DDOS: The Drone Depth and Obstacle Segmentation Dataset
Supplementary Material

Benedikt Kolbeinsson
Imperial College London
bk915@imperial.ac.uk

Krystian Mikolajczyk
Imperial College London
k.mikolajczyk@imperial.ac.uk

https://huggingface.co/datasets/benediktkol/DDOS

A. Datasheet

In light of the growing recognition of the pivotal role that datasets play in shaping the behavior and outcomes of machine learning models, this section adheres to the framework proposed in the *Datasheets for Datasets* paper [14]. Acknowledging the potential consequences of mismatches between training or evaluation datasets and real-world deployment contexts, as well as the risk of perpetuating societal biases within machine learning models, we embrace the call for increased transparency and accountability in documenting the provenance, creation, and use of machine learning datasets [39]. By adopting this standardized reporting scheme, we aim to provide a comprehensive understanding of our dataset's motivation, composition, collection process, and recommended uses. This adherence to the datasheets for datasets framework aligns with the broader objective of enhancing transparency, mitigating biases, fostering reproducibility, and aiding researchers and practitioners in selecting datasets tailored to their specific tasks. In the following subsections, we systematically address the key questions outlined in the datasheets for datasets, providing a thorough account of our dataset's characteristics and attributes.

A.1. Motivation

For what purpose was the dataset created? The Drone Depth and Obstacle Segmentation (DDOS) dataset, was created to address the limitations posed by the scarcity of annotated aerial datasets, specifically for training and evaluating models in depth and semantic segmentation tasks. The primary objective is to focus on the detection and segmentation of thin structures like wires, cables, and fences in aerial views, which are critical for ensuring the safe operation of drones. The dataset aims to fill the gap in existing datasets that predominantly concentrate on common structures and lack representation of fine spatial characteristics of thin structures.

Who created the dataset? The dataset was created by Benedikt Kolbeinsson and Krystian Mikolajczyk.

A.2. Composition

What do the instances that comprise the dataset represent? The instances in the dataset represent individual drone flights which are composed of sequences of observations (images, depth maps, segmentation, etc.) captured during each flight.

How many instances are there in total? The dataset consists of a total of 340 drone flights, and each flight comprises 100 sequential observations. Therefore, there are a total of 34 000 observations (340 flights $\times$ 100 observations per flight).

Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set? No, there exists many more possible flight paths in the environments used as well as in other environments.

What data does each instance consist of? Each flight consists of 100 sequential observations, comprising of a high-resolution image captured by a monocular camera affixed to the front of the drone, corresponding depth maps, pixel-level object segmentation masks, optical flow information and surface normals. As well as coordinates, pose and speed information and environment information including weather. All image modalities maintain a resolution of 1280×720, and the depth maps cover a range from 0 to 100 m.

Is there a label or target associated with each instance? Yes, DDOS features pixel-wise object segmentation masks with ten distinct classes, allowing for detailed analysis of diverse obstacles and environmental elements. These classes are: *ultra thin*, *thin*, *small mesh*, *large mesh*, *trees*, *buildings*, *vehicles*, *animals*, *other*, and *background*. For instance, the *ultra thin* class covers objects like wires and cables, while the *thin* class encompasses streetlights and poles. The *small mesh* class includes objects like fences and nets, and the *large mesh* class involves structures similar to pylons and radio masts. In addition, corresponding depth maps, optical flow information and surface normals are included.

Is any information missing from individual instances? No.

Are relationships between individual instances made explicit? Yes, the flight coordinates are available.

Are there recommended data splits? Yes, the dataset is partitioned into training, validation, and testing subsets, encompassing 300, 20, and 20 flights, respectively.

Are there any errors, sources of noise, or redundancies in the dataset? The data is simulated and no artificial noise is added.

Is the dataset self-contained, or does it link to or otherwise rely on external resources? Yes, DDOS is self-contained.

Does the dataset contain data that might be considered confidential? No.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety? No.

A.3. Collection Process

How was the data associated with each instance acquired? The data was acquired through simulated drone flights using AirSim [32], a drone simulator.

What mechanisms or procedures were used to collect the data? DDOS was generated using AirSim and data was saved using built-in APIs.

If the dataset is a sample from a larger set, what was the sampling strategy? During the simulation process, flights with severe crashes were discarded.

