

## Proceedings in UAS-Assisted Bridge Inspections: RTK-Based Photogrammetric Reconstruction and Spatial Filtering

E.T. Bartczak\*, M. Bassier and M. Vergauwen

Department of Civil Engineering  
Faculty of Engineering Technology  
Geomatics Research Group  
Gebroeders De Smetstraat 1, B-9000 Gent, Belgium  
erkkitobias.bartczak@kuleuven.be

**KEY WORDS:** Unmanned aerial systems, bridge, inspection, damage detection, photogrammetry, deep learning.

### ABSTRACT:

Traditional bridge inspections present considerable challenges in terms of efficiency and accuracy. However, recent advancements in Unmanned Aerial Systems (UASs) and deep learning for object detection have opened up new avenues for automatic bridge damage detection. We present a comprehensive framework leveraging these technologies for automated damage detection in UAS imagery, followed by accurate mapping of the damage predictions on photogrammetric models. In this work, we propose a photogrammetric procedure to retrieve geolocated bridge models solely based on Real-Time Kinematics (RTK) information. Within the damage detection step, we conduct extensive testing and optimization of model hyperparameters using YOLOv8 and Slicing Aided Hyper Interference (SAHI). Next, we map the predictions onto the 3D model using ray casting, allowing to group and filter the predictions by their area and position. Finally, a Graphical User Interface (GUI) allows bridge inspectors to identify false positive predictions, generate new training data, and directly measure damage dimensions in the images. Validation on a concrete box girder bridge resulted in a photogrammetric model with a mean error of 1.3 cm, negating the need for ground control points. Our model training process revealed substantial performance variations between training and test datasets, underscoring the importance of evaluating optimal hyperparameters on UAS inspection images rather than relying on the validation metrics. Lastly, we successfully map the detected damage and create new training data from the UAS inspection images. This framework significantly enhances bridge inspection accuracy and efficiency, providing a strong foundation for future developments in automated bridge inspections.

### 1. INTRODUCTION

Bridges, key elements of infrastructure, link cities, enable trade, and ensure public safety. With traffic loads often surpassing original expectations and environmental factors exerting constant pressure, many of these crucial structures are approaching their intended lifespan's limit. To ensure the continued health and functionality of these structures, it is vital to perform thorough inspections and detect potential damages at the earliest possible stage. Traditional bridge inspections struggle with cost-effectiveness, safety, and accuracy. Conventional techniques like under-bridge vehicles or rappelling disrupt traffic and are time-consuming, making timely inspections a challenge for authorities.

Amidst the ongoing technological evolution, the application of UASs in bridge inspections emerges as a notable advancement. These systems present an alternative that not only improves cost-efficiency and safety but also substantially mitigates traffic disruptions, and facilitates access to previously unreachable sections of bridge structures. The integration of machine learning, specifically deep learning algorithms, has further propelled this technological shift, enhancing the proficiency of UAS-based inspections. However, several challenges remain to enable accurate automated UAS-based bridge inspections.

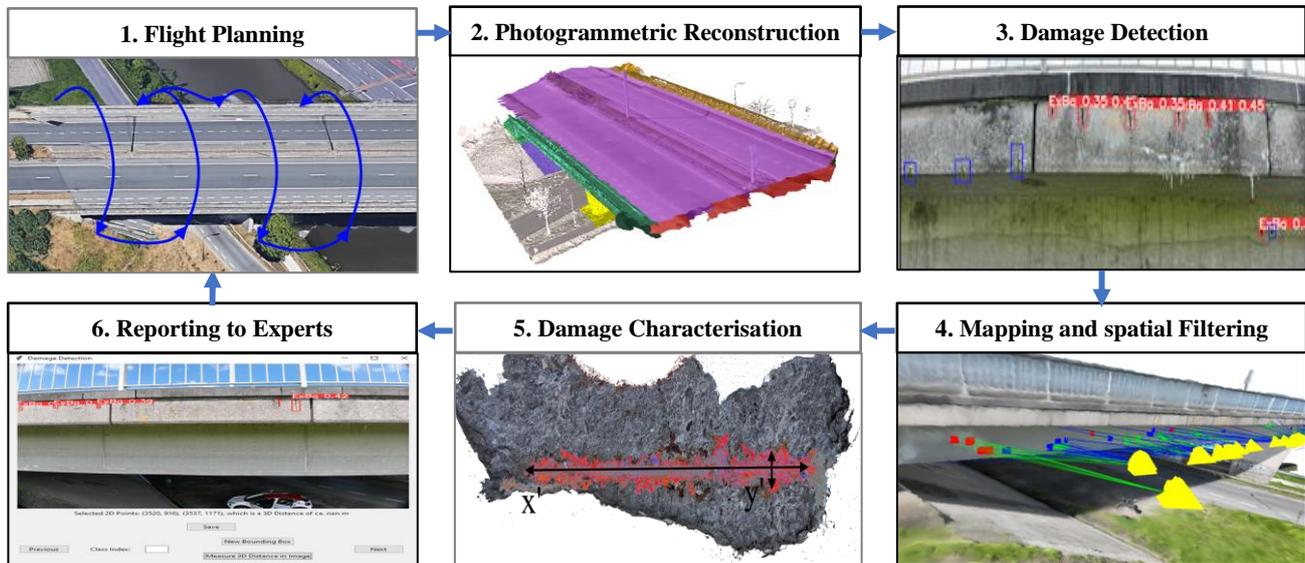
To further maximize the effectiveness of the UAS inspection procedure, it is imperative that the photogrammetric model is highly precise. Minimizing field work is a critical factor; hence, the reliance on ground control points should be minimized, and the focus should shift towards leveraging the RTK data from the UAS for photogrammetric reconstruction.

