

Detecting Vietnam War bomb craters in declassified historical KH-9 satellite imagery

Philipp Barthelme^{a,*}, Eoghan Darbyshire^b, Dominick V. Spracklen^c, Gary R. Watmough^{a,d,e}

^a School of Geosciences, University of Edinburgh, Edinburgh, United Kingdom

^b Conflict and Environment Observatory, Mytholmroyd, United Kingdom

^c School of Earth and Environment, University of Leeds, Leeds, United Kingdom

^d Global Academy of Agriculture and Food Systems, University of Edinburgh, United Kingdom

^e Novel Data Ecosystems for Sustainability (NODES) Research Group. Advancing Systems Analysis Program, International Institute for Applied Systems Analysis (IIASA), Austria

ARTICLE INFO

Keywords:

Unexploded ordnance
Bomb crater
Satellite imagery
KH-9 Hexagon
Vietnam War
Mine action

ABSTRACT

Thousands of people are injured every year from explosive remnants of war which include unexploded ordnance (UXO) and abandoned ordnance. UXO has negative long-term impacts on livelihoods and ecosystems in contaminated areas. Exact locations of remaining UXO are often unknown as survey and clearance activities can be dangerous, expensive and time-consuming. In Vietnam, Lao PDR and Cambodia, about 20% of the land remains contaminated by UXO from the Vietnam War. Recently declassified historical KH-9 satellite imagery, taken during and immediately after the Vietnam War, now provides an opportunity to map this remaining contamination. KH-9 imagery was acquired and orthorectified for two study areas in Southeast Asia. Bomb craters were manually labeled in a subset of the imagery to train convolutional neural networks (CNNs) for automated crater detection. The CNNs achieved a F1-Score of 0.61 and identified more than 500,000 bomb craters across the two study areas. The detected craters provided more precise information on the impact locations of bombs than target locations available from declassified U.S. bombing records. This could allow for a more precise localization of suspected hazardous areas during non-technical surveys as well as a more fine-grained determination of residual risk of UXO. The method is directly transferable to other areas in Southeast Asia and is cost-effective due to the low cost of the KH-9 imagery and the use of open-source software. The results also show the potential of integrating crater detection into data-driven decision making in mine action across more recent conflicts.

1. Introduction

Unexploded ordnance (UXO) refers to explosive munitions, including bombs, artillery projectiles and cluster submunitions that have been deployed during military conflicts but did not explode. UXO continues to present significant humanitarian and environmental challenges. In 2022 alone, the United Nations Mine Action Service (UNMAS) reported more than 3000 casualties from explosive remnants of war, which include UXO and abandoned explosive ordnance, across 15 countries and numbers of UXO are increasing due to ongoing conflicts such as in Ukraine (Cluster Munition Coalition, 2023; UNMAS, 2019). UXO has negative long-term impacts on public health, livelihoods and ecosystems in contaminated areas (Frost et al., 2017; Hofmann and Juergensen, 2017; Lin et al., 2020; Nguyen, 2020; Ounmany and Andriess, 2018).

Moreover, the removal of UXO remains technically challenging, expensive and hazardous, particularly in conflict and post-conflict environments where access to reliable data on contamination is limited.

Mainland Southeast Asia has one of the highest UXO contamination rates in the world, mainly originating from the aerial bombardment by the U.S. military during the Vietnam War, also known as the American War in Vietnam or the Second Indochina War which took place between 1955 and 1975 (Martin et al., 2019). During the war, the U.S. Air Force dropped approximately eight million tons of bombs on the countries of Vietnam, Cambodia and Lao PDR (Anderson, 2002; High et al., 2013). Today, about 20% of the land in these countries is thought to still be contaminated by UXO (Martin et al., 2019). However, the exact locations and extents of contaminated areas mostly remain unknown, despite being essential for an efficient allocation of limited resources for

* Corresponding author.

E-mail address: philipp.barthelme@ed.ac.uk (P. Barthelme).

<https://doi.org/10.1016/j.srs.2024.100143>

Received 4 April 2024; Received in revised form 31 May 2024; Accepted 6 June 2024

Available online 7 June 2024

2666-0172/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

UXO clearance. Non-technical survey is commonly used as a first step when assessing explosive ordnance contamination. It typically consists of a set of non-technical methods, such as desk assessments, physical visits to field locations and the analysis of historical records to identify contaminated land and categorize it into suspected or confirmed hazardous areas. It is cheaper than technical survey and clearance, as it does not rely on expensive technical assets to be deployed to the field. An accurate non-technical survey can therefore ensure a more efficient allocation of these limited technical assets to the most affected areas (Bold and Avenell, 2021; Lin et al., 2020; UNMAS, 2019).

U.S. bombing records are one of the most comprehensive data sources used for non-technical survey in Southeast Asia. In 2016, the United States Department of Defense released these records to the public as part of the Theater History of Operations (THOR) data, an attempt to record all air operations by the United States since World War I. The THOR data includes the geographical coordinates of target locations, the type and number of weapons dropped on each target and the time of the attack. The bombing records have been a valuable data source for non-technical surveys (Bold and Avenell, 2021) and for research into the political, economic and health impacts of the Vietnam War (Le et al., 2022; Le and Nguyen, 2020; Yamada and Yamada, 2021). However, High et al. (2013) suggest the bombing data should only be used as one source among many, after identifying multiple issues, including missing, corrupted and actively falsified records. An overview of THOR bombing targets in Southeast Asia during the Vietnam War is shown in Fig. 1a.

Remote sensing data can provide a valuable alternative data source where bombing records are unavailable or inaccurate (Bennett et al., 2022). Lin et al. (2020) used recent, very high resolution (<1 m) satellite imagery to detect bomb craters from the Vietnam War in Cambodian agricultural land. However, detecting bomb craters from past conflicts

in more recent satellite images can be challenging as the appearance of bomb craters changes over time due to erosion, vegetation growth and human intervention (Lin et al., 2020). Historical aerial wartime imagery has been used as an alternative to detect and analyze World War II bomb craters in Europe (Clermont et al., 2019; Kruse et al., 2019; Waga et al., 2022), but its availability is often restricted to small areas. Declassified historical U.S. satellite imagery (USGS EROS Center, 2018), taken during and immediately after the Vietnam War, now presents an opportunity to overcome some of these challenges. The KH-4a/b CORONA missions provide high resolution imagery (1.8–2.8 m) between 1963 and 1972 which, since its declassification in 1995, has been used in a variety of applications that range from the discovery of archaeological sites to land cover change detection (Deshpande et al., 2021; Lasaponara et al., 2018; Nita et al., 2018). Recently, it was used to classify land affected by bombing in a part of Quang Tri province, Vietnam (Munteanu et al., 2024). The KH-9 HEXAGON stereo-panoramic imagery provides almost complete coverage of the Earth’s land area between 1971 and 1984 at a spatial resolution of 0.6–1.2 m. Due to its recent declassification in 2011 and the technical challenges associated with orthorectifying the imagery (Zhou et al., 2021), researchers have only recently begun to explore its use in a diverse range of applications such as archaeology (Hammer et al., 2022) and glaciology (Ghuffar et al., 2023).

Previous studies on the automatic detection and counting of bomb craters in remotely sensed imagery have relied on methods developed for detecting extra-terrestrial craters on planetary surfaces (Clermont et al., 2019; Lin et al., 2020). In this field, convolutional neural networks (CNNs) are increasingly replacing applications that rely on the extraction of manually specified features such as crater shape and shadows. U-Nets, a type of CNN architecture originally developed for segmenting medical imagery (Ronneberger et al., 2015), have been successfully

