

# Non-Uniform Spatial Alignment Errors in sUAS Imagery From Wide-Area Disasters

Thomas Manzini<sup>†\*</sup>, Priyankari Perali<sup>†\*</sup>, Raisa Karnik\*, Mihir Godbole\*, Hasnat Abdullah\*, Robin Murphy\*

**Abstract**—This work presents the first quantitative study of alignment errors between small uncrewed aerial systems (sUAS) geospatial imagery and a priori building polygons and finds that alignment errors are non-uniform and irregular. The work also introduces a publicly available dataset of imagery, building polygons, and human-generated and curated adjustments that can be used to evaluate existing strategies for aligning building polygons with sUAS imagery. There are no efforts that have aligned pre-existing spatial data with sUAS imagery, and thus, there is no clear state of practice. However, this effort and analysis show that the translational alignment errors present in this type of data, averaging 82px and an intersection over the union of 0.65, which would induce further errors and biases in downstream machine learning systems unless addressed. This study identifies and analyzes the translational alignment errors of 21,619 building polygons in fifty-one orthomosaic images, covering 16787.2 Acres (26.23 square miles), constructed from sUAS raw imagery from nine wide-area disasters (Hurricane Ian, Hurricane Harvey, Hurricane Michael, Hurricane Ida, Hurricane Idalia, Hurricane Laura, the Mayfield Tornado, the Musset Bayou Fire, and the Kilauea Eruption). The analysis finds no uniformity among the angle and distance metrics of the building polygon alignments as they present an average degree variance of 0.4 and an average pixel distance variance of 0.45. This work alerts the sUAS community to the problem of spatial alignment and that a simple linear transform, often used to align satellite imagery, will not be sufficient to align spatial data in sUAS orthomosaic imagery.

## I. INTRODUCTION

Following a major disaster, teams managing the response will align on a selection of building footprint polygons to represent the structures in the impacted areas. These building polygons become a shared reference for disaster operations in the impacted area. Following the selection of pre-disaster building polygons, and in some cases in parallel, additional post-disaster imagery is gathered by satellite, manned aviation assets, and small uncrewed aerial systems (sUAS) sources. This post-disaster imagery is then used to determine the conditions in the disaster area, and in particular, the conditions of the buildings represented by the polygons that have been decided upon. However, when this post-disaster imagery arrives, alignment errors are present. An example of one such alignment error is presented in Figure 1.

These alignment errors impede systems that rely on sUAS imagery. For example, machine learning systems that are trained with sUAS imagery to automate building damage assessment in disaster response will incorrectly classify



Fig. 1: A building footprint polygon (green) out of alignment with the building’s structure (image). The adjustment to align the polygon to the building’s structure shown in blue.

incorrect damage labels to buildings due to misalignment. Further, robot systems that leverage similar spatial data for navigation will have to operate with increased risk.

In order to build systems that can integrate closely with response efforts, it is important that alignment errors between geospatial sUAS imagery and existing geospatial data, like building polygons, is understood. With this in mind, this work presents the first analysis of spatial alignment errors between sUAS geospatial imagery and building polygons by studying 51 orthomosaics collected from sUAS at 9 wide-area disasters in the United States. This is done by quantitatively measuring the alignment errors of 21,619 building polygons imaged by sUAS. Alignment errors were measured through human-generated “adjustments,” a line segment that describes the angle and distance by which a building polygon needs to be translated in order for that polygon to align with the imagery, as shown in Figure 1. The adjustment terminology is motivated by the polar coordinate system.

The contributions of this work are as follows.

- The first explicit measurement of the distribution of adjustments necessary to align building polygons with sUAS orthomosaic imagery.
- The first dataset of ground truth adjustments that can be used to train or evaluate alignment techniques in sUAS orthomosaic imagery.

The remaining paper discusses the need for sUAS imagery

<sup>†</sup> Indicates that the authors made equal contributions to this work.