In regards to the automatic detection of bridge damages, one key limitation is the discrepancy of training images obtained from decades of conventional bridge inspections and the UAS image quality obtained from inspection flights. The optimal combination of these hyperparameters is not a one-size-fits-all scenario, but rather, it should be fine-tuned based on the UAS inspection images. In the pursuit of exhaustive damage detection, even predictions with low confidence should be factored into consideration. While this approach may lead to an increased number of false positives, it significantly reduces the risk of overlooking potential damage. However, further advancements in filtering the false positive predictions are necessary to assure efficient and reliable inspections. In this research, we emphasise on identifying bridge damages, particularly exposed rebars, by adopting a three-staged strategy: (1) the development of a photogrammetric model based exclusively on the RTK data from UASs, (2) the analysis of performance changes while fine-tuning the hyperparameters of a state-of-the-art object detection model, and (3) the localization and filtration of potential damage instances.

The structure of this paper is as follows: Chapter 2 offers an overview of existing literature in this field. Chapter 3 presents our methodology, shedding light on the process of creating a UAS-based photogrammetric model of the bridge, training the damage detection model, and implementing post-prediction steps to locate and filter potential damages. We validate our proposed method through a case study of a complex multiple concrete box girder bridge in Chapter 4. The subsequent chapters, 5 and 6, deliver our findings and subsequent discussions, ultimately evaluating the efficacy of our proposed methodology.

---

\* Corresponding author



**Fig. 1.** Overall framework for comprehensive UAS-based bridge inspections. All images are part of the case study of this work.

## 2. RELATED WORK

In this section, we provide an overview of studies related to automated routine bridge inspections. We focus on (1) the role of UASs in creating accurate 3D models of bridges, and (2) the challenges often encountered in object detection using machine learning, particularly when applied to UAS bridge inspections.

### 2.1 UAS based photogrammetric bridge models

Advancements in photogrammetry have presented promising alternatives for surveying and digitizing existing bridges. Particularly, UASs have emerged as a potent tool, offering numerous advantages over traditional terrestrial laser scanning (TLS) for creating 3D models of bridges. Chen et al. (2019) highlighted that UAS photogrammetry is not only cost-effective but also reduces labour on-site. Furthermore, they reported a more equal distribution of point density than the TLS as well as an increased coverage of the reconstruction for extensive flight routes. The reported geometric accuracy of the resulting 3D reconstructions ranges from 5 mm (Hallermann and Morgenthal 2016) to 32.2 mm (Chen et al. 2019) compared to check points and TLS. However, reported challenges of using UAS photogrammetry for bridge digitization include sensitivity to weather conditions, and difficulties in capturing slender objects like railings or surfaces with minimal features (Otero and Gagliardo 2015). Another challenge in UAS photogrammetry is the geo-localisation of the reconstruction and a frequently occurring convex deformation of the models, the so-called dome effect (James and Robson 2014). It has been found that the use of Ground Control Points (GCPs) can reduce the deformations by introducing local constrictions (Tscharf 2020), but presents a time consuming effort (Rock et al. 2011). Using the UAS-RTK information has the potential of minimizing the need for deploying the GCPs, while achieving comparable geometric accuracy of the resulting models. In a land survey, Štroner et al. (2021) achieve a accuracy of the resulting model of up to 3 cm, proving the potential of this approach. However, the process and the achievable accuracy of solely based on RTK-UAS imagery has still to be investigated for the use of bridge inspections.

### 2.2 Damage detection in UAS images

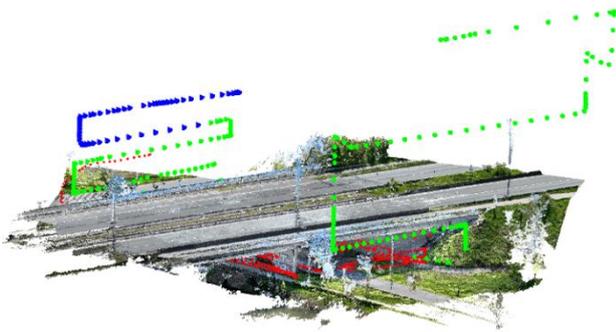
In the past years, various machine learning architectures have been developed to detect various damages in steel, concrete and masonry structures. Recent research, underscores the remarkable

capabilities of contemporary object detection models, illustrating their success in accurately identifying cracks as narrow as 0.2 mm (Ding et al. 2023). An extensive overview of available detection models for bridge inspections is given in Toriumi et al. (2021).

However, applying object detection to bridge damage detection poses several challenges. Large training datasets for this purpose exist, sourced from a lengthy history of conventional bridge inspection documentation. These datasets typically comprise images from a diverse range of camera types, such as high-resolution cameras, smartphones, and tablets. Consequently, the camera quality, perspective angle, and ground sampling resolution significantly differ from the imagery produced by UAS inspections. This discrepancy is known in machine learning as the so-called domain shift and creates a substantial challenge (Bukhsh et al. 2021). Therefore, these models frequently underperform in the context of UAS bridge damage detection due to a lack of specialized training data. Another challenge is related to the small object size of the damages within the images (Mittal et al. 2020). A sliding window approach has been proposed to tackle this challenge. This method has demonstrated a performance increase of up to 14.5% on several benchmark datasets and is compatible across a wide range of commonly used object detection architectures (Akyon et al. 2022). However, these challenges have not yet been specifically investigated in the context of UAS bridge damage detection.

In various studies, bridge damage detection was successfully applied. Liang et al. (2023) achieved an Average Precision (AP) of 0.67 to automatically detect exposed bars in UAS images. Lastly, ray casting has been proposed to map the damage detection onto 3D models (Lin et al. 2021). This allows to group the predictions, leading to an improvement of up to 3.5% in AP. However, the impact of further spatial filtering techniques on the overall damage detection performance has yet to be investigated.

