution of a fraction of a megapixel and the largest had tens of kilometers of recorded path. Comparing them with more recent surveys [6] and [1], the standard have risen significantly. Unfortunately many large datasets are not freely available [20]. The following paragraphs cover only the state of the art represented by those opened to public interest.

A. Special Purpose Datasets

The survey would be incomplete without a note on specialized datasets for distinct tasks in automated driving. Long solved is the traffic sign recognition problem [21] with dedicated datasets such as [22], [23]. Similar is the traffic light detection and interpretation [24]. The other datasets are more focused on mobile devices and augmented reality, therefore they do not cover the traffic from the perspective of the vehicle, but rather of a pedestrian as e.g. [15]. Another distinguishable group are the synthetic datasets such as SYNTHIA [25] or P.F.B. [26], which are specific due to precisely computed yet somewhat simplifying output data. The WildDash dataset [27], is focused on data, where the image segmentation algorithms fail. Contrary to its larger counterparts, this dataset is not meant for learning of the algorithms, but mainly for testing of their robustness.

B. Vision-focused datasets

Many datasets are designed to serve mostly for machine vision research, especially segmentation of real-life scenes and recognition of objects of interest. Large amount of very short recordings with a few manually annotated frames and crude GPS positioning is usually all what is available. The key contribution consists in large variety of scenes (often

Dataset	Cameras	LiDARs	GNSS
KITTI	4x 1.4 MP, 10 FPS	2000x64 pt, 10 Hz	RTK GPS
Málaga	2x 0.8 MP, 20 FPS	2x 1080 pt, 75 Hz 3x 1080 pt, 40 Hz	GPS
Oxford	3x 1.2 MP, 16 FPS 3x 1 MP, 11 FPS	2x 540 pt, 50 Hz 680x4 pt, 12.5 Hz	GPS
ApolloScape	4x 9.2 MP, 30 FPS	2x 1000000 pt/s	RTK GPS
Audi	6x 2.3 MP, 30 FPS	5x 1875x16 pt, 10 Hz	GPS
Waymo	3x 2.4 MP, 10 FPS 2x 2 MP, 10 FPS	5x 354000 pt/s	GPS

TABLE I

AN OVERVIEW OF THE DATASETS FOR GENERAL MAPPING.

acquired though crowd sourcing) and reliable reference data. The CityScapes dataset [28], Mapillary Vistas dataset [29] and the Berkeley DeepDrive dataset [30] are the main examples.

C. General mapping datasets

Arguably the most important group of datasets strongly focuses on sensory quality and variety, which allows large amount of applications in research. The vision subsystems are usually of higher quality and accompanied with laser scanners and inertial units. The most notable these days are the KITTI [31], the Málaga urban [19], the Oxford RobotCar [32], the ApolloScape [33], Audi [34] and Waymo [35] datasets summarized in Tab. I.

The size of the datasets varies a lot, the KITTI and Málaga recordings cover only tens of kilometers, while the Oxford and the ApolloScape sets map 1000+ kilometers of roads and the Audi and Waymo datasets focus on short segments from different environments. Treatment of timing and synchronization is unique in each case, the Oxford dataset timing has an excellent sub-millisecond precision, authors of the ApolloScape set claim just synchronization with no details explained and the KITTI and Málaga provide less precisely timed and mostly not synchronized data. The Audi dataset is precisely synchronized with GPS time and the Waymo dataset exhibits timestamping deviations in a millisecond scale.

D. Summary

The previous overview covers many datasets relevant to autonomous driving, but only a few are of the kind, which we are presenting. Removing the vision only datasets [28], [29], [30] and [27], we are left with six projects providing similar data as we do. The KITTI [31], Málaga [19] and Oxford [32] datasets offer lower quality sensory data, the Audi [34] and Waymo [35] are similar to ours, while the ApolloScape [33] is clearly the richest dataset available in sense of camera resolution and point cloud density. On the other hand, none of these six references contain a differential RTK GPS or thermal camera and, although synchronization and timestamping is mentioned everywhere, only the Oxford dataset presents detailed treatment of the topic using special software. We have achieved the same accuracy with more