Who was involved in the data collection process? Data collection scripts were written by Benedikt Kolbeinsson.

Over what timeframe was the data collected? The simulation process took two days.

Were any ethical review processes conducted? No.

A.4. Preprocessing / cleaning / labeling

Was any preprocessing / cleaning / labeling of the data done? During the simulation, labels such as depth and semantic segmentation are automatically recorded. Flights with severe crashes were discarded.

Was the "raw" data saved in addition to the preprocessed / cleaned / labeled data? The processed data is a lossless function of the raw data. The only removed data are flights with severe crashes and are not saved.

Is the software that was used to preprocess / clean / label the data available? Yes, AirSim is open source.

A.5. Uses

What (other) tasks could the dataset be used for? DDOS is valuable for training and evaluating algorithms related to obstacle and object segmentation, depth estimation, and drone navigation.

Is there anything about the composition of the dataset or the way it was collected and preprocessed / cleaned / labeled that might impact future uses? No.

Are there tasks for which the dataset should not be used? Yes, DDOS should not be used for malicious purposes.

A.6. Distribution

Will the dataset be distributed to third parties outside of the entity on behalf of which the dataset was created? Yes, DDOS is hosted on Hugging Face and is available at: huggingface.co/datasets/benediktkol/DDOS

How will the dataset be distributed? DDOS is openly available on Hugging Face: huggingface.co/datasets/benediktkol/DDOS

Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)? Yes, DDOS is openly licensed under CC BY-NC 4.0.

Have any third parties imposed IP-based or other restrictions on the data associated with the instances? No.

Do any export controls or other regulatory restrictions apply to the dataset or to individual instances? No.