This research builds upon the project initiated by Bartczak et al. (2023) and focuses on the advancement of a comprehensive bridge inspection procedure. The envisioned framework, as detailed in Fig. 1, advocates for recurring UAS inspection flights, automated damage detection, damage mapping, and includes an expert-in-the-loop review stage in data management. This paper showcases our advancements in stages 2 (RTK-based photogrammetric reconstruction), 3 (model training), 4 (spatial filtering), and 6 (Reporting to experts).



**Fig. 2.** RTK-based point cloud with camera positions: green and blue cones represent cameras with optimal GPS signal, while red cones indicate impaired GPS signal.

### 3. METHODOLOGY

This chapter details our refined methodology for UAS-based bridge inspection. It primarily focuses on the facilitation of photogrammetric processing using RTK data (3.1), the exploration of ideal hyperparameters and datasets for training damage detection models (3.2) and a spatial filtering technique designed to enhance detection performance (3.3).

#### 3.1 UAS-RTK based photogrammetric reconstruction

In this phase of our framework, we integrate the image acquisition stage with photogrammetric processing, proposing an RTK-based method to scale and geolocate the 3D models, thereby reducing onsite effort.

Given the inconsistent reliability of RTK information for camera positions under the deck, we adopt a purely feature-based alignment. Therefore, we first remove all GPS data from the cameras to ensure dependable image registration. This approach capitalizes on the rich feature set of typical bridge imagery and employs sequential image pairing to avoid misalignments and inaccuracies potentially introduced by additional constrictions of the RTK information. In a subsequent photogrammetric stage, we align only the top-view images of the bridge (**Fig. 2**), this time

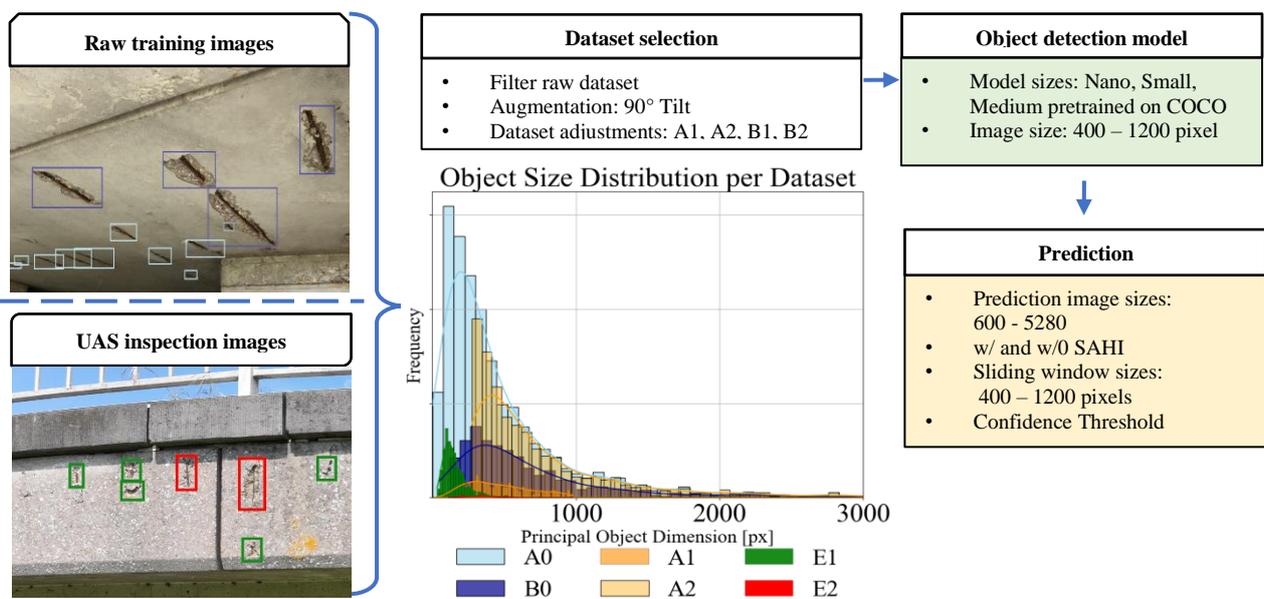
incorporating the RTK information. To safeguard traffic, this flight path is maintained at a substantial distance from the bridge, consequently providing a more reliable GPS signal. The photogrammetric reconstructions from both stages are subsequently aligned based on the identical cameras. Finally, we segment the geolocated 3D bridge model into the structural components, such as ground pillars, underdeck, pier cap, and others.

To assess the efficacy of this proposed method, we design a case study to generate multiple photogrammetric models using the same images but including varying levels of RTK information. The fundamental assumption of our research is that this process will yield more accurate 3D models. To verify this, we compare the accuracy of the resulting dense point clouds against a TLS using a Cloud-to-Cloud (C2C) comparison. This TLS is georeferenced using six Leica targets surveyed with an Emlid Reach RS2+. In order to evaluate absolute accuracy, no further alignment between point clouds is performed.