Sensor	Type	Details	Freq.	Output data
4x RGB camera	DFK33-GX174	75° FoV 1920x1200 px	10	h265 video
Thermal camera	FLIR Tau 2	70° FoV 640x512 px	30	h265 video
2x LiDAR	Velodyne HDL-32e	32 beams ~2000 pts/turn	10	point cloud
GNSS receiver	Trimble BX982	RTK accuracy Direct heading	20 20	global pose time
IMU	Xsens MTi-G-710	Combined non-electrical quantities sensor	400 400 100 400 400 50	accelerometer gyroscope magnetometer temperature global pose time pressure

TABLE II
AN OVERVIEW OF THE SENSORS INSTALLED ON THE ATLAS
MEASUREMENT PLATFORM.

stable results due to hardware precautions similar to the Audi solution. Although we offer only raw data, taking into account current state of the art, the Brno urban dataset has features reaching beyond that margin.

III. THE DATA ACQUISITION PLATFORM

The dataset has been recorded using an extensible sensory platform called ATLAS built in our laboratory. The system is composed of a communication network, data processing computer, synchronization unit and sensors themselves. The whole apparatus was designed with precision and modularity in mind, allowing wide range of experimental setups, while maintaining quality of the recordings. The following paragraphs will go through the current state of the ATLAS platform (see Fig. 1) and the Conclusion will summarize the upcoming extensions.

A. Data Gathering Infrastructure

Because we deal with large throughput of raw data, the central control and recording computer is built on a CPU with 64 PCIe lanes and a NMVe SSD disk. A Nvidia GTX 1080Ti graphic card is present as well for video processing and compression. The backbone of the ATLAS recording framework is an Ethernet network. It is based on IP communication protocol and all the sensors and the acquisition PC are interconnected via a high-speed switch with up to 18 Gb/s bandwidth for connection to the acquisition PC. The only exception is the IMU which is connected to the PC through a virtual serial link via USB interface.

It should be noted, that neither the recording PC, nor the Ethernet network exhibit real-time capabilities for precise and reliable timing and this functionality is solved independently as will be discussed in Sec. III-D.

B. Sensory Equipment

Sensors are the clearly a crucial part of the system. Table II summarizes all the devices currently installed and their most important parameters. In the next paragraphs, we will briefly elaborate on each sensor category to better present the data we provide.

As seen in Sec. II, RGB cameras are a must in autonomous vehicle applications. The ATLAS platform has two cameras installed for stereo vision in the front and two lateral cameras with wider field of view (FoV) for better coverage of the crossroads, walkways and other road users passing around. The frontal cameras are installed wide apart (~70 cm) for better accuracy of distance estimation. With the full setup we cover more than 220° of the car's surroundings.

The next important sensors are LiDARs. We employ two 3D Velodyne scanners mounted with a slight tilt around the forward-pointing axis of the car as can be seen from the photograph in Fig. 1. The reason is twofold: first, the scanners better cover the area sideways of the car and second, the rays from both LiDARs cover the area in front of the car in a grid-like manner, which improves further object reconstruction possibilities.

Thermal camera usage is unique among existing datasets for automated driving. The device we employ senses infrared radiation in range of 7.5-13.5 μm , which corresponds to peak wavelengths emitted by objects in usual temperatures of -40 to +80 °C. The resolution is low in comparison to RGB cameras, but opacity of many objects (e.g. smoke, fog, thin foil) differs for infrared and visible light, so even with 640x512 pixels per frame, the information gain is substantial. The camera is mounted in the most important forward-looking direction.

The fourth kind of sensor mounted on the ATLAS platform is an inertial measurement unit. Besides accelerometers and gyroscopes, the device contains a combined GNSS receiver (GPS, GLONASS, Galileo, BeiDou) and provides additional environmental measurements such as temperature, atmospheric pressure and magnetic field, which have limited usage in automated driving, but we have decided to publish them as well for completeness.

Last but not least, there is a dedicated combined GNSS receiver to obtain the most precise global localization available. The RTK functionality allows centimetre level accuracy. Additionally, the receiver allows connection of two antennas at the same time allowing to directly obtain a heading vector in global coordinates. This feature was not present in any dataset we know about and can be very valuable as a reference in map building and localization applications. Of course, the receiver provides diagnostics of a reliability of the measurements as well as precise time for other sensors. This feature has a key role in our solution of synchronization and timestamping described in Sec. III-D.