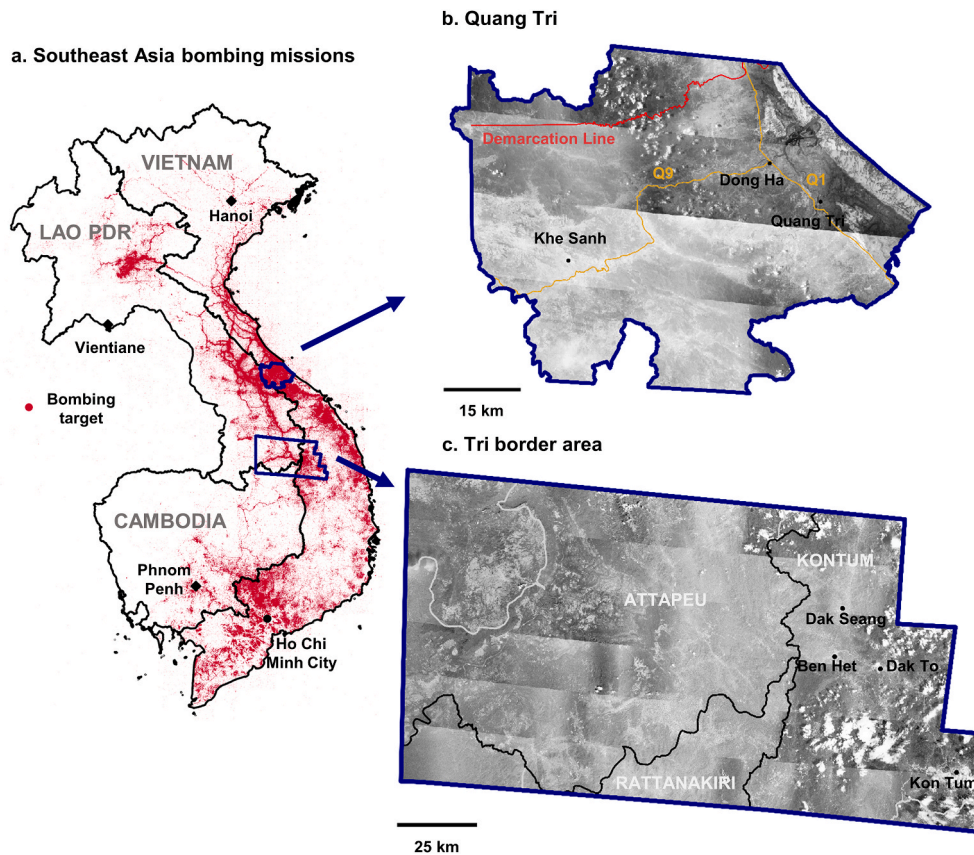


Fig. 1. (a) THOR bombing targets over Southeast Asia during the Vietnam War. (b) and (c) show the two study areas, including points of interest during the war. The cities of Dong Ha, Quang Tri, Dak To and Kon Tum were locations of larger military bases whereas Khe Sanh, Dak Seang and Ben Het contained smaller military camps.

applied to segment extra-terrestrial craters (Chen et al., 2023; Silburt et al., 2019) and more recently to detect artillery craters in Ukraine (Duncan et al., 2023). To achieve instance segmentation, a method for identifying individual instances of an object, methods are often adjusted by introducing a boundary class (Duncan et al., 2023) or by using template matching on the semantic segmentation product (Chen et al., 2023; Silburt et al., 2019).

In this study, we investigated whether Vietnam War-era bomb craters can be automatically detected in declassified KH-9 imagery using machine learning methods and we analyzed whether these detected craters add value to non-technical surveys in Southeast Asia beyond existing data sources. Additionally, we hypothesized that variations in crater appearance in the imagery affect model performance, which could lead to biased crater predictions, thereby impacting prioritization of clearance activities. Our study was structured in the following way. First, we acquired KH-9 imagery for two study areas in Southeast Asia and orthorectified the imagery using open-source tools. We manually labeled craters on a subset of the imagery. To ensure objective labeling we defined key appearance characteristics of craters visible in the imagery which we used to categorize each crater. This crater categorization was later used to identify any biases in model performance. To automate crater detection, we used an instance segmentation workflow using CNNs with a U-Net architecture. Model performance was analyzed, revealing differences by crater appearance. Finally, we compared detected crater locations to U.S. bombing records, identifying multiple issues with the bombing records in the process. Our results show that craters visible in the KH-9 imagery provide more precise information about where bombs landed than target locations of declassified bombing records. Additionally, our findings demonstrate how methods to automatically detect these craters can improve data-driven decision making within the mine action sector in Southeast Asia.

2. Materials and methods

2.1. Study areas

Two study areas across Southeast Asia were selected (Fig. 1). The first study area covers a total of 4148 km² of Quang Tri (QT) province, the most heavily bombed province in Vietnam during the war (Miguel and Roland, 2011), as it contained the 17th parallel, the dividing line between North and South Vietnam at the time. The second study area, here referred to as the tri-border area (TBA), is located around the meeting point of the borders of Vietnam, Lao PDR, and Cambodia. Encompassing 17,285 km² of predominantly mountainous and densely vegetated land, the TBA contained sections of the Ho Chi Minh Trail, the principal supply route for the North Vietnamese Army, including a vital entry point of the trail into South Vietnam in Kon Tum province. The KH-9 images were taken on November 4, 1972 (TBA) and on March 20, 1973 (QT province).

2.2. Processing the KH-9 imagery

A total of 20 KH-9 images, forming 10 stereo pairs of forward and aft looking cameras, were used for the study. The U.S. Geological Survey provided photogrammetric film scans of the archived KH-9 film sources at a resolution of 7 μm and a cost of 30\$ per image. Previously digitized images, now including all images used in this study, are available at no cost via the Earth Explorer platform.

The film scans were provided in multiple sections and were not georeferenced. The open-source Nasa Ames Stereo Pipeline (ASP) (Beyer et al., 2021) was used to process and orthorectify the imagery. The ASP implements a rigorous camera model including motion compensation (Sohn et al., 2004) for the panoramic cameras used by the KH-9 satellites. We adapted the example workflow described in section 8.26 of the ASP manual (Beyer et al., 2021), as we integrated manual ground control points (GCPs) to improve accuracy.

First, image parts were stitched together and cropped to the image extent using the *image_mosaic* and *historical_helper* tools in ASP. QGIS (QGIS Association, 2023) and Google Earth imagery were used to identify approximately 15 ground control points (GCPs) per image. The GCPs, in combination with locations of image corners provided in the image metadata, were used to initialize intrinsic and extrinsic camera parameters for each individual image. In a next step, tie points between the images of each stereo pair were computed. The tie points, in combination with the GCPs, were used to optimize the initial camera parameters using a joint bundle adjustment for each stereo pair. Each image was subsequently orthorectified by projecting it onto a digital elevation model (NASA Shuttle Radar Topography Mission (SRTM), 2013) using the optimized camera parameters and the *mapproject* tool in ASP. While the original image resolution can change between and within images, to ensure consistency during crater prediction, all images were orthorectified at the same resolution of 1 m per pixel.

The resulting images were cropped to the study areas introduced in Section 2.1. The QT imagery was mosaicked into one image, while the TBA images, being larger in size, were not mosaicked. A further 60 validation GCPs (QT province: 20, TBA: 40) were collected for validation of the orthorectification process and showed a mean absolute horizontal error of 7.0 m (25th percentile: 3.4 m, median: 5.8 m, 75th percentile: 8.8 m) for QT province and 17.5 m (25th percentile: 8.6 m, median: 13.7 m, 75th percentile: 21.4 m) for the TBA.

2.3. Labeling of bomb craters

The processed KH-9 imagery was divided into image tiles with a width and height of 256 pixels. From the QT imagery, 1000 random image tiles were chosen and divided into sets of 600 for training, 200 for validation, and 200 for testing. From the TBA imagery, 1400 tiles were selected, with 600 allocated for training, 200 for validation, and 600 for testing. The decision to increase the number of test tiles for the TBA was driven by its lower density of bomb craters, aiming to ensure a more representative test score.

Craters visible in the selected image tiles were manually labeled if they were larger than 25 pixels (equivalent to 25 m²). Smaller ground features were excluded as they were difficult to reliably identify given the image resolution and quality. The threshold of 25 pixels was selected based on visual inspection. Each labeled crater was assigned one of five classes based on its appearance in the imagery which varied substantially (Fig. 2).

Labeling proved particularly challenging in mountainous areas with heavy vegetation and in areas featuring houses, trees or graves that could resemble craters in the imagery. Where necessary, the context visible in the KH-9 and current satellite imagery was used to make a better-informed decision. Notably, the crater prevalence was much lower for the TBA where 964 craters were identified compared to 10,132 craters in QT province. Additional details on the crater labeling and the different crater classes are provided in the Supplementary Materials.

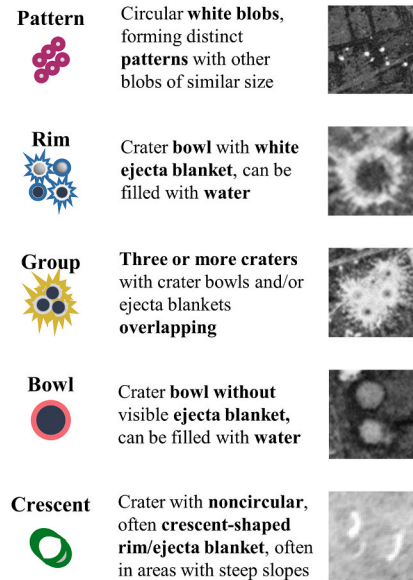
2.4. Detection of bomb craters

An instance segmentation workflow was used to predict individual craters in the imagery. The instance segmentation was implemented as a semantic segmentation problem by adding a boundary class, an approach commonly used in biomedical applications such as nucleus segmentation (Caicedo et al., 2019), where large amounts of densely packed objects have to be separated. For this approach, the area of each labeled crater was expanded by two pixels which were assigned to the new boundary class. All crater pixels that were touching neighboring craters were also labeled as boundary pixels.

2.4.1. Neural network architecture and training

A U-Net with a Resnet50 backbone, pre-trained on the Imagenet dataset, was used for the segmentation (Deng et al., 2009; He et al.,

a. Overview crater types



b. Crater types in context




Fig. 2. (a) Different crater types defined based on their appearance characteristics in the KH-9 imagery. (b) Examples of the crater types in context for an area in Quang Tri province.

2015; Ronneberger et al., 2015). While multiple improvements to the standard U-Net architecture have been suggested, in most settings they only lead to minor or no accuracy improvements at a larger computational cost (Gut et al., 2022; Kugelman et al., 2022; Wang and Miao, 2022). Therefore, instead of comparing different model architectures, the analysis in this paper focuses on different bomb crater appearances, which have a large impact on model accuracy, and the comparison of detected craters with historical bombing records.