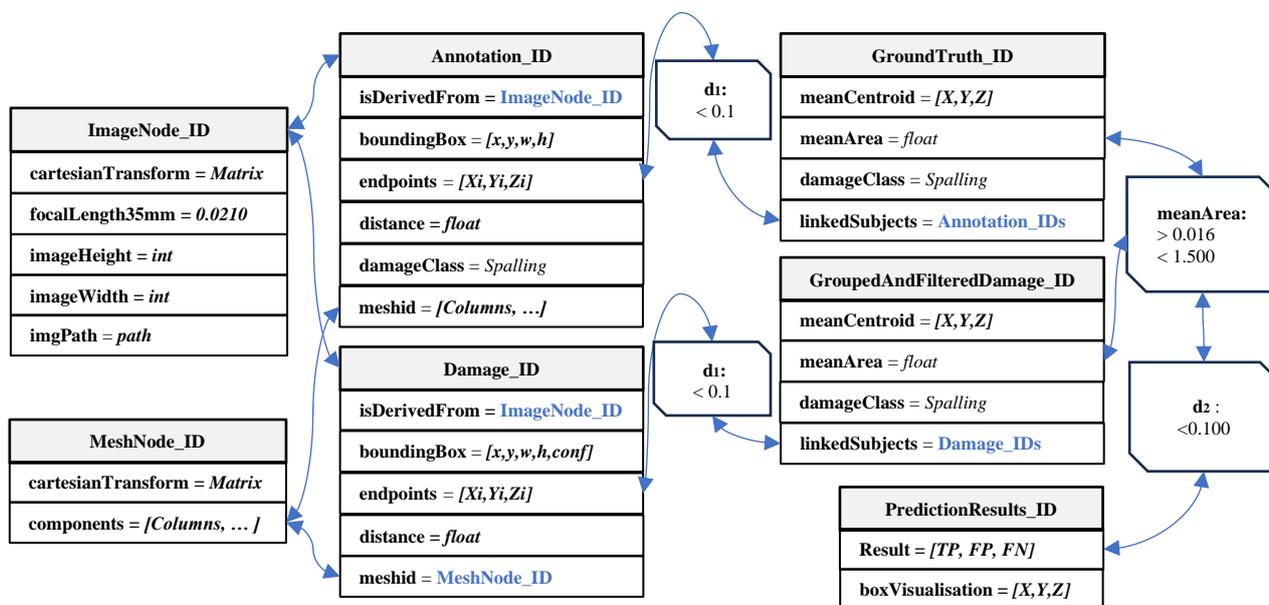
#### 3.2 Object detection model training for damage detection

One key challenge in deploying object detection models for UAS-based bridge inspections is mitigating the domain shift between inspection documentation images, used for training, and the UAS-acquired images used in model deployment. The optimal combination of various model parameters is crucial to superior model performance. We conduct a systematic study to understand the impact of different hyperparameter settings, iteratively adjusting the variables as outlined in **Fig. 3**. Our goal is to identify the optimal parameter combination to maximize our damage detection models' performance.

In the initial data processing stage, a quality control filter is applied to the raw images, ensuring they meet a basic threshold of quality and consistency. Roboflow (2023) is used for damage annotation, with stricter filtering conditions applied to establish two baseline datasets; (A0) featuring original annotations, and (B0) comprising only the largest and most visible damages. To address the issue of object size disparity between the training data and UAS inspection imagery, we analyse images from an inspection flight along the pier caps. We adjust the training



**Fig. 3.** Discrepancy between the training dataset and the inspection images. The histogram shows the distribution object sizes in the different datasets used for training and evaluation. The model training stage concludes the investigated settings in this work.



**Fig. 4.** Visualization of ontology. The Euclidean distances  $d_1$  and  $d_2$  as well as the  $meanArea$  are used to group and filter the predictions, reducing the number of instances and false positive predictions.

datasets, including (A/B1) only the intersection of object sizes and (A/B2) only the largest 50 % of the objects. Additionally, data augmentation, is employed separately via a 90° rotation of the images. In the training phase, multiple object detection models are trained, with hyperparameters such as image size and model size iterated. In the prediction phase, SAHI is used additionally to handle the small object problem, iterating between various window sizes and confidence thresholds.

The model performance evaluation involves assessing recall and precision metrics. Since we opt on minimizing missed damages, we primarily focus on the F2-score over the AP. All models are evaluated first on the validation images in the training dataset and second on the inspection images (E1). Since the inspection data includes very minute damages, we created a third dataset only including the largest objects (E2). The best performing models are subsequently deployed in the last step of our proposed framework.

### 3.3 Damage mapping and spatial analysis

The raw predictions of the damage detection model, may include large numbers of false positive predictions. To improve the performance metrics, we map the predictions in the images onto the 3D model derived from the photogrammetric process. Concretely, we utilize the Geomapi library (2023) to cast rays originating from the camera positions, targeted at the centre point of each prediction bounding box, and determine the intersection point on the 3D model. This process serves multiple purposes: it eliminates false positive predictions that don't intersect with the bridge model, groups potential damages based on their location, allows for spatial filtering over the damage candidate area, and extracts semantic information, such as the specific bridge component impacted.

In practice, our framework adopts an Resource Description Framework (RDF) architecture comprised of several nodes, as illustrated in Fig. 4. In our spatial filter approach, we first group potential damage sites based on the distance  $d_1$ , between the ray-casted prediction centre points. We then approximate the damage area by dividing the convex hull of all 3D points, including the

corner points, into triangles and summing their areas, a method based on Heron's formula. The centroid of the damage area is also computed as the mean position of all the centre points. Finally, we apply a threshold filter to the surface area to filter out irrelevant or unrealistically large damages. We process both the ground truth annotations and predictions through this pipeline, compare the results using a distance tolerance  $d_2$ , and then calculate the final F2-score for evaluation and visualization

For visualization and further inspection, a user interface is utilized. This GUI presents all images associated with a specific damage prediction simultaneously, allowing inspectors to gauge the likelihood and severity of potential damage more effectively. Furthermore, inspectors can refine the labels of an image, thereby creating more accurate training data for future object detection processes. The GUI is also equipped with functionality to perform ray casting based on specific points selected within the image. This feature enables the calculation of 3D distances directly from inspection images, providing inspectors with the tools necessary to accurately assess the dimensions of identified damage within the image. Furthermore, it allows to swiftly annotate new training images, improving the damage detection process.