C. Calibration

So far, we have spoken about sensor poses within the ATLAS platform in a somewhat vague manner. The reason comes from difficulties with exact measurements. Each sensor has its own frame of coordinates, whose origin mostly lies within the device and is usually tied with the chassis by a few dimensions with certain tolerance. Although mutual position could be acquired with decent accuracy, the orientation measurement is very sensitive and even a small error could result in faulty alignment of data from multiple

sources. Additionally, even if we could measure exactly, there are still manufacturing tolerances, which cannot be dealt with this way.

For this reason, we have decided to perform thorough calibration of sensors. Some of the methods used are suited for estimation of the intrinsic parameters of the device, while others are designed to obtain their mutual pose. Methods used, their settings, calibration data and the best estimates of the desired parameters are provided along the dataset on its website (<https://github.com/RoboticsBUT/Brno-Urban-Dataset-Calibrations>). We expect the sensory equipment to change over time and update this material accordingly.

Currently we employ the Kalibr tool [36] for RGB cameras and IMU mutual pose calibration and method [37] for RGB cameras and LiDARs mutual pose estimation. The thermal camera is then added to the model using method [8] from our laboratory. The last remaining devices are the antennas of the GNSS receiver, which turned out to be quite problematic to relate to the other sensors in some calibration procedure. The manufacturer fortunately provides precise location of their phase centre with respect to the chassis, so in this case we had to stick to manual measurements.

D. Synchronization and Timestamping

As already discussed in Sec. II, exact timing in data acquisition systems is, to some extent, replaceable by dedicated algorithms. We prefer to prevent the problems instead of fix them and the ATLAS platform was built with precise timing in mind. The scheme in Fig. 2 showing data flow in the system also plots the synchronization and time distribution lanes (dashed arrows).

The key source of precise time in our system is the GPS signal. Even the most basic receiver needs to maintain time precision in a nanosecond scale to provide usable positioning. Obviously, the Trimble RTK receiver has access to it, but the Velodyne laser scanners and Xsens IMU are equipped with small antennas as well, therefore precise time is available for them directly. For devices with no receiver of the GPS signal, we have designed a synchronization unit, which is clocked by a precise clock source from the Trimble receiver. It can either capture input trigger, pair it with an exact time and send a packet to the recording computer, or generate an output trigger signal with given frequency and send a timestamp corresponding to each firing. The unit is built around a simple microcontroller without operating system or nested interrupts, which allows to maintain transparent timing of its routines and guarantee an upper limit on timestamp-signal mismatch. Taking into account propagation delays in hardware etc., the error is well below 1 ms. With all that precautions, we can completely bypass the system time of the control computer and stamp the data with timing from a trustworthy source.

Currently we use the synchronization unit for triggering of all RGB cameras with a common signal. The thermal camera unfortunately can not be directly triggered by an

Fig. 2. Time synchronization and data-flow diagram of the ATLAS platform.

external signal, neither provides an output trigger, but contains precise clock allowing additional corrections. Resulting list of timestamps is then paired with a timestamp sequence from the synchronization unit providing interpolated timing of each frame.

An obvious drawback of this approach is a strong dependence on the GPS signal availability. For this reason, all devices drawing the GPS time employ a graceful fallback to local clock source with known accuracy. In the worst case scenario the timing errors break the 1 ms limit after tens of minutes without GPS signal, which is enough to pass tunnels and other problematic locations and regain the exact time. We take a great care not to exceed this limit in any of our recordings.

IV. DATASET

As mentioned above, the dataset was recorded in Brno, Czech Republic. The ATLAS platform is not yet entirely waterproof, so the range of weather conditions captured is limited. On the other hand, thanks to mid-size of the town, the environmental diversity spans from natural, countryside-looking locations to city center with historical buildings, public transportation and especially large amount of traffic and pedestrians.