\*All authors are with the Department of Computer Science and Engineering, Texas A&M University. Contact: {tmanzini, perali, raisak, amigo2000, hasnat.md.abdullah, robin.r.murphy}@tamu.edu

<sup>1</sup>All imagery, building polygons, and adjustments will be released upon publication at WEBSITE FORTHCOMING




Fig. 2: 5 orthomosaics that were evaluated in this analysis from 5 of the 9 disasters considered. Starting in the top left and proceeding clockwise, the imagery was collected from: Hurricane Ian (2022), Kilauea Eruption (2018), The Musset Bayou Fire (2020), Hurricane Idalia (2023), and Hurricane Ida (2021). Pixel and spatial scales vary across imagery.

building polygon alignments and reviews previous work (Section II), provides an overview of the analysis approach (Section III), presents the analysis results (Section IV), discusses the analysis results and implications (Section V), and concludes with summary remarks (Section VI).

## II. PREVIOUS AND RELATED WORK

In order to evaluate the previous and related work and understand the motivation for this paper, the importance of the alignment problem and the sources of alignment are explained first, followed by a review of the literature. The alignment problem in sUAS orthomosaic imagery is similar to research in improving photogrammetry accuracy, however those methods assume ground control points which are not always available during disasters.

Before discussing the literature, the terminology must be established. Different communities may refer to this work as “registration” [24], [9], “alignment” [14], [25], [26], “image shifting” [11], or a study of “reconstruction accuracy,” “spatial accuracy,” or “positional precision” [5], [7], [15], [16], [4], [12]. Due to the sparse and non-visual nature of building polygon data and the way it is being utilized in this work, “alignment” represents the most appropriate term.

### A. Why is addressing alignment important?

The robotics and machine learning (ML) literature to date has not explicitly considered the issue of aligning predefined spatial data with sUAS imagery, instead focusing on semantic segmentation and object recognition [1], [22], [28]. This creates friction when working with existing spatial imagery because of the following three reasons. First, if sUAS and ML systems are expected to provide operational benefit to a wide-area disaster response, integration with the existing workflows, and thus predefined building polygons, is crucial. Second, ML systems developed for use in geospatial imagery can lose substantial performance when faced with alignment errors. A 2-meter translation error during training can induce as much as a 10% accuracy loss during inference [17]. When

introducing translation noise during inference, translation errors have been shown to decrease performance by at least 0.546 F-score (an 83% relative decrease in F-score) [25]. These dynamics are dependent on ML model architecture and objective, but addressing alignment errors consistently improves model performance. Third, some robot systems leverage geospatial data for navigation, and alignment errors can increase risk in these operations[8], [13]. These details mean that robot and ML systems need to understand the alignment errors that are present in predefined spatial data.

### B. What are the sources of alignment error?

Alignment errors between sUAS imagery and spatial data derive from five sources: satellite imagery acquisition, polygon generation, scale variance between drone and satellite, drone GPS noise, and errors induced by the orthomosaic generation process.

Inconsistencies in the satellite positioning (radiometric distortion), and errors in GPS signals or other sensor systems allow for inaccuracies within satellite imagery [6], [19]. The errors in satellite imagery allow for errors later in building polygon generation.

From the satellite imagery, the polygons are generated through semantic segmentation followed by polygonization [2] or they are manually generated [26], [11]. In either case, an opportunity for error arises.

Building polygons generated based on satellite imagery present positional inaccuracies when translated onto drone imagery. Satellites and drones do not capture imagery at the same ground sampling distances (GSD). Satellites capture imagery at a wide variety of GSDs. For disaster response, it is common to use GSDs in the 1 m/px to 30 cm/px range [18], [11]; however, lower resolutions are also available for other applications. Drones, on the other hand, can generate imagery at a far lower GSD, and the imagery in this work, has an average of 3.74 cm/px. With the lower resolution of satellite imagery, there will be an intrinsic lack of precision when compared to the higher resolution drone imagery which has the potential to induce error.

Drone imagery is subject to the drone’s GPS noise which notably arises from atmospheric conditions [21], [27] and interference with urban structures [20]. When there is GPS error, the overlaid building polygons will be based on positional inaccuracies, creating misalignment.