The model was implemented using the *pytorch* and *segmentation_models_pytorch* packages in Python (Iakubovskii, 2019; Paszke et al., 2019). An initial model was trained using data from both study areas before the model was fine-tuned for each study area independently, using only the training data for the respective study area. Min-max scaling was applied to individual image tiles. During model training, the images were augmented by applying random vertical and horizontal flips as well as random brightness and contrast adjustments. As the pre-trained model expected color images with three channels as input, whereas the KH-9 images are grayscale with a single channel, an additional layer was added in front of the pre-trained model to map from one to three channels.

The models were trained on a Nvidia RTX 2060 GPU using an Adam optimizer (Kingma and Ba, 2014) and a focal loss function (Lin et al., 2017), which assigned more weight to training examples that were not well classified. Focal loss has been shown to work well for imbalanced data (Lin et al., 2017; Mulyanto et al., 2021) which was a problem here as more than 99% of all labeled pixels were background pixels. A focal loss alpha value of 1 was used for the background class and 3 for all other classes based on the model performance on the validation images. A batch size of 8 and a learning rate of 1e-3 was used during initial model training and the learning rate was reduced to 1e-5 for the fine-tuning of each study area. Early stopping was used to stop model training if the validation loss did not decrease for 50 epochs in a row.

2.4.2. Semantic segmentation evaluation

The segmentation results were evaluated using a pixel-to-pixel comparison on the test images. We used precision, recall and F1-score which are commonly applied in settings of class imbalance and which are defined as:

$$Precision = \frac{TP}{TP + FP} \quad (1)$$

$$Recall = \frac{TP}{TP + FN} \quad (2)$$

$$F_1 = \frac{2TP}{2TP + FN + FP} \quad (3)$$

where TP denotes true positives, FN denotes false negatives and FP denotes false positives. We calculated these metrics for each individual crater class, and additionally calculated one combined score that only considers whether a pixel had been correctly identified as a crater pixel, even if the crater class of predicted and labeled pixels differed.

2.4.3. Instance segmentation

Multiple post-processing steps were applied to transform the semantic segmentation output into individual crater instances (Fig. 3). Connected crater pixels were considered as one crater instance even if they belonged to different crater classes. Pixels of class *Boundary* were treated as background pixels at this stage. Each predicted crater instance was assigned the majority class of its pixels. All predicted crater instances smaller than 25 pixels were removed.

The accuracy of crater instances was evaluated using the metrics described in Section 2.4.2. A predicted crater (A) was considered correct if it had an *Intersection over Union (IoU)* of 0.5 or more with a labeled crater (B), where *IoU* is defined as:

$$IoU(A, B) = \frac{A \cap B}{A \cup B} \quad (4)$$

Accuracy scores were calculated for each individual crater type and for a combined crater class that only considered whether a crater instance had been correctly identified even if the crater class of the predicted and labeled craters differed.

2.4.4. Model prediction

The trained models were applied to the entire study areas using a sliding window approach with an overlap of 64 pixels. Only the center 192×192 pixels of each predicted 256×256 image tile were retained to avoid artefacts and improve performance at tile edges. When

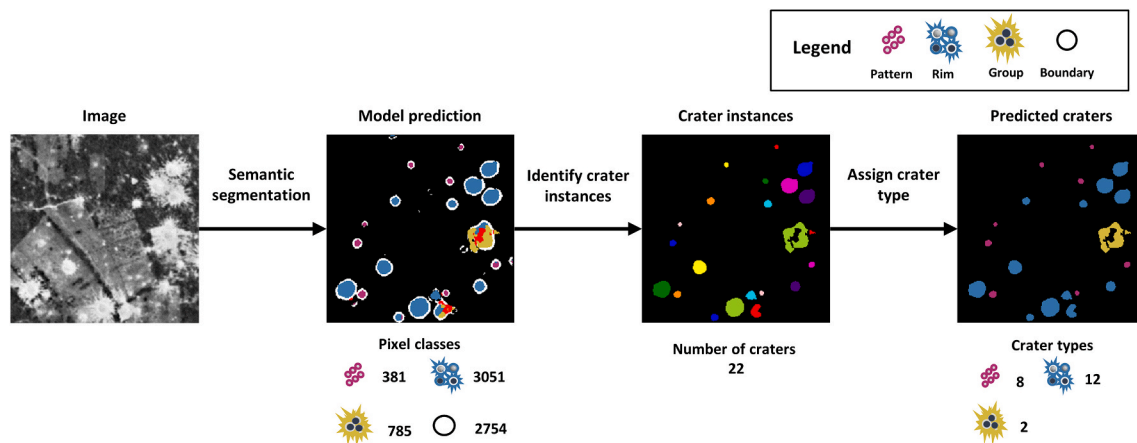


Fig. 3. Crater prediction and post-processing workflow.

identifying individual crater instances on the predicted segmentation masks the tile size of 1024×1024 with an overlap of 512 pixels was used to avoid mistakenly separating large craters that crossed one or more image tiles. As the images, and therefore crater predictions, in the second study area were overlapping, we only kept predicted craters of one of the images in these overlapping areas. This was not necessary for QT province, where the KH-9 images were mosaicked before crater detection, resulting in the same outcome.

2.5. THOR bombing data

The THOR bombing data was used as a comparison for the crater predictions as declassified bombing records are one of the main data sources currently used for non-technical survey in Southeast Asia. Multiple issues with the THOR bombing data were identified and addressed before the analysis. The target coordinates were labeled as using the WGS84 datum while analysis suggested they were provided in the Indian 1960 datum. We transformed the coordinates from EPSG:4131 to EPSG:4326 which led to a shift of about 500 m. Moreover, B-52 bombing missions were counted twice from 1971 onwards as they seemed to be included in both the SACCOACT and SEADAB databases which were both used for THOR. Therefore, all B-52 bombing missions originating from the SACCOACT database that occurred after 1971 were removed. The analysis was limited to large aircraft bombs which would result in craters larger than 25 m^2 , and the provided weapon type information was used to identify the corresponding records. Where no information on weapon type was available, as was the case for a large amount of SEADAB records, the weapon type was imputed based on the provided weight of the weapon used. More details on the identified issues and applied corrections, including checks to the robustness of our results based on alternative processing, are provided in the Supplementary Materials (Text S2, Figs. S1–S2, Tables S2–S5).

All records of bombing that occurred after the respective KH-9 images were taken were dropped. The resulting records are referred to as *total bombing*. The data were further split into (1) bombs dropped within the year before the respective KH-9 images were taken (*previous year bombing*) and (2) bombs dropped more than a year before the imagery was taken (*bombing before previous year*). The resulting numbers of bombs dropped were directly compared to the number of detected craters in each study area. Additionally, aggregated counts of detected craters and bombs dropped for grid cells of various cell sizes between 100 m and 4 km were compared using the Spearman correlation coefficient r (Schober et al., 2018). To allow for a direct comparison between the number of detected craters and the number of bombs dropped during *previous year bombing*, excluding the influence of older craters, a distinct analysis was undertaken. This analysis focused on grid cells ($2 \text{ km} \times 2 \text{ km}$) within QT province, where *previous year bombing* constituted at least

90% of *total bombing*.

3. Results and discussion

3.1. Model evaluation

The trained models achieved an F1-Score of 0.61 (precision: 0.67, recall: 0.56) when predicting craters of all types across the test sets and predicted a total of 541,398 craters (QT: 442,157, TBA: 99,241) across the full study areas (Fig. 4). The model performance differed between the two study areas with an F1-Score of 0.64 for QT province and 0.44 for the TBA (Table 1). The model performed better in image tiles with a high number of predicted craters with an F1-Score of 0.66 (QT: 0.66, TBA: 0.71) compared to an F1-Score of 0.32 (QT: 0.43, TBA: 0.25) in images with a low number of predicted craters (Table S6). We present detailed metrics by study area and crater types in Table 1. Most of the predicted craters were of type *Pattern* (QT: 229,467, TBA: 67,985), *Rim* (QT: 91,112, TBA: 9995) and *Crescent* (QT: 71,364, TBA: 16,836). The model only predicted a small number of craters of type *Group* (QT: 9,645, TBA: 46) and *Bowl* (QT: 40,569, TBA: 4379). Fig. 5 shows the detected crater locations by crater type for the QT study area.

3.1.1. Model performance across study areas

Model performance is comparable between image tiles of both study areas with a high number of predicted craters (Table S6). The difference in F1-Scores for the two study areas is therefore likely due to the lower prevalence of craters in the TBA where only 1 in 1434 pixels were crater pixels compared to 1 in 90 for QT province. The lower prevalence results in a smaller number of labeled craters and a higher influence of every false positive crater prediction on the evaluation metrics, which has been identified as a challenge in previous research on bomb crater detection (Clermont et al., 2019; Lin et al., 2020). As a random sample of images was used in each study area, the test data had a realistic class distribution, and the results for the TBA reflect the difficulty of predicting bomb craters over large, mostly unaffected areas. Further, land cover can influence the accuracy of predictions; in the TBA, land cover mostly consisted of heavily vegetated land and mountainous terrain, with only small amounts of agricultural land in which bomb craters are generally easier to identify and segment (Duncan et al., 2023).