A. Content

It is a good practice to sort the data according to its content. The time of recording serves mostly as a unique identifier and a brief description is good to get a quick overview of the recording, but both are cumbersome to use, if a whole database of all recordings is needed to be searched through. For this reason, we have employed a system of tags, which allow us to highlight the most important content and enable easy filtration of the recordings summarized in Tab. III.

So far, we have made available 67 recordings of total length of 375.7 km and duration of more than 10 hours.

Fig. 3. Example of Brno Dataset data visualization. In the first six rows there are the RGB camera images, in the seventh and eighth row there are the thermal (IR) camera images, and the LiDAR data visualizations are on the bottom.

Tag category	Tag	Recordings	Distance	Duration
Weather	Sunny	42	245.1	6:23
	Partly-cloudy	25	130.6	4:32
Daytime	Morning	15	60.1	1:48
	Noon	26	175.6	4:04
	Afternoon	21	96.4	3:37
	Evening	5	43.6	1:24
Environment	City	36	181.9	5:56
	Suburb	21	71.0	2:34
	Country	6	48.1	1:16
	Highway	4	74.7	1:08

TABLE III

STATISTICS OF THE DATASET CONTENT. LENGTH IS MEASURED IN KILOMETERS AND DURATION IN HOURS:MINUTES FORMAT.

Fig. 4. Map visualization of the roads driven through during the dataset recording. Green - RTK positioning + heading, Red - non RTK positioning or missing the heading data. Brno coordinates: 49.2002211N, 16.6078411E.

B. Data Structure

The structure of a single recording follows the scheme in Fig. 5. RGB and thermal camera data are distributed as H265 video and the LiDAR scans are compressed into .zip archives to reduce their size as much as possible. Other data such as timestamps and calibration files occupy a negligible amount of memory and are stored in human readable .txt and .yaml files.

C. Software and Development Tools

The recording session runs fully on the Robot Operation System (ROS). This allows us to create highly scalable solution which is compatible with many other projects using ROS backend as a base line for development of robotic applications. To satisfy a wider audience and for more comfortable usage, we publish the data in an easily readable raw format and provide a script for conversion into the ROS bag.

We also provide a set of several Python-OpenCV based scripts that help to process the video data into separated files for each frame, to draw the trajectory into the Google Maps, etc. The software is available from <https://github.com/RoboticsBUT/Brno-Urban-Dataset-Tools>, or as a submodule of the dataset git repository in the `tools/` folder.

Fig. 5. Data structure of a recording with separated folders for every sensor and common data.

V. CONCLUSION AND FUTURE WORK

In this paper, we have presented a new dataset for autonomous driving research recorded in Brno - Czech Republic. We provide state of the art sensory measurements with three key additions exceeding other datasets available. First, most routes where GNSS signal was available are accompanied with centimeter accurate global position and heading from differential GNSS receiver. Second important feature is synchronization and timestamping of the data with sub-millisecond precision allowing simpler data fusion and evaluation of the algorithms, where temporal shifts in measurements are part of the stochastic model. The last addition is the infrared camera significantly increasing detection and recognition capabilities of the ATLAS measurement platform.

At the time of writing of this paper, 67 recordings with total length of 375.7 km and duration of more than 10 hours are available. The recordings are tagged with respect to the environment, weather and other events encountered and provided in an easily usable format through the project page: <https://github.com/RoboticsBUT/Brno-Urban-Dataset>. Data collection is a long term process and we expect the dataset to grow over time with various new recordings required by our research. We plan to densely cover a smaller region in Brno for map-building applications and to obtain more data acquired with problematic lightning and weather conditions or the winter sessions. We are also considering waterproofing of the ATLAS platform to cover full range of weather conditions encountered in central Europe.