## 4. CASE STUDY

To validate the proposed framework, we conduct a case study, inspecting a 30x30 m section of a concrete highway bridge in Ghent, Belgium. This site is chosen for the real-world challenge it presents, featuring a complex four box-girder cross section, high traffic conditions and visible damages, including exposed bars along the pier caps. The inspection is carried out using a DJI Mavic 3 Enterprise.

### 4.1 Photogrammetric processing

To evaluate our proposed processing strategy, we conduct a photogrammetric flight gathering 1523 images in 1.5 hours. We then conducted a series of iterative reconstructions of the bridge, each time modifying only the amount of RTK data. As a baseline, we processed all images with their associated RTK data. In the subsequent iteration, this was reduced to 267 positions, captured

**Table 1.** Geometric accuracy of photogrammetric dense point clouds compared to TLS

Geolocation procedure	Number of RTK Images	Mean Distance [cm]	Standard Deviation [cm]
Including RTK data	1523	26.1	22.5
Including RTK data	54	1.5	2.0
Including RTK data	267	109.1	108.2
Post alignment	54	1.9	2.0
Post alignment	267	1.3	1.8

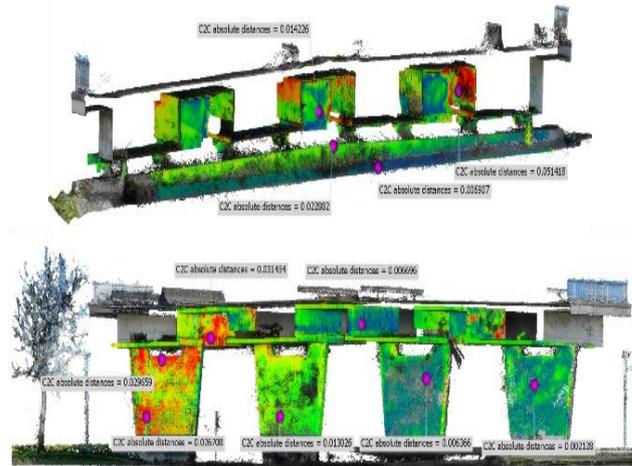
from above or adjacent to the bridge (Fig. 2). This number was further decreased to 54 in the following iteration, utilizing images from a single side of the structure. As proposed in this work, for the following iterations we exclude all RTK information during the photogrammetric process. Instead, we perform a separate alignment using the 267 images with RTK data, resulting in a distinct reconstruction and registered camera positions. The initial relative alignment, is then scaled using the camera positions from this separate alignment. Across these iterations, all models were processed under identical photogrammetric settings. The results are presented in Table 1.

In the course of the experiment, we observe significant misalignments and disrupted geometries. This is especially apparent in the third iteration where we use 267 images taken from above the bridge for camera alignment. In this case, the dense point cloud of the bridge deck is reconstructed separately and disoriented from the rest of the model. The second iteration, which utilized 54 georeferenced images results in a consistent point cloud reconstruction, but fails to register all cameras successfully. Ultimately, our proposed procedure achieved the best geometric accuracy with a mean distance of 1.3 cm, and successfully aligned all images.

Our investigation into photogrammetric processing outlines a strategy for generating highly accurate 3D bridge models. We conclude that our approach leads to slightly better geometric accuracy of the resulting point cloud while assuring that all cameras are registered successfully. To further analyse the results, we examine two cross-sections, as depicted in Fig. 5. The largest discrepancy is observed in challenging areas with insufficient camera coverage, such as between the box girders at the abutment, showing a maximum local variation of 5.1 cm. Notably, these regions also present difficulties for the TLS method, which exhibits reduced coverage and locally missing segments. In fact, conducting a TLS survey for the top part of the bridge is not feasible due to high traffic. In fact, conducting a TLS survey for the top part of the bridge is not feasible due to high traffic. This highlights the practical advantages of the UAS RTK surveying method in situations where traditional methods are not applicable.

#### 4.2 Model Optimization for Damage Detection

In our case study, we conduct an experiment aimed at automatically detecting bridge damages. Specifically, the chosen site provides a total 26 exposed bars of sizes ranging approximately from 1.5 to 30 cm. We train and evaluate several YOLOv8 models, following the procedure outlined in Section 3.2. Fig. 6 provides the F2-score of ten trained models across varying image sizes and training datasets.



**Fig. 5.** Cross-sections of final dense point cloud at the abutment (top) and the columns (bottom) the colour range is set between 0 (blue) to 6 cm (red).

#### Baseline Datasets:

The labelling process results in a second baseline dataset B0 containing only the most visible objects. Although the trained models show performance increase in detecting the larger objects of up to 30% using a training image size of 800 pixels, the F2-score on all objects is significantly lower.

#### Data Augmentation:

We find that a 90° image tilt enhances the performance of models trained on dataset B, resulting in an average F2-score increase of 7%. Conversely, models trained on dataset A show an average F2-score decrease of about 11% under the same conditions. This discrepancy may be due to the greater impact of augmentation on smaller datasets.

#### Dataset Adjustment:

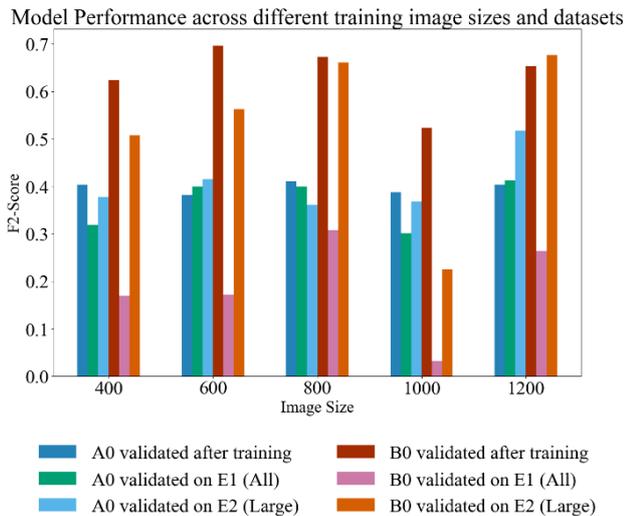
The adjustments to the training datasets have a modest impact on performance for models trained on dataset B. In some instances, the F2-score increases up to 7% for the detection of larger evaluation objects (E2). However, these adjustments do not significantly alter the overall performance consistently.