The lower F1-Score in areas with lower crater prevalence is less problematic when trying to identify areas with a high likelihood of unexploded bombs, which tend to be areas with a high number of bomb craters. When considering crater predictions as correct if any overlap with a labeled crater exists ($\text{IOU} > 0$), which is still sufficient to identify risk areas, the F1-Score increased to 0.70 (QT: 0.72, TBA: 0.53). Additionally, the model favored precision over recall (Table 1), which also holds for low prevalence areas (Table S6). This reduced false positive

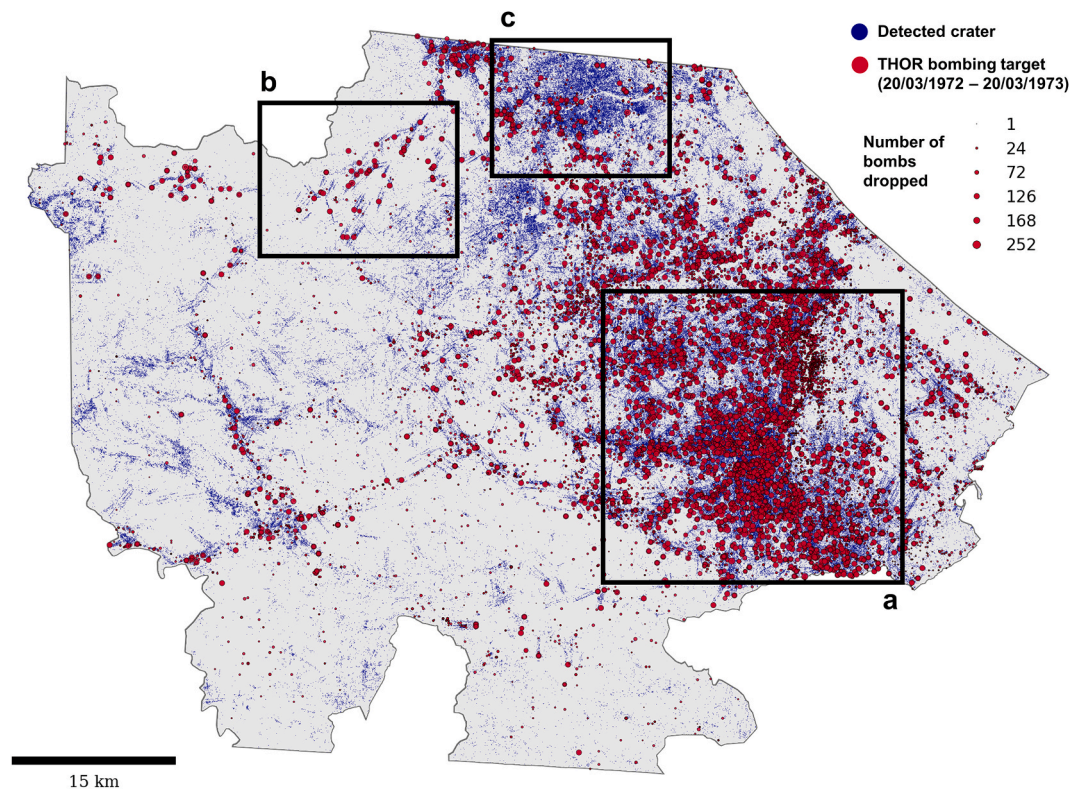


Fig. 4. Comparison of predicted bomb craters (blue) and THOR bombing targets (red) during the year preceding the KH-9 image acquisition in Quang Tri province. (a) Shows a high density of bomb craters and bombs dropped close to Quang Tri city. (b) Shows multiple lines of bomb craters matching B-52 bombing targets recorded in THOR. (c) Shows an area with large amounts of craters but little bombing during the year before the KH-9 images were taken, indicating the craters originated from earlier in the war. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

Table 1

Bomb crater detection results showing F1-score (precision/recall) and the number of labeled craters N in the test data. For the Craters category, all crater classes are considered as one combined crater class.

	Craters	Pattern	Rim	Group	Crescent	Bowl	Boundary	Background
Quang Tri								
Pixels	0.65 (0.70/0.61)	0.63 (0.62/0.64)	0.59 (0.60/0.58)	0.13 (0.45/0.07)	0.35 (0.31/0.40)	0.27 (0.35/0.22)	0.45 (0.44/0.46)	0.99 (0.99/0.99)
Craters (IOU >0.5)	0.64 (0.68/0.60)	0.70 (0.68/0.73)	0.55 (0.58/0.52)	0.07 (0.33/0.04)	0.30 (0.25/0.37)	0.24 (0.30/0.20)	–	–
	N=1712	N = 748	N = 449	N = 247	N = 111	N = 157		
Tri-border area								
Pixels	0.41 (0.61/0.31)	0.53 (0.56/0.50)	0.38 (0.67/0.27)	0.00 (0.00/0.00)	0.17 (0.21/0.14)	0.06 (0.47/0.03)	0.29 (0.39/0.23)	1.00 (1.00/1.00)
Craters (IOU >0.5)	0.44 (0.58/0.35)	0.58 (0.59/0.57)	0.37 (0.64/0.26)	0.00 (0.00/0.00)	0.20 (0.25/0.17)	0.03 (0.20/0.02)	–	–
	N=314	N = 142	N = 54	N = 15	N = 36	N = 67		

predictions in areas where no bombing happened.

3.1.2. Model performance across crater types

Model performance varied between the different crater types (Table 1), with higher F1-scores for craters of type *Pattern* and *Rim* compared with *Group*, *Crescent* and *Bowl*. This could be attributed to the lower prevalence for these crater types, resulting in fewer training data for the model to learn from. However, we also identified issues with some of the crater types that are unlikely to be resolved by increasing the training data. Understanding these limitations is important as crater appearance depends on spatial characteristics such as land cover and terrain which could introduce spatial bias into the crater predictions. Previous work has often circumvented these issues by restricting the analysis to agricultural land where craters are easier to identify (Duncan et al., 2023; Lin et al., 2020). However, this is not an option for the

Vietnam War where much of the bombing happened in mountainous and densely vegetated regions.

Visual inspection and the pixel level accuracy assessment highlighted that the model detected *Group* craters in the correct areas (Fig. 5). However, the model did not accurately separate individual crater instances, a challenge that we also encountered during crater labeling. One way to address this could be to use an area-based approach that treats overlapping craters as one object and uses the total covered area instead of the crater count as a metric. *Crescent* craters were often located in areas with steep slopes and dense vegetation which meant that the appearance of these craters varied substantially, making reliable labeling difficult. *Bowl* craters were often old craters that had eroded and blended into the surroundings, which made labeling and detection challenging. These craters often occurred along rivers and canals where they were filled with water and only visible as dark circular blobs that

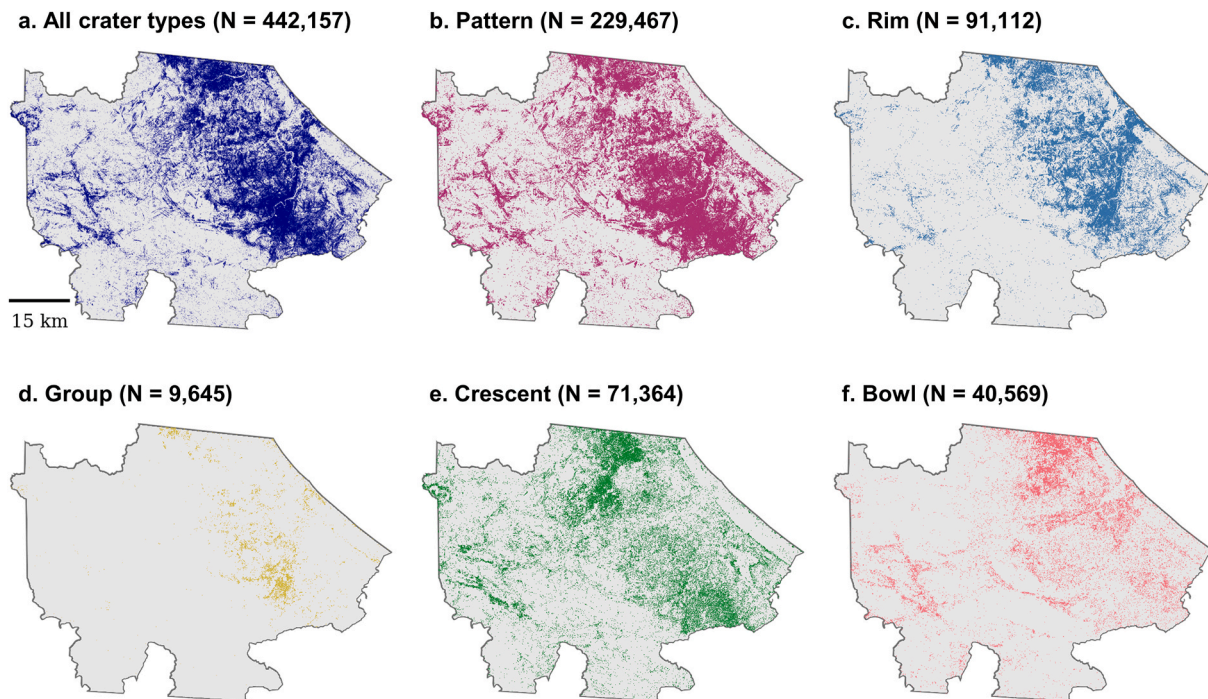


Fig. 5. Predicted craters by crater class in Quang Tri province. The total number of detected craters N is provided for each class. There is a clear difference in the distribution of the crater classes. Rim and Bowl craters were mostly located in the paddy fields closer to the coast, where the Rim craters seem to match better with previous year bombing. Group craters were rare and only predicted in very specific locations that have seen the heaviest bombing. Pattern and Crescent craters were spread across the whole study area.

could be confused with other ground features like trees. Therefore, *Crescent* or *Bowl* craters might be easier to detect in images taken closer to the date of the bombing.