Image Depth Segmentation

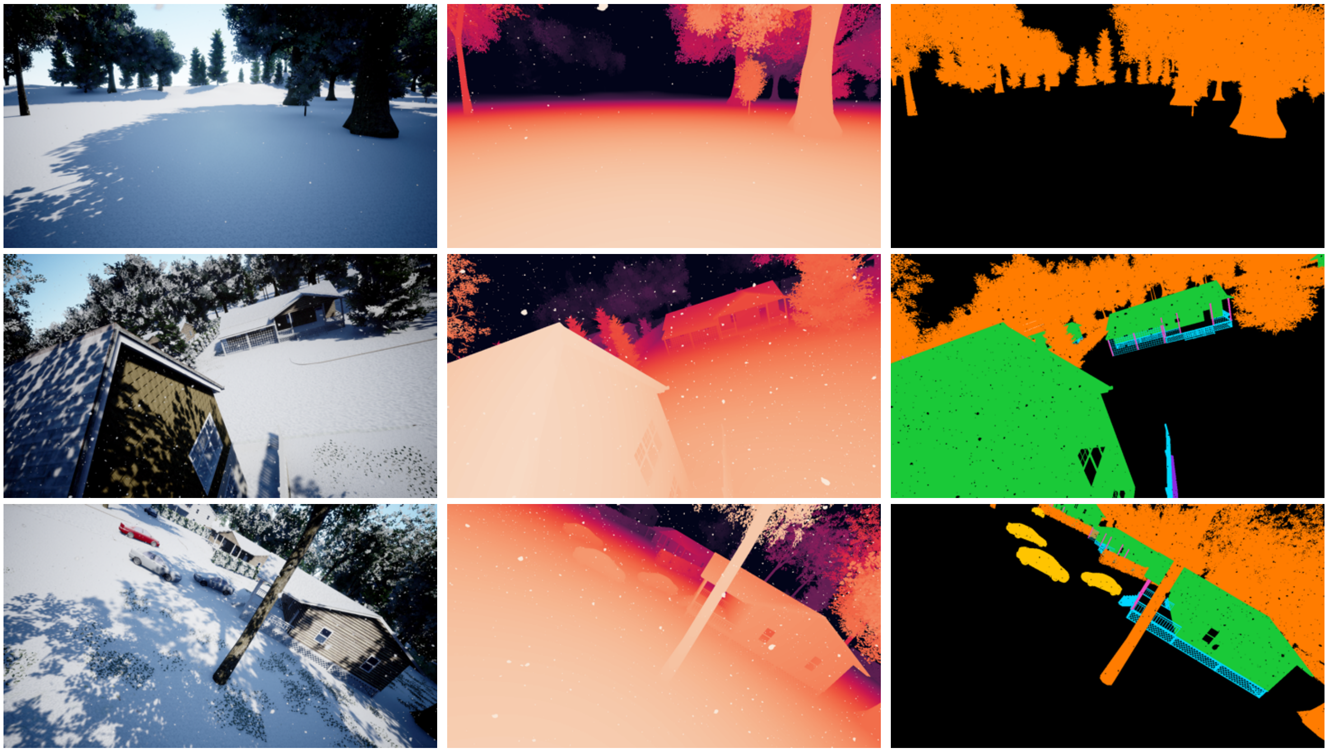

Figure 1. **Low altitude examples from DDOS.** The DDOS dataset encompasses flights featuring diverse flight characteristics, including examples of low altitude maneuvers and aggressive turns under snowy conditions.

A.7. Maintenance

Who will be supporting / hosting / maintaining the dataset? DDOS is hosted on Hugging Face

How can the owner / curator / manager of the dataset be contacted? Contact can be made on Hugging Face: huggingface.co/datasets/benediktkol/DDOS

Will the dataset be updated? There is no current plan to augment the dataset.

Will older versions of the dataset continue to be supported / hosted / maintained? Yes.

If others want to extend / augment / build on / contribute to the dataset, is there a mechanism for them to do so? There is no specific mechanism for others to extend / augment / build on / contribute to the dataset.

B. Additional Examples

In this section, we present further examples from the DDOS dataset, as illustrated in Figures 1 and 2. These examples are specifically selected to highlight the dataset's diversity and the intricate details captured within. For clarity and emphasis on these finer aspects, the visualizations are confined to the RGB images, accompanied by their respective depth maps and semantic segmentations. Notably, Figure 2 offers a glimpse into the diverse perspectives encompassed within DDOS. Conversely, Figure 1 is dedicated to showcasing scenarios captured during low altitude flights in snowy conditions, underscoring the dataset's versatility and the challenging environments it encompasses.

DDOS, serves as a comprehensive aerial resource for the research community, particularly in the domains of depth estimation and segmentation. Its utility is especially evident in scenarios involving aerial perspectives, as encountered by drones, offering valuable insights for discerning thin structures within the visual field.