#### Model parameter size:

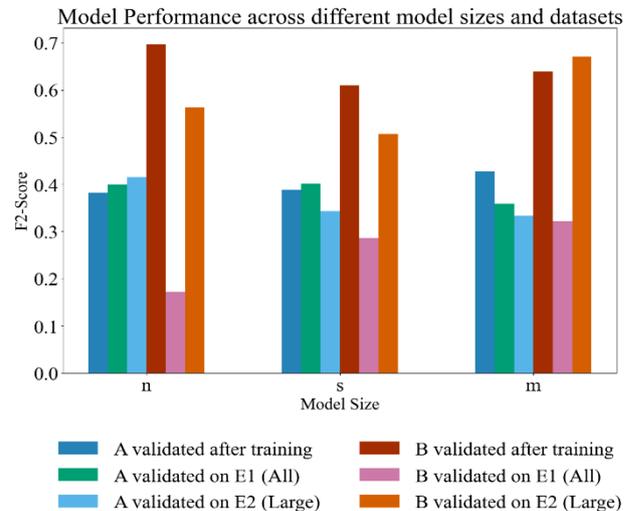
We iterate the model sizes using the nano, small and medium sized models provided by *Ultralytics* (Jocher et al. 2023), which are pretrained on the COCO dataset. For smaller image sizes, the model size does not significantly affect performance. The highest F2-score is achieved by a medium model at 800 pixels, which is the largest model and image size combination feasible on our GPU. Here, the performance is increased by 2.1% compared to the smaller model size. Comparing the effect on the Datasets A0 and B0, we observe a decrease in performance for dataset A0 when evaluating on the UAS dataset, while the performance of dataset B models increases (Fig. 7).

#### Training image size:

Analysis of our experiments suggests that larger training image size generally leads to improved model performance of up to 9.1%. However, models with a training image size of 1000 pixels significantly underperform, presumably due to outliers. The maximum processable image size on our Geforce RTX 2080 TI GPU Server was 1200 pixels.



**Figure 6:** Model performance comparison based on F2-score across varying image sizes (400, 600, 800, 1000, and 1200 pixels) for baseline datasets A and B. Each model is evaluated on the validation dataset, all objects (E1) and large objects (E2) in the UAS inspection datasets.



**Figure 7:** Model performance comparison based on F2-score across varying model sizes (Nano, Small and Medium) for baseline datasets A0 and B0, trained on an image size of 600 pixels.

### Prediction with SAHI:

Considering the small size of the damages in the images, we incorporate SAHI on a selection of the higher-performing models. We compare performance metrics with and without SAHI integration in the prediction stage. Notably, we find that SAHI generally enhances recall, although this is offset by significantly lower precision metrics. However, we do not see a consistently higher F2-score compared to predictions without SAHI. The best performance occurs when the sliding window size matches the training image size of the models

Our comprehensive analysis of object detection models reveals a complex relationship between model performance and various factors such as training dataset choice, object size distribution, and hyperparameter tuning. Notably, aligning the object sizes within the training dataset with those in the inspection images does not improve model performance consistently. Instead, we observe better performance from models trained on larger datasets featuring varied object sizes. Furthermore, while increasing image size during training generally improves results, we found an optimal image size of 800 pixels for our UAS image evaluation. The prediction stage emphasizes the importance of selecting appropriate parameters for image size, as the sliding window approach leads to both a higher true positive and false

**Table 2.** Parameters of selected model.

Parameter	Setting
Training Dataset	B0
Training image size	800
Model size	Small
Augmentation	yes
Dataset adjustment	Not applied
Recall	0.48
Precision	0.40
F2-Score	0.41
Epochs	137
Batch	16
Confidence	0.4
SAHI window size	800x600

positive prediction rate. Given these findings, we select the model shown in **Table 2** for further processing.

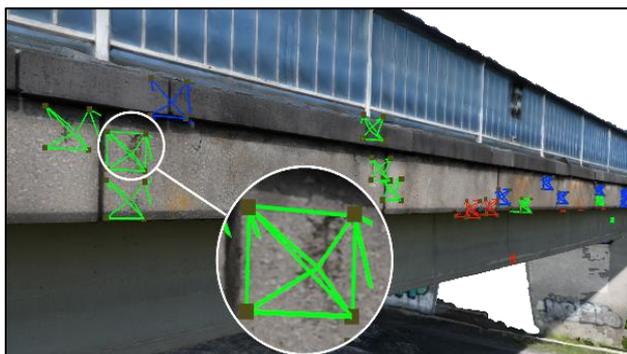
### 4.3 Spatial analysis

After deploying the selected damage detection model on the inspection flight data, we use our method as explained in 3.2 to map the damages onto the mesh as shown in **Fig. 8**. Concretely, we use the chosen model to predict damages on 121 images from an inspection flight of approx. 5 m distance to the pier cap, resulting in 474 predictions. The proposed spatial filtering procedure successfully filters out 17 damages with a smaller area than 1.6 cm<sup>2</sup> and 6 false predictions with an area larger than 1.5 m<sup>2</sup>. Filtering out the false positive predictions and insignificant annotations, results in an increased F2-score of 0.60, with a successful recall of 0.93. This result underscores the significance of spatial filtering techniques in automated damage detection procedures. Finally, we analyse the missed detections utilizing the damage area predicate of the RDF.