3.2. Comparison of detected bomb craters with THOR bombing data

The THOR bombing records show that around 1 million bombs (QT: 654,730, TBA: 321,504) were dropped across the two study areas during the year preceding the KH-9 image acquisition (*previous year bombing*) and more than 3 million bombs (QT: 2.23 million, TBA: 1.13 million) during the entire conflict before the respective KH-9 images were taken (*total bombing*). Comparisons between detected craters and number of dropped bombs over grid cells of $2\text{ km} \times 2\text{ km}$ (Table 2) indicated that craters were positively correlated with *previous year bombing* (QT: $r = 0.76$, TBA: $r = 0.51$) and *total bombing* (QT: $r = 0.58$, TBA: $r = 0.51$) and correlation coefficients increased with grid cell size (Fig. 6).

A visual comparison showed that detected craters were located close to THOR target locations and were often organized in lines of craters characteristic for the B-52 bombing strikes (Fig. 4). The predicted crater locations are overlaid with aggregated bombing data for QT province (grid size: $2\text{ km} \times 2\text{ km}$) and the TBA (grid size: $4\text{ km} \times 4\text{ km}$) in Fig. 7. In grid cells in QT province for which more than 90% of *total bombing*

happened during the year before KH-9 image acquisition, the model detected a total of 157,846 craters, accounting for 46% of the 344,135 bombs dropped during *previous year bombing* (44% of *total bombing*).

3.2.1. Spatial precision

Craters identified in the KH-9 imagery can offer more precise information about impact locations of bombs compared to target locations from bomb strikes in the THOR data. Bomb strikes, which were less accurate at the time of the Vietnam War, did not always exactly hit their targets. Additionally, each bomb strike in the THOR data is confined to a single target location, but often encompasses tens or hundreds of dropped bombs. Fig. 8, depicting three target locations of B-52 bombing missions, shows resulting craters for the bombing of one target location spanning across several kilometers. Only few of these craters lie within a hundred-meter radius of the target location. This could explain the lower correlations between detected bomb craters and dropped bombs for smaller grid sizes of 100–500 m compared to grid sizes larger than 1000 m (Fig. 6).

While some uncertainty about the exact locations of bomb craters visible in the KH-9 imagery remains due to positional errors when georeferencing the KH-9 imagery, these are on a much smaller scale of on average less than 20 m. This is negligible for our application as

Table 2

Spearman correlation coefficients between detected craters and number of bombs dropped (THOR) aggregated across grid cells of $2\text{ km} \times 2\text{ km}$. For the Craters category detected craters of all crater classes were aggregated.

	Craters	Pattern	Rim	Group	Crescent	Bowl	Number of bombs dropped
Quang Tri							
Previous year bombing	0.76	0.74	0.78	0.68	0.52	0.62	654,730
Total bombing	0.58	0.55	0.46	0.33	0.61	0.41	2,232,280
Number of detected craters	442,157	229,467	91,112	9645	71,364	40,569	–
Tri-border area							
Previous year bombing	0.51	0.52	0.42	0.11	0.47	0.34	321,504
Total bombing	0.51	0.51	0.40	0.09	0.44	0.34	1,133,025
Number of detected craters	99,241	67,985	9995	46	16,836	4379	–

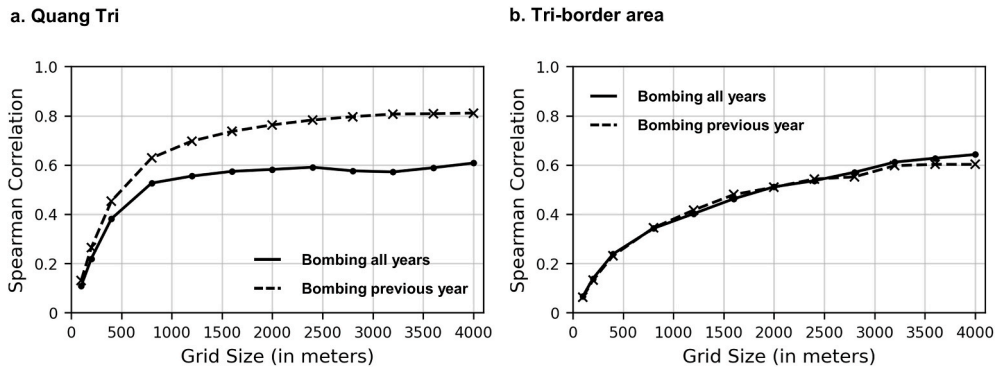


Fig. 6. Spearman correlations of the aggregated number of detected craters against bombs dropped (THOR) within grid cells of multiple sizes.

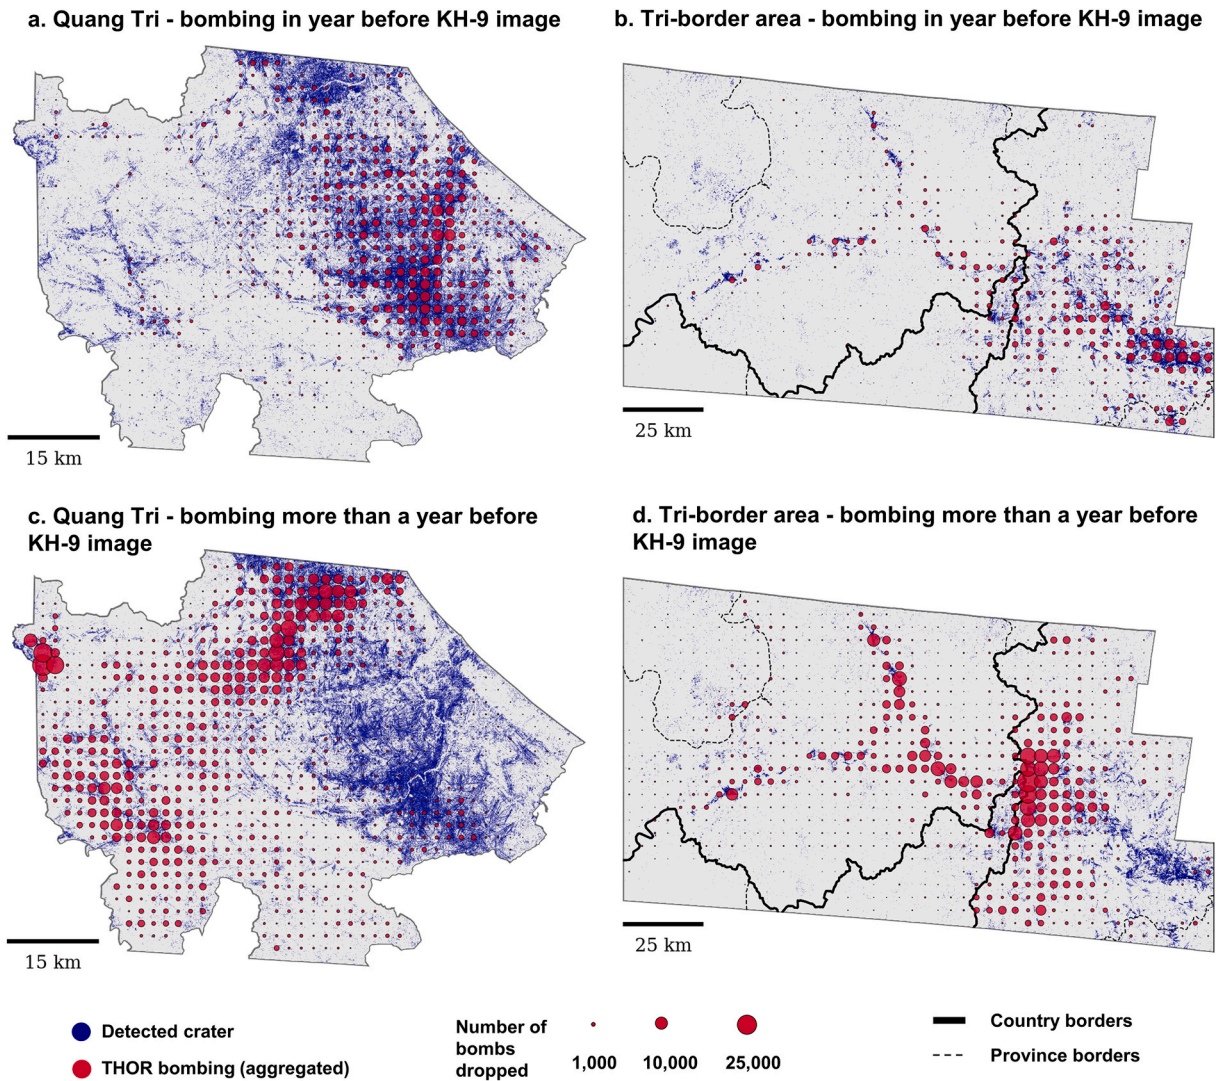


Fig. 7. Comparison of predicted bomb craters (blue) and THOR bombs dropped (red) aggregated across grid cells of $2\text{ km} \times 2\text{ km}$ for Quang Tri province and $4\text{ km} \times 4\text{ km}$ for the tri-border area. Bombs dropped during the year preceding the KH-9 image acquisition are shown in (a) and (b) whereas (c) and (d) show all bombing that happened more than a year before the image acquisition. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

detected bomb craters only serve as an indicator of risk for an area and not as exact locations of unexploded bombs, because bombs that created a crater already exploded. As bombs were dropped in groups it is still reasonable to assume that unexploded bombs would be located close to bomb craters from the same strike that often form distinct patterns