The orthomosaic generation process is another area where errors may be introduced in drone imagery. After drone imagery is collected, it is used to construct an orthomosaic, shown in Figure 2. During construction, the orthomosaic generation software may manipulate each image by rotating, scaling, skewing and translating. This process represents another opportunity for errors to arise [21], [7].

### C. Related Work in Photogrammetry Accuracy

While there has been substantial work in developing and measuring the accuracy of data products generated by photogrammetry software [5], [7] and further work measuring the accuracy when imagery is captured by sUAS [15],

[16], [4], [12], this paper is interested in the degree to which the orthomosaics align with programmatically generated building polygons obtained from satellite imagery. The common approach in these papers is to use ground control points (GCPs)-points on the earth whose position is known precisely-to measure the accuracy of, and in some cases further refine, the generated data product or orthomosaic [15], [16], [4], [12], [21]. While the content of this paper can be seen as a neighbor to this area of work, it is different as this work makes no use of GCPs and is not interested in the underlying accuracy of orthomosaics.

#### D. Related Work in Polygons and Geospatial ML

Prior works [10], [17], [25], [26] have already identified that coincident satellite imagery will have spatial errors that need to be managed when training machine learning (ML) models; however, none of these have explicitly measured the distribution of alignment errors, which is the focus of this paper. The two studies that measured the impact of alignment quantitatively [25], [17] only measured the downstream impact on ML model performance rather than the alignment phenomenon itself.

Four strategies have been considered for the management of alignment errors. First, and most simply, in satellite imagery, a uniform translation to all building polygons has been considered [10]. A second case was able to use “off-the-shelf image registration algorithms” to automatically align spatial polygons with imagery [9]. In a third case, [25], more complex techniques, CNN models and Markov Random Fields, were explored. Finally, another effort subverted this problem by maintaining an internal set of polygons that they constructed and then queried using nearest neighbor search with a 20-meter threshold [23].

### III. APPROACH

The approach taken in this work is to manually adjust the overlaid building polygons such that they align with the buildings in the imagery. Following this, the alignments’ angles and distances are aggregated, and misalignment is quantified using the statistical measures, mean and variance, the correlational measures, spearman  $\rho$  and Pearson R, and the intersection over union (IoU) of the aligned and unaligned building polygons. This section details the orthomosaic imagery, the building polygons used, and the adjustment generation process, with the analysis following in Section IV.

#### A. sUAS Orthomosaic Imagery

The imagery used within this analysis to align the building polygons consisted of 51 orthomosaics collected from sUAS at 9 wide-area disasters in the United States. orthomosaic (ortho) imagery (shown in Figure 2) are collections of raw images taken from an aerial camera that have been stitched together to produce a single image that is pixel aligned with longitude and latitude. This imagery includes 6 hurricanes- Hurricane Ian, Hurricane Harvey, Hurricane Michael, Hurricane Ida, Hurricane Idalia, and Hurricane




Fig. 3: Alignment vector field of the 1001-San-Carlos-Island orthomosaic. Note the diversity and non-linearity of the vector field.

Laura- the Mayfield Tornado, the Musset Bayou Fire, and the Kilauea Eruption.

The imagery that was used in the creation of these orthos was collected by 9 distinct models of drone. In descending order of prevalence these were the DJI Mavic 2 (16 orthos), SenseFly eBee X (12 orthos), DJI Matrice 30 (6 orthos), DJI Mavic Pro (5 orthos), DJI Phantom 4 (4 orthos), DJI Matrice 600 (3 orthos), DJI Matrice 300 (3 orthos), Wingtra WingtraOne Gen II (2 orthos), Parrot Anafi (1 ortho). In two instances, two different drone models were used to collect a single ortho. Orthos were generated using Pix4D React (98%)[21], and Agisoft Metashape (2%)[3].

#### B. Building Polygons

In this analysis, the orthos were overlaid with building polygons provided by the Microsoft Building Footprints Dataset [2]. The building polygons were generated through semantic segmentation on US satellite imagery.