REFERENCES

- [1] C. Badue, R. Guidolini, R. V. Carneiro, P. Azevedo, V. B. Cardoso, A. Forechi, L. F. R. Jesus, R. F. Berriel, T. M. Paixão, F. Mutz, T. Oliveira-Santos, and A. F. De Souza, "Self-Driving Cars: A Survey," *ArXiv*, jan 2019.
- [2] A. Faisal, T. Yigitcanlar, M. Kamruzzaman, and G. Currie, "Understanding autonomous vehicles: A systematic literature review on capability, impact, planning and policy," *Journal of Transport and Land Use*, vol. 12, no. 1, pp. 45–72, jan 2019.
- [3] T. Litman, "Autonomous vehicle implementation predictions : implications for transport planning Implications for Transport Planning," Victoria Transport Policy Institute, Tech. Rep., 2019.
- [4] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, dec 2016.
- [5] E. C. McKiernan, P. E. Bourne, C. T. Brown, S. Buck, A. Kenall, J. Lin, D. McDougall, B. A. Nosek, K. Ram, C. K. Soderberg, J. R. Spies, K. Thaney, A. Updegrove, K. H. Woo, and T. Yarkoni, "How open science helps researchers succeed," *eLife*, vol. 5, pp. 1–19, jul 2016.
- [6] H. Yin and C. Berger, "When to use what data set for your self-driving car algorithm: An overview of publicly available driving datasets," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, vol. 2018-March. IEEE, oct 2017, pp. 1–8.
- [7] M. Bieshaar, S. Zernetsch, G. Reitberger, B. Sick, E. Fuchs, and K. Doll, "Detecting Intentions of Vulnerable Road Users Based on Collective Intelligence," in *AAET 2017 Automatisiertes und vernetztes Fahren*, Braunschweig, Germany, 2017.
- [8] L. Zalud and P. Kocmanova, "Fusion of thermal imaging and CCD camera-based data for stereovision visual telepresence," in *2013 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, vol. 2. IEEE, oct 2013, pp. 1–6.
- [9] J. Schneider, C. Eling, L. Klingbeil, H. Kuhlmann, W. Forstner, and C. Stachniss, "Fast and effective online pose estimation and mapping for UAVs," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2016-June. IEEE, may 2016, pp. 4784–4791.
- [10] G. Schall, D. Wagner, G. Reitmayr, E. Taichmann, M. Wieser, D. Schmalstieg, and B. Hofmann-Wellenhof, "Global pose estimation using multi-sensor fusion for outdoor Augmented Reality," in *2009 8th IEEE International Symposium on Mixed and Augmented Reality*. IEEE, oct 2009, pp. 153–162.
- [11] C. Albl, Z. Kukulova, A. Fitzgibbon, J. Heller, M. Smid, and T. Pajdla, "On the Two-View Geometry of Unsynchronized Cameras," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2017-Janua. IEEE, jul 2017, pp. 5593–5602.
- [12] R. Mhirri, P. Vasseur, S. Mousset, R. Bouteau, and A. Bensrhair, "Visual odometry with unsynchronized multi-cameras setup for intelligent vehicle application," in *2014 IEEE Intelligent Vehicles Symposium Proceedings*. IEEE, jun 2014, pp. 1339–1344.
- [13] J. Levinson and S. Thrun, "Automatic Online Calibration of Cameras and Lasers," in *Robotics: Science and Systems IX*. Robotics: Science and Systems Foundation, jun 2013.
- [14] C. Le Gentil, T. Vidal-Calleja, and S. Huang, "3D Lidar-IMU Calibration Based on Upsampled Preintegrated Measurements for Motion Distortion Correction," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2018, pp. 2149–2155.
- [15] S. Cortés, A. Solin, E. Rahtu, and J. Kannala, "ADVIO: An Authentic Dataset for Visual-Inertial Odometry," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2018, vol. 11214 LNCS, pp. 425–440.
- [16] L. Wang, Y. Zhang, and J. Wang, "Map-Based Localization Method for Autonomous Vehicles Using 3D-LIDAR," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 276–281, jul 2017.
- [17] R. W. Wolcott and R. M. Eustice, "Visual localization within LIDAR maps for automated urban driving," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 33, no. 5. IEEE, sep 2014, pp. 176–183.
- [18] H. Deilamsalehy and T. C. Havens, "Sensor fused three-dimensional localization using IMU, camera and LiDAR," in *2016 IEEE SENSORS*. IEEE, oct 2016, pp. 1–3.