(Fig. 8). In the instance illustrated in Fig. 8, our estimation indicates that identifying the impact crater locations from the B-52 bomb strikes reduces the potential area for locating unexploded bombs from those strikes to about 9% of the area derived from the THOR target locations alone.




Fig. 8. KH-9 imagery for an area in Kon Tum province showing three target locations of B-52 bombing missions that occurred during the month preceding the image acquisition. Overlaid on the imagery are estimated risk areas, delineating areas where unexploded bombs resulting from the bombing strikes could be located. These risk zones were determined by a 2.5 km radius around the THOR target locations (red) and rectangles drawn around the visible impact craters (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

Moreover, the KH-9 imagery can be useful to identify and correct errors in the THOR data. The imagery in Fig. 8 revealed a discrepancy with the THOR data, where no nearby craters were visible for one target location. According to the THOR records, this mission had been diverted with some bombs supposedly dropped on the target visible in Fig. 8 and the remainder on a second target. However, the KH-9 imagery suggests it is more likely that all bombs were dropped at the second target and none at the first. This highlights the advantages of having multiple independent data sources that can be cross-referenced for a more thorough analysis.

3.2.2. Temporal analysis

Bomb craters can become increasingly difficult to detect from space over time. In Southeast Asia, with its dense rainforests and regular flooding, craters can quickly become covered up by vegetation, deformed by erosion or filled up by humans (Lin et al., 2020). Our analysis underscores that these effects impact detection results even after short periods of time, not only limiting the utility of current satellite imagery but emphasizing the need for additional imagery taken during the earlier stages of the war.

Our models detected a high concentration of craters near the cities of Quang Tri and Kon Tum, which were subject to heavy bombing during the year preceding the KH-9 image acquisition. Comparatively fewer craters were detected in areas targeted earlier during the war (Fig. 7). This is reflected by a higher correlation of detected craters with *previous year bombing* compared to *total bombing* in QT province, albeit not for the TBA (Table 2). Notably, we encountered challenges labeling and detecting partly eroded craters formed by bombs dropped earlier in the war, which we associated with the crater types *Bowl* and *Crescent*. To mitigate this bias towards areas bombed later in the war, we propose the use of the CORONA imagery captured during the earlier phases of the conflict (Munteanu et al., 2024).

Even in cases where bombings occurred near the time of image acquisition, not every dropped bomb recorded in THOR resulted in a crater detected by our model. In areas in QT province where bombing almost exclusively happened in the year before image acquisition, our model detected 150,895 craters, equivalent to 46% of bombs dropped that year. Several factors contribute to the lower number of detected craters, including: (1) bombs that left no craters, either because they

exploded on water or failed to explode altogether; (2) craters that initially formed but vanished within less than a year due to human activities, natural events like landslides or consecutive bombing of the same location; (3) craters that were obscured in the imagery by clouds, vegetation, or flooding; and (4) craters that were visible in the imagery but not detected by our models.

While some of these limitations can be addressed, many are inherent to the approach. However, their impacts can be mitigated if they are recognized and dealt with correctly. Typically, bombing strikes involved dropping numerous bombs on a single target, and identifying half of the resulting craters can provide a sufficiently accurate representation of the affected area. The main challenge lies in recognizing and compensating for factors that introduce bias, such as crater visibility and model performance variations across different soil and land cover types. Further research is needed to investigate these factors and should incorporate multiple data sources including the THOR bombing data, historical land cover maps and confirmed locations of UXO. Future studies should make use of these data sources in a stratified random sampling approach to identify biases based on a more representative subset of areas in Vietnam, Lao PDR and Cambodia.

3.3. Implications for mine action

The use of KH-9 imagery and derived crater locations could offer significant advantages to the mine action sector in Southeast Asia, extending beyond the capabilities of existing data sources used for non-technical surveys. Notably, our analysis revealed shortcomings of the THOR bombing data, emphasizing its lower precision compared to the detected crater locations. Additionally, the THOR data excludes weapons used by ground forces on both sides, such as artillery projectiles. Reports from local population carry a subjective element and are susceptible to recall bias, particularly when recounting events that happened 50 years ago. Additionally, their utility may be limited in previously unpopulated areas or where significant population shifts have happened since the war. Similarly, visual observations of UXO are invariably biased towards more populated areas. In contrast, the KH-9 imagery offers a more objective perspective, presenting an opportunity to address and overcome some of these challenges.

Despite the discussed benefits, the KH-9 imagery comes with its own

biases and limitations. Due to their danger, mine action in Southeast Asia focuses on contamination with cluster submunitions which are only about the size of a tennis ball (McCosker et al., 2020). While patterns of smaller craters, that might be linked to artillery fire or cluster bomb strikes, were visible in certain areas of the imagery, these craters would have been too small to be detected by our models. However, even where impact craters are not directly visible in the KH-9 imagery, the presence of other objects, such as larger craters or military infrastructure, could be indicators for the presence of cluster submunitions. More research is needed to explore this possibility and should make use of existing clearance data. Additionally, despite the current focus on cluster submunitions, there are increasing efforts to understand and manage the residual risk from other weapon types (Stauffer and Mestre, 2020). The number of craters, as detected by our models, could be a valuable indicator to help determine the residual risk level for an area at a more fine-grained level than would be possible using only the bombing records.

One of the key strengths of the KH-9 imagery lies in its cost-effectiveness and ease of integration into existing workflows. Each image, covering a large area, only costs \$30 on first request and previously requested images are freely available. The main limitation is the additional processing needed to orthorectify the images, including the time-consuming manual creation of ground control points. However, as demonstrated in our research, open-source tools can be used for this processing which reduces the cost. Products derived from our analysis can easily be integrated into existing mine action tools through imagery base layers for the KH-9 imagery and risk maps derived from detected bomb craters. The availability of the imagery for large parts of Southeast Asia makes it a useful tool for detailed analysis at both large (Fig. 7) and small (Fig. 8) scales.

3.4. Implications for sustainable development

Our work is directly aligned with Goal 16.1 of the Sustainable Development Goals (SDGs), which aims to “significantly reduce all forms of violence and related death rates everywhere”. Additionally, mine action has been shown to have a direct impact on 12 out of the 17 SDGs (Hofmann and Juergensen, 2017). Notably, Lao PDR and Cambodia went as far as introducing an 18th SDG that specifically addresses the legacy of unexploded ordnance. The craters detected in this study allow for a detailed analysis of the impact of bombing on post-conflict land-use changes, which have previously been linked to deforestation (SDG 13, SDG 15), reduced agricultural productivity (SDG 2) and hindered infrastructure development (SDG 1, SDG 9, SDG 11) (Clerici et al., 2020; Lin, 2022; Martin et al., 2019; Munteanu et al., 2024; Ounmany and Andriess, 2018).

In addition to supporting mine action, our work extends to other domains. While bomb craters have been identified as biodiversity hotspots (SDG 15) (Vad et al., 2017), they could also present potential public health hazards (SDG 3), as the stagnant water they collect can become breeding sites for mosquito larvae (Wimberly et al., 2021). Moreover, sediment buildup within these craters may contain concentrated levels of dioxins from herbicide spraying during the Vietnam War, posing a risk to individuals (SDG 3), particularly when the craters are repurposed as fish ponds (Olson and Morton, 2019). Bomb craters have been shown to alter hydrology and soil development in affected areas (Certini et al., 2013; Hupy and Koehler, 2012; Kiernan, 2015), but it remains unclear whether this could relate to the prevalence of landslides and flooding (SDG 13, SDG 15). More research is needed to understand these effects, and the bomb crater locations identified in this study could serve as a valuable resource for such investigations.