#### C. Manual Adjustment Method

The manual adjustment method consisted of annotating the imagery with adjustments and using the annotated adjustments to perform alignments on the building polygons. An *adjustment* refers to a line that is drawn on an image by an annotator. Each adjustment is a line drawn from a building polygon vertex to the true location of the vertex in the image, an example of this is shown in Figure 1. Whereas an *alignment* refers to the translation done to the building polygon itself. During the annotation process, adjustments were provided for 33% of all building polygons. The remaining, unannotated, building polygons were aligned based on the following logic.

When an alignment is needed in a given area, the nearest adjustment is taken and applied. This process forms a vector field, an example is shown in Figure 3. More formally, this vector field is generated using a piecewise linear combination of all of the adjustments by selecting the nearest adjustment

TABLE I: Table with mean alignment angles (degrees), distance (pixels), distance (centimeters) and average variance alignment angles (degrees), distance (pixels), distance (centimeters) for all 51 orthomosaics and individual disaster events. The value for N is the number of alignments across all 51 orthomosaics. The disaster events are ordered alphabetically.

| Alignment Aggregate Mean and Variances (N=21619) |              |               |               |              |               |               |
|--------------------------------------------------|--------------|---------------|---------------|--------------|---------------|---------------|
|                                                  | Mean         |               |               | Variance     |               |               |
|                                                  | Angle (deg.) | Distance (px) | Distance (cm) | Angle (deg.) | Distance (px) | Distance (cm) |
| Overall                                          | 209.23       | 81.68         | 276.55        | 0.40         | 0.45          | 1.77          |
| Hurricane Harvey                                 | 205.29       | 130.59        | 303.69        | 0.37         | 0.41          | 1.08          |
| Hurricane Ian                                    | 224.47       | 75.71         | 264.28        | 0.40         | 0.42          | 1.43          |
| Hurricane Ida                                    | 193.60       | 55.86         | 166.82        | 0.26         | 0.48          | 1.47          |
| Hurricane Idalia                                 | 229.37       | 17.94         | 227.79        | 0.30         | 0.57          | 7.26          |
| Hurricane Laura                                  | 202.94       | 127.19        | 493.15        | 0.23         | 0.44          | 1.77          |
| Hurricane Michael                                | 236.30       | 117.37        | 243.89        | 0.38         | 0.50          | 0.90          |
| Kilauea Eruption                                 | 147.36       | 65.63         | 462.50        | 0.89         | 0.57          | 4.01          |
| Mayfield Tornado                                 | 194.43       | 80.41         | 217.64        | 0.41         | 0.47          | 1.31          |
| Musset Bayou Fire                                | 181.39       | 101.62        | 295.63        | 0.36         | 0.49          | 1.38          |

to populate the vector field. Following adjustment annotations, the generated vector field was reviewed to ensure that all alignments were accurate.

#### IV. ANALYSIS OF MISALIGNMENT

As mentioned in Section III, the building polygon alignments' angles and distances, were measured and aggregated, and their misalignments were quantified through statistical, correlation, and IoU measures. The distances of the alignments were computed for both the pixel and centimeter distances of the adjustments. This section presents the analysis of the results with a discussion following in Section V.

##### A. Mean

The mean distance and angle of aligned building polygons were used to estimate the average case for alignment errors in orthomosaic imagery. The results of which are presented in Table I. The analysis finds that across all 51 orthos the mean alignment angle was 209.23 degrees, and the mean alignment distance was 81.68 px or 276.55 cm. These means are further analyzed by dividing the dataset by the 9 disasters. The 9 disasters' mean alignment angles range from 147.36 degrees (Kilauea Eruption) to 236.30 degrees (Hurricane Michael). Their mean alignment pixel distances range from 17.94 px (Hurricane Idalia) to 130.59 px (Hurricane Harvey). As for their mean alignment centimeter distance, they range from 166.82 cm (Hurricane Ida) to 493.15 cm (Hurricane Laura).