To visualize the final detection results, we use our GUI as shown in **Fig. 9**. With a manual effort of approx. 15 minutes for the remaining 57 instances, the last false positive predictions are filtered, and the annotations corrected, allowing to use the inspection images as training data to improve the detection model. Additionally, we were able to approximately measure the size of the detected damages. Lastly, we include the final results in a report, detailing damage dimensions, location and damaged bridge component.

**Table 3.** Results of spatial analysis.

Parameter	Number of predictions	Ground Truth
Total predictions	474	26
Not intersecting mesh	12	0
Remaining after grouping	80	26
Remaining after spatial filter	57	15



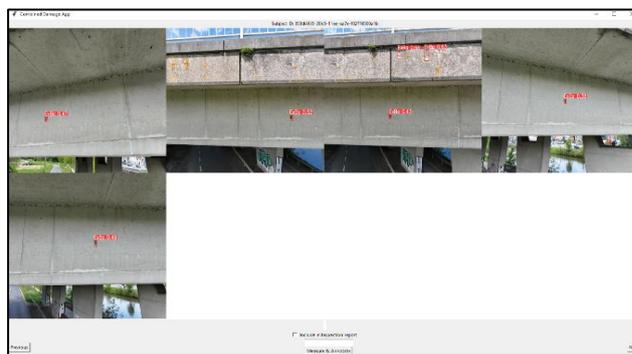
**Fig. 8.** Damages mapped onto the rendered 3D model. The boxes represent the true positive (green), false positive (blue) and false negative (red) predictions. The detail shows the precise mapping of an exposed bar.

## 5. DISCUSSION

In the following section we discuss the results of our experiments in the larger context of the overall framework and reflect.

The initial phase of our project involves photogrammetric processing to create highly accurate and geolocated dense point clouds. Impressively, these precise models were achieved without the reliance on GCPs, which traditionally constitute a significant portion of photogrammetry efforts. We observed no systematic deformations due to errors in camera orientation, countering previous studies' findings, such as those reported by Štroner et al. (2020). However, the resulting mesh models were less accurate than the evaluated dense point clouds, owing to the use of low photogrammetric settings for numerous model reconstructions. This discrepancy can potentially be rectified by increasing the computed point count. Importantly, the high-quality model texture allows for visual verification of the detection and mapping results from the automated pipeline, affirming the accuracy of our methodology (Fig 8). Furthermore, the procedure of aligning new images e.g., from new inspection epochs, to the existing project, allows to retrieve highly accurate camera positions even in GPS restricted environments. The ability to accurately register new inspection flights is crucial in the context of our overall framework of recurring UAS bridge inspections, reinforcing the robustness of our methodology. All in all, RTK-UAS photogrammetry shows great potential in surveying and digitalizing bridges, especially for sites where TLS is difficult due to spatial constrictions, such as very high altitudes, inaccessible sites or high traffic on site.

In evaluating our trained damage detection models, we note a substantial performance discrepancy between the validation dataset and the UAS-derived evaluation dataset. At the same time, utilizing large amounts of the available training data from conventional bridge inspections improves the models' generalization ability. This discrepancy can potentially be reduced by incorporating annotated UAS inspection data into the validation set and iteratively retraining the model. Our GUI plays a pivotal role in this iterative refinement process. It enables bridge inspectors to effortlessly annotate UAS inspection data with each conducted inspection, contributing to a progressively enriched training dataset that bolsters model performance. Furthermore, the addition of unannotated 'negative' examples, such as joints misinterpreted as exposed bars, can aid in fine-tuning the model's predictive accuracy. Collectively, these findings underscore the crucial role of meticulously curated,



**Fig. 9.** GUI, showing the detection results of a single exposed bar in different detection images, allowing to approximate damage dimensions and create new training data.

annotated UAS inspection data, and our GUI's utility in improving the precision of bridge damage detection.

Our case study highlights the immense potential of damage mapping and spatial filtering for more effective UAS bridge inspections. By grouping filtering the predictions onto the 3D model we can significantly decrease the amount of false positive predictions, reducing the final manual effort of the inspectors. The ray casting method for the mapping process is only based on a simply pinhole camera model, but still delivers mostly accurate mapping results. However, especially for damages close to edges, the mapping accuracy could benefit from more complex projection procedures, such as including a lens distortion. Lastly, it is important to note that the case study encompasses only a small amount of ground truth damages, and therefore larger studies need to be conducted to verify the results.

## 6. CONCLUSIONS

This study presents a robust framework to facilitate semi-automated, recurrent UAS-based bridge inspections. Concretely, this work demonstrates the potential of highly accurate photogrammetric models solely based on UAS-RTK imagery, automatically detect bridge damage and map the damages onto the texturized 3D bridge model. We integrate these capabilities into a user interface that incorporates an expert-in-the-loop, thus enabling quick verification of results and reducing false-positive detections. This system not only generates additional training data for UAS image analysis but also facilitates approximate measurements of bridge damage directly from the images.

Our methodology for achieving precise photogrammetric models exclusively from UAS-RTK imagery is validated through a detailed case study. The results reveal remarkable geometric accuracy, demonstrating an average deviation of 1.3 cm and a standard deviation of 1.8 cm. while outperforming traditional TLS methods in both feasibility and efficiency. The trained YOLOv8 damage detection models are able to detect exposed bars with a larger area than 1.6 cm<sup>2</sup> with a recall rate of 93%. The predictions are successfully mapped onto the bridge model, allowing to compute position and area of the damages. The proposed spatial filtering method proves to be effective in eliminating 29% false positive predictions, further reducing manual efforts in the inspection. Finally, we visualize the results in a user interface and create 57 newly annotated training images to further improve the damage detection model.