4. Conclusions

The presence of UXO in Vietnam, Lao PDR, and Cambodia continues to pose a significant threat to both public health and economic

development. However, due to the expense and time required for detailed surveys, the exact locations of UXO often remain unknown. This study developed a workflow to orthorectify and automatically detect bomb craters in the declassified KH-9 imagery. The models achieved an overall F1-Score of 0.61 and predicted more than 500,000 bomb craters across the two study areas. The results demonstrate how the identified bomb craters can complement existing data sources such as the THOR bombing records. We estimate this could allow for more precise localization of suspected hazardous areas during non-technical surveys as well as a more fine-grained determination of residual risk of UXO in areas where extensive clearance operations are deemed too expensive. The developed methods are scalable to large regions at low cost and directly transferable to other affected areas in Southeast Asia. The instance segmentation workflow for the crater detection is also applicable to more recent conflicts including the ongoing war in Ukraine.

Funding

This work was supported by NERC through a SENSE CDT studentship (NE/T00939X/1) and by the Conflict and Environment Observatory. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising from this submission.

CRedit authorship contribution statement

Philipp Barthelme: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation. **Eoghan Darbyshire:** Writing – review & editing, Supervision, Conceptualization. **Dominick V. Spracklen:** Writing – review & editing, Supervision, Conceptualization. **Gary R. Watmough:** Writing – review & editing, Supervision, Conceptualization.

Declaration of AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT (GPT-3.5) in order to improve the readability of the paper and to support the writing and documentation of code. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code used for the analysis in this study is available via <https://github.com/pbarthelme/detecting-vietnam-war-bomb-craters> and archived at <https://doi.org/10.5281/zenodo.10709375>. The predicted bomb craters and trained models are available at <https://doi.org/10.5281/zenodo.10629987>. The georeferenced KH-9 images are available via <https://doi.org/10.7488/ds/7682> for Quang Tri province and via <https://doi.org/10.7488/ds/7683> for the tri-border area. The scanned KH-9 images (not georeferenced) are available as part of the Declassified Satellite Imagery - 3 collection courtesy of the U.S. Geological Survey (USGS EROS Center, 2018). They can be accessed via the EarthExplorer website <https://earthexplorer.usgs.gov/> as part of the Data Set: Declassified Data > Declass 3 collection (Entity IDs for filtering the specific images used in the study are provided in the Supplementary Materials), download requires setting up a free account with USGS EROS web services. The THOR bombing records are available at <https://data.world/datamil/vietnam-war-thor-data> (File name: thor_data_vietnam.

csv, last accessed 5. September 2023), download requires setting up a free account with data.world. Version 4.1 of the GADM administrative units used for creating some of the figures in this study are freely available for academic and other non-commercial use at www.gadm.org (last accessed: 2. February 2024). The SRTM GL1 dataset used for the orthorectification of the KH-9 imagery is available at OpenTopography via <https://doi.org/10.5069/G9445JDF> (NASA Shuttle Radar Topography Mission (SRTM), 2013). Version 3.0.0 of the Nasa Ames Stereo Pipeline used for orthorectification of the KH-9 imagery is preserved at <https://doi.org/10.5281/ZENODO.5140581>, available via Apache License 2.0 and developed openly at <https://github.com/NeoGeographyToolkit/StereoPipeline> (Beyer et al., 2021). Version 3.16.9 of QGIS used for the creation of ground control points is preserved at <https://download.qgis.org/downloads/qgis-3.16.9.tar.bz2> (last accessed 2. February 2024), available via GNU-General-Public-License and developed openly at <https://github.com/qgis/QGIS> (QGIS Association, 2023). The U-Net was implemented using the Python packages pytorch-cuda v11.7 (Paszke et al., 2019) and segmentation-models-pytorch v0.3.3 (Iakubovskii, 2019).

Acknowledgments

We thank the United States Geological Survey for providing the KH-9 imagery and the United States Department of Defense for providing the THOR bombing data. We thank Tim Horner for his valuable insights into mine action in Vietnam and Kinh Bac Dang for his support in identifying ground features in the KH-9 imagery. This work has made use of the resources provided by the Edinburgh Compute and Data Facility (<http://www.ecdf.ed.ac.uk/>).

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.srs.2024.100143>.

References

- Anderson, D.L., 2002. The Vietnam War from A to Z. In: *The Columbia Guide to the Vietnam War*. Columbia University Press, p. 92–92.
- Bennett, M.M., Van Den Hoek, J., Zhao, B., Prishchepov, A.V., 2022. Improving satellite monitoring of armed conflicts. *Earth's Future* 10, e2022EF002904. <https://doi.org/10.1029/2022EF002904>.
- Bold, M., Avenell, D., 2021. Data-driven decision-making in Southeast Asia. *The Journal of Conventional Weapons Destruction* 25.
- Caicedo, J.C., Roth, J., Goodman, A., Becker, T., Karhohs, K.W., Broisin, M., Molnar, C., McQuin, C., Singh, S., Theis, F.J., Carpenter, A.E., 2019. Evaluation of deep learning strategies for nucleus segmentation in fluorescence images. *Cytometry* 95, 952–965. <https://doi.org/10.1002/cyto.a.23863>.
- Certini, G., Scalenghe, R., Woods, W.I., 2013. The impact of warfare on the soil environment. *Earth Sci. Rev.* 127, 1–15. <https://doi.org/10.1016/j.earscirev.2013.08.009>.
- Chen, D., Hu, F., Mathiopoulos, P.T., Zhang, Z., Peethambaran, J., 2023. MC-UNet: martian crater segmentation at semantic and instance levels using U-Net-Based convolutional neural network. *Rem. Sens.* 15 <https://doi.org/10.3390/rs15010266>.
- Clerici, N., Armenteras, D., Kareiva, P., Botero, R., Ramírez-Delgado, J.P., Forero-Medina, G., Ochoa, J., Pedraza, C., Schneider, L., Lora, C., Gómez, C., Linares, M., Hirashiki, C., Biggs, D., 2020. Deforestation in Colombian protected areas increased during post-conflict periods. *Sci. Rep.* 10, 4971. <https://doi.org/10.1038/s41598-020-61861-y>.
- Clermont, D., Kruse, C., Rottensteiner, F., Heipke, C., 2019. Supervised detection of bomb craters in historical aerial images using convolutional neural networks. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences - ISPRS Archives* 42, 67–74. <https://doi.org/10.5194/isprs-archives-XLII-2-W16-67-2019>.
- Cluster Munition Coalition, 2023. Cluster Munition Monitor 2023. (ICBL-CMC: August 2023). <https://www.the-monitor.org/en-gb/reports/2023/cluster-munition-monitor-2023.aspx>. (Accessed 5 October 2023).
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. ImageNet: a large-scale hierarchical image database. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>.
- Deshpande, P., Belwalkar, A., Dikshit, O., Tripathi, S., 2021. Historical land cover classification from CORONA imagery using convolutional neural networks and geometric moments. *Int. J. Rem. Sens.* 42, 5144–5171. <https://doi.org/10.1080/01431161.2021.1910365>.
- Duncan, E.C., Skakun, S., Kariryaa, A., Prishchepov, A.V., 2023. Detection and mapping of artillery craters with very high spatial resolution satellite imagery and deep learning. *Science of Remote Sensing* 7, 100092. <https://doi.org/10.1016/j.srs.2023.100092>. –100092.
- Frost, A., Boyle, P., Autier, P., King, C., Zwijnenburg, W., Hewitson, D., Sullivan, R., 2017. The effect of explosive remnants of war on global public health: a systematic mixed-studies review using narrative synthesis. *Lancet Public Health* 2, e286–e296. [https://doi.org/10.1016/S2468-2667\(17\)30099-3](https://doi.org/10.1016/S2468-2667(17)30099-3).
- Ghuffar, S., King, O., Guillet, G., Rupnik, E., Bolch, T., 2023. Brief communication: glacier mapping and change estimation using very high-resolution declassified Hexagon KH-9 panoramic stereo imagery (1971–1984). *Cryosphere* 17, 1299–1306. <https://doi.org/10.5194/tc-17-1299-2023>.
- Gut, D., Tabor, Z., Szymkowski, M., Rozynek, M., Kucybała, I., Wojciechowski, W., 2022. Benchmarking of deep architectures for segmentation of medical images. *IEEE Trans. Med. Imag.* 41, 3231–3241. <https://doi.org/10.1109/TMI.2022.3180435>.
- Hammer, E., FitzPatrick, M., Ur, J., 2022. Succeeding CORONA: declassified HEXAGON intelligence imagery for archaeological and historical research. *Antiquity* 96, 679–695. <https://doi.org/10.15184/aqy.2022.22>.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*. <https://doi.org/10.48550/arXiv.1512.03385>.
- High, H., Curran, J.R., Robinson, G., 2013. Electronic records of the air war over Southeast Asia: a database analysis. *J. Vietnamese Stud.* 8, 86–124. <https://doi.org/10.1525/vs.2014.8.4.86>.
- Hofmann, U., Juehgens, O., 2017. Mine Action and the Sustainable Development Goals. *GICHD & UNDP*. <https://www.undp.org/publications/mine-action-and-sustainable-development-goals>. (Accessed 2 May 2024).
- Hupy, J.P., Koehler, T., 2012. Modern warfare as a significant form of zoogeomorphic disturbance upon the landscape. *Geomorphology, Special Issue Zoogeomorphology and Ecosystem Engineering Proceedings of the 42nd Binghamton Symposium in Geomorphology* 169–182. <https://doi.org/10.1016/j.geomorph.2011.05.024> held 21–23 October 2011 157–158.
- Iakubovskii, P., 2019. Segmentation models pytorch (Version 0.3.3). https://github.com/qubvel/segmentation_models.pytorch. (Accessed 5 July 2023).
- Kiernan, K., 2015. Nature, severity and persistence of geomorphological damage caused by armed conflict. *Land Degrad. Dev.* 26, 380–396. <https://doi.org/10.1002/ldr.2216>.
- Kingma, D.P., Ba, J., 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*. <https://doi.org/10.48550/arXiv.1412.6980>.
- Kruse, C., Rottensteiner, F., Heipke, C., 2019. Marked point processes for the automatic detection of bomb craters in aerial wartime images. *Int. Arch. Photogram. Rem. Sens. Spatial Inf. Sci. XLII-2-W13*, 51–60. <https://doi.org/10.5194/isprs-archives-XLII-2-W13-51-2019>.
- Kugelman, J., Allman, J., Read, S.A., Vincent, S.J., Tong, J., Kalloniatis, M., Chen, F.K., Collins, M.J., Alonso-Caneiro, D., 2022. A comparison of deep learning U-Net architectures for posterior segment OCT retinal layer segmentation. *Sci. Rep.* 12, 14888. <https://doi.org/10.1038/s41598-022-18646-2>.
- Lasaponara, R., Yang, R., Chen, F., Li, X., Masini, N., 2018. Corona satellite pictures for archaeological studies: a review and application to the lost forbidden city of the han-wei dynasties. *Surv. Geophys.* 6 (39), 1303–1322. <https://doi.org/10.1007/S10712-018-9490-2>, 2018 39.
- Le, D.T., Pham, T.M., Polachek, S., 2022. The long-term health impact of agent orange: evidence from the Vietnam War. *World Dev.* 155, 105813. <https://doi.org/10.1016/J.WORLDDEV.2022.105813>. –105813.
- Le, K., Nguyen, M., 2020. Aerial bombardment and educational attainment. *Int. Rev. Appl. Econ.* 34, 361–383. <https://doi.org/10.1080/02692171.2020.1736012>.
- Lin, E., 2022. How war changes land: soil fertility, unexploded bombs, and the underdevelopment of Cambodia. *Am. J. Polit. Sci.* 66, 222–237. <https://doi.org/10.1111/ajps.12577>.
- Lin, E., Qin, R., Edgerton, J., Kong, D., 2020. Crater detection from commercial satellite imagery to estimate unexploded ordnance in Cambodian agricultural land. *PLoS One* 15, e0229826. <https://doi.org/10.1371/journal.pone.0229826>.
- Lin, T.-Y., Goyal, P., Girshick, R.B., He, K., Dollár, P., 2017. Focal loss for dense object detection. *arXiv preprint arXiv:1708.02002*. <https://doi.org/10.48550/arXiv.1708.02002>.
- Martin, M.F., Dolven, B., Feickert, A., Lum, T., 2019. War legacy issues in Southeast Asia: unexploded ordnance (UXO). *Current politics and economics of South, Southeastern, and Central Asia* 28, 199–230.
- McCosker, K., Stoa, J., Harrison, K., 2020. Cluster munition remnant survey in Southeast Asia. *The Journal of Conventional Weapons Destruction* 23.
- Miguel, E., Roland, G., 2011. The long-run impact of bombing Vietnam. *J. Dev. Econ.* 96, 1–15. <https://doi.org/10.1016/j.jdevco.2010.07.004>.
- Mulyanto, M., Faisal, M., Prakosa, S.W., Leu, J.-S., 2021. Effectiveness of focal loss for minority classification in network intrusion detection systems. *Symmetry* 13. <https://doi.org/10.3390/sym13010004>.
- Munteanu, C., Kraemer, B.M., Hansen, H.H., Miguel, S., Milner-Gulland, E.J., Nita, M., Ogashawara, I., Radeloff, V.C., Roverelli, S., Shumilova, O.O., Storch, I., Kuemmerle, T., 2024. The potential of historical spy-satellite imagery to support research in ecology and conservation. *BioScience* biae002. <https://doi.org/10.1093/biosci/biae002>.
- NASA Shuttle Radar Topography Mission (SRTM), 2013. Shuttle Radar Topography Mission (SRTM) global. <https://doi.org/10.5069/G9445JDF>.
- Nguyen, L.B., 2020. Land cover change detection in northwestern Vietnam using Landsat images and Google Earth Engine. *J. Water Land Dev.* 46, 162–169. <https://doi.org/10.24425/JWLD.2020.134209>.
- Nita, M.D., Munteanu, C., Gutman, G., Abrudan, I.V., Radeloff, V.C., 2018. Widespread forest cutting in the aftermath of World War II captured by broad-scale historical