##### B. Variance

The variances of distances and angles of aligned building polygons were used to characterize the width of the distribution of alignment errors in orthomosaic imagery. The average variances for the angles and distances (pixels and centimeters) across all 51 orthos are shown in Table I. The overall average angle variance is 0.40 degrees and the average distance variance is 0.45 px or 1.77 cm. The variances are further analyzed for each of the 9 disaster events. The analysis finds the angle variance range from 0.23 degrees (Hurricane Laura) to 0.89 degrees (Kilauea Eruption). As for the average distance variance, they range from 0.41 px (Hurricane Harvey) to 0.57 px (Hurricane Idalia and Kilauea

Eruption). The average distance centimeter variances range from 0.90 cm (Hurricane Micheal) to 4.01 cm (Kilauea Eruption).

##### C. Correlation

The correlation between the angle and distance of alignments was utilized to determine if there was a connection between the distribution of angles and distances within an individual ortho. The correlation between these two variables was measured using both the Pearson R and Spearman  $\rho$  statistics. The average Pearson R and average Spearman  $\rho$  across all orthos was 0.06 (Max 0.47, Min -0.57) and 0.067 (Max 0.74, Min -0.72), respectively. Orthos with  $\leq 2$  adjustments were ignored in this evaluation. This result is somewhat unsurprising, as a positive correlation would indicate a spiraling pattern, which has not been observed.

##### D. Intersection over Union

The intersection over the union (IoU) of the aligned and unaligned building polygons was explored to contextualize the scale of adjustments with the scale of the building polygons. The IoUs for all building polygons were taken and aggregated per ortho. The average IoUs ranged from 0.24 to 0.84, averaging 0.65.

#### V. DISCUSSION

The analysis results presented in Section IV are indicative of varying distributions, reasoning a need for a in-depth interpretation of them, and push for future work to address the needs of alignment. This section further discusses the analysis results and implications for future work.

##### A. Interpretation of Results

Considering the results presented in Section IV, the evident trend is a lack of uniformity in the types of errors that are encountered. It is noted that the spatial distance errors, when grouped by disaster, are within the same order of magnitude, both in terms of pixels and centimeters. However, these metrics are being computed at the disaster level, and it may be more insightful to inspect these results at the level of each ortho, shown in Figure 4. From this, it is observed that




Fig. 4: Violin plots which describe the distributions of angles and distance metrics of alignments from 9 selected orthomosaics. Note the non-uniform, non-linear, and varying nature of these distributions. Row 1 plots the distribution of alignment angles. Row 2 plots the distribution of alignment distances in terms of pixels. orthomosaics were selected for visualization by taking the orthomosaic with the largest count of buildings for each of the 9 disasters from which imagery was collected. In each violin plot, the topmost and bottommost line mark the extremes within the data, with the middle line representing the mean of the distribution. The value for N is the number of alignments represented.

these adjustments are non-uniform within the orthomosaic as well. While some of these distributions appear normal, there are others that do not. Moving to Figure 5, these errors are spatially non-linear and discontinuous in both adjustment distance and angle.

#### B. Implications for ML with Spatial Data and sUAS Imagery

As discussed in Section III, building performant ML systems that leverage spatial data depends upon the accurate alignment of that spatial data with the imagery for effective training and inference. Spatial alignment needs to be considered and managed when working with spatial data and sUAS imagery like this. Despite being captured with in situ GPS measurements and at a far higher resolution than satellite imagery, sUAS imagery is not immune to these issues.

There are two potential directions for techniques that can be leveraged to manage these alignment errors. First, preprocessing techniques that manage alignment before presentation to the ML model. Second, trained models that manage alignment errors explicitly so preprocessing is no longer necessary. Both paths are considered viable options. However, whatever methods are developed to manage the

alignment issue in sUAS imagery need to be non-linear in nature in order to capture the types of non-linearities observed in this study.