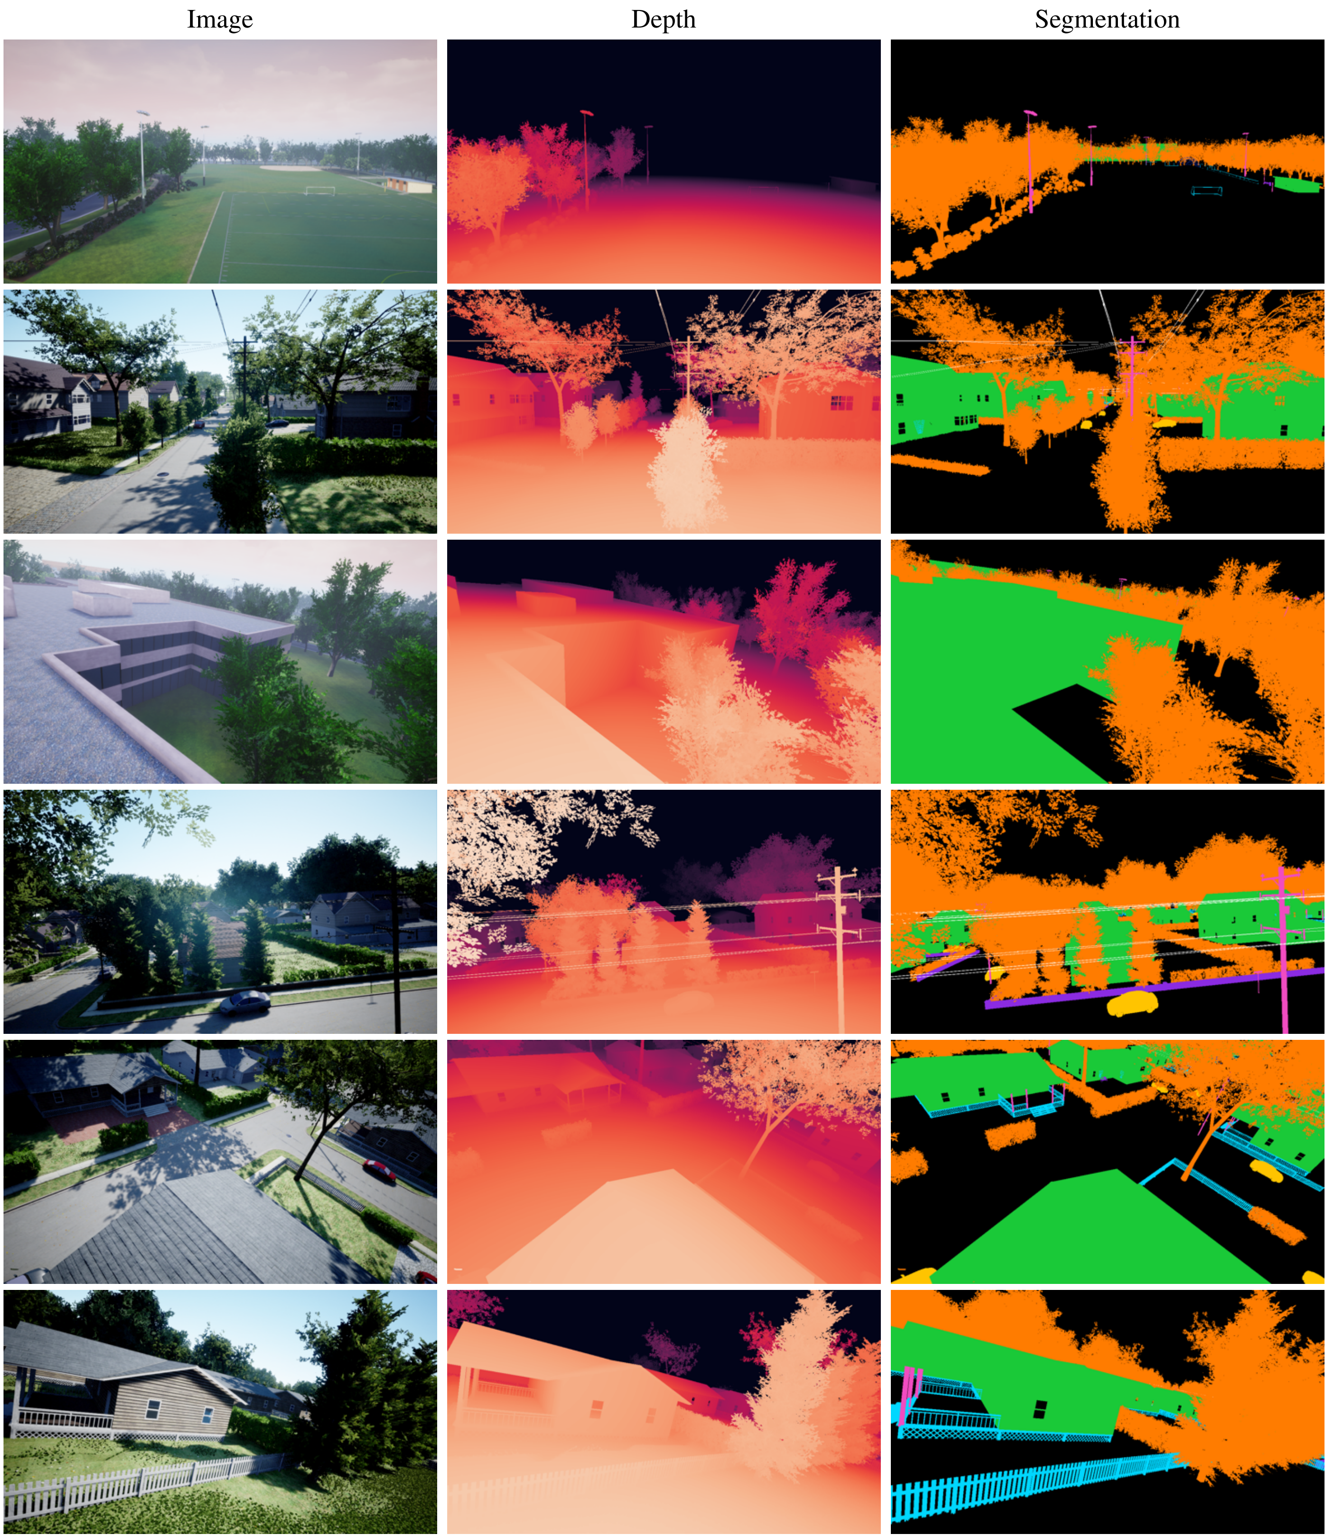

Figure 2. **Diverse perspectives in DDOS.** This selection highlights various aerial views from the DDOS dataset, with each frame presenting an RGB image, its depth map, and semantic segmentation. The imagery captures a range of features, from varied vegetation to complex architectural structures. Optical flow and surface normals, while part of the dataset, are not included in this visualization. Viewers are advised to examine these images digitally.

References

[1] Rabab Abdelfattah, Xiaofeng Wang, and Song Wang. Ttpla: An aerial-image dataset for detection and segmentation of transmission towers and power lines. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 8

[2] Stuart M Adams and Carol J Friedland. A survey of unmanned aerial vehicle (uav) usage for imagery collection in disaster research and management. In *9th international workshop on remote sensing for disaster response*, pages 1–8, 2011. 7

[3] Babankumar Bansod, Rangoli Singh, Ritula Thakur, and Gaurav Singhal. A comparision between satellite based and drone based remote sensing technology to achieve sustainable development: A review. *Journal of Agriculture and Environment for International Development (JAEID)*, 111(2): 383–407, 2017. 7

[4] Taha Benarbia and Kyandoghere Kyamakya. A literature review of drone-based package delivery logistics systems and their implementation feasibility. *Sustainability*, 14(1):360, 2021. 7

[5] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021. 13

[6] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020. 8

[7] Joshua Candamo, Rangachar Kasturi, Dmitry Goldgof, and Sudeep Sarkar. Detection of Thin Lines using Low-Quality Video from Low-Altitude Aircraft in Urban Settings. *IEEE Transactions on Aerospace and Electronic Systems*, 45(3): 937–949, 2009. 8