- Corona spy satellite photography. *Rem. Sens. Environ.* 204, 322–332. <https://doi.org/10.1016/J.RSE.2017.10.021>.
- Olson, K.R., Morton, L.W., 2019. Long-term fate of agent orange and dioxin TCDD contaminated soils and sediments in Vietnam hotspots. *Open J. Soil Sci.* 9, 1. <https://doi.org/10.4236/ojss.2019.91001>.
- Ounmany, K., Andriess, E., 2018. The legacy of the Vietnam War: making a living amid unexploded ordnance in Xieng Khouang province, northern Laos. *Asian Stud. Rev.* 42, 439–458. <https://doi.org/10.1080/10357823.2018.1481365>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gímelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E.Z., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: an imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*. <https://doi.org/10.48550/arXiv.1912.01703>.
- QGIS Association, 2023. QGIS Geographic Information System.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241. https://doi.org/10.1007/978-3-319-24574-4_28.
- Schober, P., Boer, C., Schwarte, L.A., 2018. Correlation coefficients: appropriate use and interpretation. *Anesth. Analg.* 126, 1763–1768. <https://doi.org/10.1213/ANE.0000000000002864>.
- Silburt, A., Ali-Dib, M., Zhu, C., Jackson, A., Valencia, D., Kissin, Y., Tamayo, D., Menou, K., 2019. Lunar crater identification via deep learning. *Icarus* 317, 27–38. <https://doi.org/10.1016/j.icarus.2018.06.022>.
- Sohn, H.-G., Kim, G.-H., Yom, J.-H., 2004. Mathematical modelling of historical reconnaissance CORONA KH-4B Imagery. *Photogramm. Rec.* 19, 51–66. <https://doi.org/10.1046/j.0031-868X.2003.00257.x>.
- Stauffer, K., Mestre, C., 2020. Long-term risk management tools for protocols for residual explosive ordnance mitigation: a pretest in Vietnam. *The Journal of Conventional Weapons Destruction* 23.
- UNMAS, 2019. IMAS 08.10 Non-technical survey. <https://www.mineactionstandards.org/standards/08-10/>. (Accessed 2 May 2024).
- USGS EROS Center, 2018. Declassified satellite imagery - 3. <https://doi.org/10.5066/F7WD3Z10>.
- Vad, C.F., Péntek, A.L., Cozma, N.J., Földi, A., Tóth, A., Tóth, B., Böde, N.A., Móra, A., Ptacnik, R., Ács, É., Zsuga, K., Horváth, Z., 2017. Wartime scars or reservoirs of biodiversity? The value of bomb crater ponds in aquatic conservation. *Biol. Conserv.* 209, 253–262. <https://doi.org/10.1016/j.biocon.2017.02.025>.
- Waga, J.M., Szypuła, B., Fajer, M., 2022. The archaeology of unexploded world war II bomb sites in the Kozle Basin, southern Poland. *Int. J. Hist. Archaeol.* <https://doi.org/10.1007/s10761-022-00672-5>.
- Wang, H., Miao, F., 2022. Building extraction from remote sensing images using deep residual U-Net. *European Journal of Remote Sensing* 55, 71–85. <https://doi.org/10.1080/22797254.2021.2018944>.
- Wimberly, M.C., Beurs, K.M. de, Loboda, T.V., Pan, W.K., 2021. Satellite observations and malaria: new opportunities for research and applications. *Trends Parasitol.* 37, 525–537. <https://doi.org/10.1016/j.pt.2021.03.003>.
- Yamada, T., Yamada, H., 2021. The long-term causal effect of U.S. bombing missions on economic development: evidence from the Ho Chi Minh Trail and Xieng Khouang Province in Lao P.D.R. *J. Dev. Econ.* 150, 102611. <https://doi.org/10.1016/J.JDEVECO.2020.102611>. –102611.
- Zhou, Y., Chen, G., Qiao, X., Lu, L., 2021. Mining high-resolution KH-9 panoramic imagery to determine earthquake deformation: methods and applications. *IEEE Trans. Geosci. Rem. Sens.* <https://doi.org/10.1109/TGRS.2021.3116441>.
- Beyer, R., Alexandrov, O., McMichael, S., 2021. NeoGeographyToolkit/StereoPipeline 3.0.0. Zenodo. <https://doi.org/10.5281/zenodo.5140581>.