## VI. CONCLUSIONS

This work presented the first quantitative study, explicitly measuring the distributions of the building polygon misalignments, and releases the first publicly available dataset of ground truth alignments for sUAS imagery. Based on the 51 orthomosaics, spanning 9 wide-area disasters and consisting of 21,619 building polygons, involved in this analysis, the study found that, with an average angle alignment variance of 0.4 degrees and an average distance alignment variance of 0.45 pixels, the alignment distributions are non-uniform between and within the orthomosaics collected. Given the impacts on machine learning systems and robot navigation of these misalignments, future work is needed if sUAS imagery is expected to be leveraged in these contexts, and especially in the disaster setting.

## ACKNOWLEDGMENT

This material is based upon work supported by the AI Research Institutes Program funded by the National Science




Fig. 5: Color map visualization for the angles and distance of alignments in the 20230830-SteinhatcheeRiver orthomosaic. Note the discontinuous and non-uniform distribution of adjustments. [left] orthomosaic's imagery. [center] Color map of the alignment's distance (pixels). Dark colors indicate regions that have alignments with distance near 0 with brightness increasing with distance. [right] Color map of the alignment's angle (degrees). Dark colors indicate regions of alignments with angles near 180 degrees, and bright colors indicate alignments near 0 or 360 degrees.

Foundation under AI Institute for Societal Decision Making (AI-SDM), Award No. 2229881 and under "Datasets for Uncrewed Aerial System (UAS) and Remote Responder Performance from Hurricane Ian" Award No. 2306453. Acknowledgment is given to CRASAR and David Merrick for the acquisition of the imagery and supporting information.