- [19] J. L. Blanco-Claraco, F. Á. Moreno-Dueñas, and J. González-Jiménez, "The Málaga urban dataset: High-rate stereo and LiDAR in a realistic urban scenario," *International Journal of Robotics Research*, vol. 33, no. 2, pp. 207–214, 2014.
- [20] S. Wang, M. Bai, G. Mattyus, H. Chu, W. Luo, B. Yang, J. Liang, J. Cheverie, S. Fidler, and R. Urtasun, "TorontoCity: Seeing the World with a Million Eyes," in *2017 IEEE International Conference on Computer Vision (ICCV)*, vol. 2017-Octob. IEEE, oct 2017, pp. 3028–3036.
- [21] A. Mogelmose, M. M. Trivedi, and T. B. Moeslund, "Vision-Based Traffic Sign Detection and Analysis for Intelligent Driver Assistance Systems: Perspectives and Survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 4, pp. 1484–1497, dec 2012.
- [22] J. Stalkamp, M. Schlipsing, J. Salmen, and C. Igel, "Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition," *Neural Networks*, vol. 32, pp. 323–332, aug 2012.
- [23] R. Timofte, K. Zimmermann, and L. Van Gool, "Multi-view traffic sign detection, recognition, and 3D localisation," *Machine Vision and Applications*, vol. 25, no. 3, pp. 633–647, apr 2014.
- [24] A. Fregin, J. Muller, U. Krebel, and K. Diermayer, "The DriveU Traffic Light Dataset: Introduction and Comparison with Existing Datasets," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 3376–3383, 2018.
- [25] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez, "The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2016-December, no. 600388. IEEE, jun 2016, pp. 3234–3243.
- [26] S. R. Richter, Z. Hayder, and V. Koltun, "Playing for Benchmarks," in *2017 IEEE International Conference on Computer Vision (ICCV)*, vol. 2017-Octob. IEEE, oct 2017, pp. 2232–2241.
- [27] O. Zendel, K. Honauer, M. Murschitz, D. Steininger, and G. F. Domínguez, "WildDash - Creating Hazard-Aware Benchmarks," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2018, vol. 11210 LNCS, pp. 407–421.
- [28] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The Cityscapes Dataset for Semantic Urban Scene Understanding," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2016-December. IEEE, jun 2016, pp. 3213–3223.
- [29] G. Neuhold, T. Ollmann, S. R. Bulò, and P. Kontschieder, "The Mapillary Vistas Dataset for Semantic Understanding of Street Scenes," in *2017 IEEE International Conference on Computer Vision (ICCV)*, vol. 2017-Octob. IEEE, oct 2017, pp. 5000–5009.
- [30] F. Yu, W. Xian, Y. Chen, F. Liu, M. Liao, V. Madhavan, and T. Darrell, "BDD100K: A Diverse Driving Video Database with Scalable Annotation Tooling," *ArXiv*, pp. 1–16, may 2018.
- [31] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, sep 2013.
- [32] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The Oxford RobotCar dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, jan 2017.
- [33] X. Huang, P. Wang, X. Cheng, D. Zhou, Q. Geng, and R. Yang, "The ApolloScape Open Dataset for Autonomous Driving and its Application," *ArXiv*, pp. 1–17, mar 2018.
- [34] J. Geyer, Y. Kassahun, M. Mahmudi, X. Ricou, R. Durgesh, A. S. Chung, L. Hauswald, V. H. Pham, M. Mühlegg, S. Dorn, T. Fernandez, M. Jänicke, S. Mirashi, C. Savani, M. Sturm, O. Vorobiov, M. Oelker, S. Garreis, and P. Schuberth, "A2D2: AEV Autonomous Driving Dataset," 2019.
- [35] P. Sun, H. Kretschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov, "Scalability in Perception for Autonomous Driving: Waymo Open Dataset," dec 2019.
- [36] P. Furgale, C. H. Tong, T. D. Barfoot, and G. Sibley, "Continuous-time batch trajectory estimation using temporal basis functions," *International Journal of Robotics Research*, vol. 34, no. 14, pp. 1688–1710, 2015.
- [37] A. Geiger, F. Moosmann, O. Car, and B. Schuster, "Automatic camera and range sensor calibration using a single shot," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, may 2012, pp. 3936–3943.