[8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 8

[9] Sharifah Mastura Syed Mohd Daud, Mohd Yusmiaidil Putera Mohd Yusof, Chong Chin Heo, Lay See Khoo, Mansharan Kaur Chainchel Singh, Mohd Shah Mahmood, and Hapizah Nawawi. Applications of drone in disaster management: A scoping review. *Science & Justice*, 62(1):30–42, 2022. 7

[10] Milan Erdelj, Enrico Natalizio, Kaushik R. Chowdhury, and Ian F. Akyildiz. Help from the Sky: Leveraging UAVs for Disaster Management. *IEEE Pervasive Computing*, 16(1): 24–32, 2017.

[11] Mario Arturo Ruiz Estrada and Abrahim Ndoma. The uses of unmanned aerial vehicles–uav's-(or drones) in social logistic: Natural disasters response and humanitarian relief aid. *Procedia Computer Science*, 149:375–383, 2019. 7

[12] Michael Fonder and Marc Van Droogenbroeck. Mid-air: A multi-modal dataset for extremely low altitude drone flights. In *Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2019. 8, 9

[13] Vipul Garg, Suman Niranjan, Victor Prybutok, Terrance Pohlen, and David Gligor. Drones in last-mile delivery: A systematic review on efficiency, accessibility, and sustainability. *Transportation Research Part D: Transport and Environment*, 123:103831, 2023. 7

[14] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021. 1

[15] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 8

[16] Yoshio Inoue. Satellite-and drone-based remote sensing of crops and soils for smart farming–a review. *Soil Science and Plant Nutrition*, 66(6):798–810, 2020. 7

[17] James R. Kellner, John Armston, Markus Birrer, K. C. Cushman, Laura Duncanson, Christoph Eck, Christoph Falleger, Benedikt Imbach, Kamil Král, Martin Krůček, Jan Trochta, Tomáš Vrška, and Carlo Zgraggen. New opportunities for forest remote sensing through ultra-high-density drone lidar. *Surveys in Geophysics*, 40:959–977, 2019. 7