#### REFERENCES

- [1] "Drone deploy ml segmentation benchmark," <https://github.com/dronedeploy/dd-ml-segmentation-benchmark>, 2019.
- [2] "Microsoft us building footprints," <https://github.com/Microsoft/USBuildingFootprints>, 2021.
- [3] Agisoft-LLC, "Agisoft metashape," <https://www.agisoft.com/>, 2024.
- [4] S. Azim, J. Rasmussen, J. Nielsen, R. Gislum, M. S. Laursen, and S. Christensen, "Manual geo-rectification to improve the spatial accuracy of ortho-mosaics based on images from consumer-grade unmanned aerial vehicles (uavs)," *Precision agriculture*, vol. 20, pp. 1199–1210, 2019.
- [5] A. Barbasiewicz, T. Widerski, and K. Daliga, "The analysis of the accuracy of spatial models using photogrammetric software: Agisoft photoscan and pix4d," in *E3S Web of Conferences*, vol. 26. EDP Sciences, 2018, p. 00012.
- [6] R. G. Congalton and K. Green, *Assessing the accuracy of remotely sensed data: principles and practices*. CRC press, 2019.
- [7] P. Dare, N. Pendlebury, and C. Frase, "Digital orthomosaics as a source of control for geometrically correction high resolution satellite imagery," in *Proceedings of the 23rd Asian Conference on Remote Sensing*, no. 173, 2002.
- [8] P. Fleischmann, T. Pfister, M. Oswald, and K. Berns, "Using openstreetmap for autonomous mobile robot navigation," in *Intelligent Autonomous Systems 14: Proceedings of the 14th International Conference IAS-14 14*. Springer, 2017, pp. 883–895.
- [9] D. Garcia, G. Mateo-Garcia, H. Bernhardt, R. Hagenseieker, I. G. L. Francos, J. Stock, G. Schumann, K. Dobbs, and F. Kalaitzis, "Pix2streams: Dynamic hydrology maps from satellite-lidar fusion," *arXiv preprint arXiv:2011.07584*, 2020.
- [10] R. Gupta, B. Goodman, N. Patel, R. Hosfelt, S. Sajevee, E. Heim, J. Doshi, K. Lucas, H. Choset, and M. Gaston, "Creating xbd: A dataset for assessing building damage from satellite imagery," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 2019, pp. 10–17.
- [11] R. Gupta, R. Hosfelt, S. Sajevee, N. Patel, B. Goodman, J. Doshi, E. Heim, H. Choset, and M. Gaston, "xbd: A dataset for assessing building damage from satellite imagery," *arXiv preprint arXiv:1911.09296*, 2019.
- [12] I.-K. Hung, D. Unger, D. Kulhavy, and Y. Zhang, "Positional precision analysis of orthomosaics derived from drone captured aerial imagery," *Drones*, vol. 3, no. 2, p. 46, 2019.
- [13] M. Kunz, A. Vierling, P. Wolf, and K. Berns, "Localization using osm landmarks," in *2021 Latin American Robotics Symposium (LARS), 2021 Brazilian Symposium on Robotics (SBR), and 2021 Workshop on Robotics in Education (WRE)*. IEEE, 2021, pp. 228–233.
- [14] P. P. Liang, A. Zadeh, and L.-P. Morency, "Foundations and recent trends in multimodal machine learning: Principles, challenges, and open questions," *arXiv preprint arXiv:2209.03430*, 2022.
- [15] N. Liba and J. Berg-Jürgens, "Accuracy of orthomosaic generated by different methods in example of uav platform must q," in *IOP Conference Series: Materials Science and Engineering*, vol. 96, no. 1. IOP Publishing, 2015, p. 012041.
- [16] M. Ludwig, C. M. Runge, N. Friess, T. L. Koch, S. Richter, S. Seyfried, L. Wraase, A. Lobo, M.-T. Sebastià, C. Reudenbach *et al.*, "Quality assessment of photogrammetric methods—a workflow for reproducible uas orthomosaics," *Remote Sensing*, vol. 12, no. 22, p. 3831, 2020.
- [17] A. Maiti, S. Oude Elberink, and G. Vosselman, "Effect of label noise in semantic segmentation of high resolution aerial images and height data," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 2, pp. 275–282, 2022.
- [18] Maxar, "Clarity and confidence how high-resolution satellite imagery can help you extract critical information at scale," <https://explore.maxar.com/Imagery-Leadership-Spatial-Resolution>, 2024.
- [19] —, "Here or there? how accuracy in satellite imagery impacts your operations and outcomes," <https://explore.maxar.com/Imagery-Leadership-Accuracy.html>, 2024.
- [20] K. Merry and P. Bettinger, "Smartphone gps accuracy study in an urban environment," *PloS one*, vol. 14, no. 7, p. e0219890, 2019.
- [21] Pix4D, "Pix4d react: 2d fast-mapping for emergency response and public safety," <https://www.pix4d.com/>, 2024.
- [22] M. Rahnemoonfar, T. Chowdhury, A. Sarkar, D. Varshney, M. Yari, and R. R. Murphy, "Floodnet: A high resolution aerial imagery dataset for post flood scene understanding," *IEEE Access*, vol. 9, pp. 89 644–89 654, 2021.
- [23] C. Robinson, S. F. Nsutezo, A. Ortiz, T. Sederholm, R. Dodhia, C. Birge, K. Richards, K. Pitcher, P. Duarte, and J. M. L. Ferres, "Rapid building damage assessment workflow: An implementation for the 2023 rolling fork, mississippi tornado event," *arXiv preprint arXiv:2306.12589*, 2023.
- [24] R. Szeliski, *Computer vision: algorithms and applications*. Springer Nature, 2022.
- [25] J. E. Vargas-Muñoz, S. Lobry, A. X. Falcão, and D. Tuia, "Correcting rural building annotations in openstreetmap using convolutional neural networks," *ISPRS journal of photogrammetry and remote sensing*, vol. 147, pp. 283–293, 2019.
- [26] J. E. Vargas-Munoz, S. Srivastava, D. Tuia, and A. X. Falcao, "Openstreetmap: Challenges and opportunities in machine learning and remote sensing," *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 1, pp. 184–199, 2020.
- [27] S. Williams, Y. Bock, and P. Fang, "Integrated satellite interferometry: Tropospheric noise, gps estimates and implications for interferometric synthetic aperture radar products," *Journal of Geophysical Research: Solid Earth*, vol. 103, no. B11, pp. 27 051–27 067, 1998.
- [28] X. Zhu, J. Liang, and A. Hauptmann, "Msnet: A multilevel instance segmentation network for natural disaster damage assessment in aerial videos," in *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 2021, pp. 2023–2032.