The results presented in this study demonstrate significant promise to facilitate UAS-based bridge inspections, making use of high-quality inspection images from otherwise hard-to-reach areas. Nevertheless, the overall framework stands to gain from further refinements, especially in the stages that are not part of this work, such as automated flight routes and an automated abstraction of damage characteristics. Continuing the outcomes of this work, we will evaluate the proposed integration of the newly annotated UAS inspection images. Furthermore, we propose refining the spatial filtering method by leveraging additional information from our ontology, e.g., the number of predictions per damage and mean confidence. This could help in better bounding box prediction and effective damage detection. As we continue to refine our framework, our focus will be on implementing the suggested improvements and conducting larger-scale bridge inspections, assessing the scalability of the method and identify possible limitations. The potential impact of this work is vast, suggesting a promising future for UAS-based bridge inspections.

**Author Contributions:** E.B. is the main author of the work and conceived the method. M.B. is the direct supervisor. M.V. is the supervisor. All authors have read and agreed to the published version of the manuscript.

**Conflicts of Interest:** There are no conflicts of interest to report.

## REFERENCES

- Akyon F.C., Onur Altinuc S., Temizel A. (2022). Slicing Aided Hyper Inference and Fine-Tuning for Small Object Detection. In: 2022 IEEE International Conference on Image Processing (ICIP). IEEE, pp 966–970
- Bartczak E.T., Bassier M., Vergauwen M. (2023) CASE STUDY FOR UAS-ASSISTED BRIDGE INSPECTIONS. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.* XLVIII-2/W3-2023:33–39. <https://doi.org/10.5194/isprs-archives-XLVIII-2-W3-2023-33-2023>
- Bukhsh ZA, Jansen N, Saeed A. (2021). Damage detection using in-domain and cross-domain transfer learning. *Neural Comput & Applic* 33:16921–16936. <https://doi.org/10.1007/s00521-021-06279-x>
- Chen S., Laefer D.F., Mangina E., Zolanvari S.M.I., Byrne J. (2019). UAV Bridge Inspection through Evaluated 3D Reconstructions. *J. Bridge Eng.* 24. [https://doi.org/10.1061/\(ASCE\)BE.1943-5592.0001343](https://doi.org/10.1061/(ASCE)BE.1943-5592.0001343)
- Ding W., Yang H., Yu K., Shu J. (2023). Crack detection and quantification for concrete structures using UAV and transformer. *Automation in Construction* 152:104929. <https://doi.org/10.1016/j.autcon.2023.104929>
- Geomapi Version 0.3.2. (2023). Retrieved March 20, 2023 from <https://geomatics.pages.gitlab.kuleuven.be/research-projects/geomapi>
- Hallermann N., Morgenthal G. (2016). From Aerial Photography to 3-Dimensional Inspection of Bridges. In: IABSE Conference, Guangzhou 2016: Bridges and Structures Sustainability - Seeking Intelligent Solutions. International Association for Bridge and Structural Engineering (IABSE) Zurich, Switzerland, pp 546–553
- James M.R., Robson S. (2014). Mitigating systematic error in topographic models derived from UAV and ground-based image networks. *Earth Surf. Process. Landforms* 39:1413–1420. <https://doi.org/10.1002/esp.3609>
- Jocher G., Chaurasia A., Qiu J. (2023). YOLOv8. YOLO by Ultralytics. <https://github.com/ultralytics/ultralytics>
- Liang H., Lee S.-C., Seo S. (2023). UAV-Based Low Altitude Remote Sensing for Concrete Bridge Multi-Category Damage Automatic Detection System. *Drones* 7:386. <https://doi.org/10.3390/drones7060386>
- Lin J.J., Ibrahim A., Sarwade S., Golparvar-Fard M. (2021). Bridge Inspection with Aerial Robots: Automating the Entire Pipeline of Visual Data Capture, 3D Mapping, Defect Detection, Analysis, and Reporting. *J. Comput. Civ. Eng.* 35. [https://doi.org/10.1061/\(ASCE\)CP.1943-5487.0000954](https://doi.org/10.1061/(ASCE)CP.1943-5487.0000954)
- Mittal P., Singh R., Sharma A. (2020). Deep learning-based object detection in low-altitude UAV datasets: A survey. *Image and Vision Computing* 104:104046. <https://doi.org/10.1016/j.imavis.2020.104046>
- Roboflow. (2023). Retrieved March 20, 2023, from <https://roboflow.com>
- Rock G., Ries J.B., Udelhoven T. (2011). Sensitivity analysis of UAV-photogrammetry for creating digital elevation models (DEM). *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.* XXXVIII-1/C22:69–73. <https://doi.org/10.5194/isprsarchives-XXXVIII-1-C22-69-2011>
- Štroner M., Urban R., Reindl T., Seidl J., Brouček J. (2020). Evaluation of the Georeferencing Accuracy of a Photogrammetric Model Using a Quadcopter with Onboard GNSS RTK. *Sensors* (Basel) 20. <https://doi.org/10.3390/s20082318>
- Štroner M., Urban R., Seidl J., Reindl T., Brouček J. (2021). Photogrammetry Using UAV-Mounted GNSS RTK: Georeferencing Strategies without GCPs. *Remote Sensing* 13:1336. <https://doi.org/10.3390/rs13071336>
- Toriumi F.Y., Bittencourt T.N., Futai M.M. (2021). UAV-based inspection of bridge and tunnel structures: an application review. *Rev. IBRACON Estrut. Mater.* 16. <https://doi.org/10.1590/s1983-41952023000100003>
- Tscharf A. (2020). UAV-gestützte Vermessung im Bergbau – Zur Frage der Genauigkeit unter Verwendung von Structure from Motion. *Berg Huettenmaenn Monatsh* 165:274–283. <https://doi.org/10.1007/s00501-020-00988-x>