[18] Zhenyu Li, Zehui Chen, Ang Li, Liangji Fang, Qinhong Jiang, Xianming Liu, Junjun Jiang, Bolei Zhou, and Hang Zhao. Simipu: Simple 2d image and 3d point cloud unsupervised pre-training for spatial-aware visual representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1500–1508, 2022. 13, 14

[19] Zhenyu Li, Xuyang Wang, Xianming Liu, and Junjun Jiang. Binsformer: Revisiting adaptive bins for monocular depth estimation. *arXiv preprint arXiv:2204.00987*, 2022. 13, 14

[20] Zhenyu Li, Zehui Chen, Xianming Liu, and Junjun Jiang. Depthformer: Exploiting long-range correlation and local information for accurate monocular depth estimation. *Machine Intelligence Research*, pages 1–18, 2023. 13, 14

[21] Ye Lyu, George Vosselman, Gui-Song Xia, Alper Yilmaz, and Michael Ying Yang. Uavid: A semantic segmentation dataset for uav imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 165:108–119, 2020. 8

[22] Ratnesh Madaan, Daniel Maturana, and Sebastian Scherer. Wire detection using synthetic data and dilated convolutional networks for unmanned aerial vehicles. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3487–3494. IEEE, 2017. 8, 9

[23] Alina Marcu, Vlad Licaret, Dragos Costea, and Marius Leordeanu. Semantics through time: Semi-supervised segmentation of aerial videos with iterative label propagation. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 8

[24] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3061–3070, 2015. 8

[25] Payal Mittal, Raman Singh, and Akashdeep Sharma. Deep learning-based object detection in low-altitude uav datasets: A survey. *Image and Vision computing*, 104:104046, 2020. 7

[26] Norzailawati Mohd Noor, Alias Abdullah, and Mazlan Hashim. Remote sensing uav/drones and its applications for urban areas: A review. In *IOP conference series: Earth and environmental science*, page 012003. IOP Publishing, 2018. 7

[27] Ishan Nigam, Chen Huang, and Deva Ramanan. Ensemble knowledge transfer for semantic segmentation. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1499–1508. IEEE, 2018. 8

[28] Yalong Pi, Nipun D Nath, and Amir H Behzadan. Convolutional neural networks for object detection in aerial imagery for disaster response and recovery. *Advanced Engineering Informatics*, 43:101009, 2020. 7

[29] Chengyi Qu, Francesco Betti Sorbelli, Rounak Singh, Prasad Calyam, and Sajal K Das. Environmentally-aware and energy-efficient multi-drone coordination and networking for disaster response. *IEEE Transactions on Network and Service Management*, 2023. 7

[30] Giulia Rizzoli, Francesco Barbato, Matteo Caligiuri, and Pietro Zanuttigh. Syndrone-multi-modal uav dataset for urban scenarios. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2210–2220, 2023. 8, 9

[31] Aamina Shah, Komali Kantamaneni, Shirish Ravan, and Luiza C Campos. A systematic review investigating the use of earth observation for the assistance of water, sanitation and hygiene in disaster response and recovery. *Sustainability*, 15(4):3290, 2023. 7

[32] Shital Shah, Debadeepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics*, 2017. 2, 9

[33] Adam Stambler, Gary Sherwin, and Patrick Rowe. Detection and Reconstruction of Wires Using Cameras for Aircraft Safety Systems. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 697–703, 2019. ISSN: 1050-4729. 8

[34] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 8

[35] Lina Tang and Guofan Shao. Drone remote sensing for forestry research and practices. *Journal of Forestry Research*, 26(4):791–797, 2015. 7

[36] Ashley Varghese, Jayavardhana Gubbi, Hrishikesh Sharma, and P Balamuralidhar. Power infrastructure monitoring and damage detection using drone captured images. In *2017 international joint conference on neural networks (IJCNN)*, pages 1681–1687. IEEE, 2017. 8

[37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 13

[38] Wenshan Wang, Delong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. Tartanair: A dataset to push the limits of visual slam. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4909–4916. IEEE, 2020. 8, 9

[39] World Economic Forum Global Future Council on Human Rights 2016–2018. How to prevent discriminatory outcomes in machine learning. https://www.weforum.org/whitepapers/how-to-prevent-discriminatory-outcomes-in-machine-learning, 2018. 1