communications earth & environment

Article

A Nature Portfolio journal

https://doi.org/10.1038/s43247-025-02183-7

An open-source tool for mapping war destruction at scale in Ukraine using Sentinel-1 time series

Check for updates

Olivier Dietrich[1] ✉, Torben Peters[1], Vivien Sainte Fare Garnot[2], Valerie Sticher[3], Thao Ton-That Whelan[4], Konrad Schindler[1] & Jan Dirk Wegner[2]

Access to detailed war impact assessments is crucial for humanitarian organizations to assist affected populations effectively. However, maintaining a comprehensive understanding of the situation on the ground is challenging, especially in widespread and prolonged conflicts. Here, we present a scalable method for estimating building damage resulting from armed conflicts. By training a machine learning model on Synthetic Aperture Radar image time series, we generate probabilistic damage estimates at the building level, leveraging existing damage assessments and open building footprints. To allow large-scale inference and ensure accessibility, we tie our method to run on Google Earth Engine. Users can adjust confidence intervals to suit their needs, enabling rapid and flexible assessments of war-related damage across large areas. We provide two publicly accessible dashboards: a Ukraine Damage Explorer to dynamically view our precomputed estimates and a Rapid Damage Mapping Tool to run our method and generate custom maps.

The Russian invasion of Ukraine in February 2022 escalated a simmering conflict, which had previously seen unofficial Russian involvement limited to the country's eastern regions, into a full-scale war between the two countries. Russian troops crossed into Ukraine from multiple fronts and advanced through the country, reorienting their geographical focus as Ukrainian forces mounted an unexpectedly fierce resistance. As the Russian troops advanced, they subjected numerous cities to heavy shelling and destroyed critical infrastructure. Millions of Ukrainians fled their homes as damage and destruction spread throughout the country. Two years into the conflict, the war has caused hundreds of thousands of casualties and inflicted billions of dollars in damages to Ukraine's infrastructure[1,2].

The fast-paced early stages of the Russia-Ukraine war serve as a stark reminder of the challenges humanitarian organizations encounter when monitoring ongoing armed conflicts in real-time. Even in less intense conflicts, like the war in Eastern Ukraine before the full-scale Russian invasion, violence can flare up unexpectedly, and minor incidents may trigger escalatory dynamics. Thus, regardless of a conflict's intensity, it is essential for humanitarian organizations to have a complete and up-to-date understanding of the situation on the ground to effectively support those most affected by the conflict. However, the task of maintaining such an overview is challenging, especially with limited resources. The difficulty is compounded when conflicts affect large geographic areas, persist over long periods of time, or occur in regions that are inaccessible due to security concerns.

To overcome these challenges, organizations have increasingly relied on satellite imagery to supplement on-the-ground monitoring. Although satellite images lack contextual details and a narrative dimension, they provide snapshots of the conditions on the ground, making it possible to observe how the situation evolves and how destruction spreads. Furthermore, satellites offer global coverage and automatic data acquisition, which is essential for monitoring large, remote, or inaccessible areas. The primary technique to perform damage assessments from satellite imagery is a manual analysis of very high-resolution (VHR) optical images. As it involves the visual comparison of expensive commercial imagery, that method is labor-intensive and costly. Despite its accuracy and reliability, it depends on the availability of VHR data and does not scale to conflict zones that are large and that require long-term monitoring.

Nowadays, the analysis of satellite imagery with machine learning offers opportunities to automate parts of the labor-intensive remote monitoring process. Multiple studies have demonstrated the efficiency and effectiveness of this approach for assessing building damage. The most accurate methods currently involve applying deep neural networks to VHR optical data[3–8], often leveraging large dedicated datasets[9]. These approaches primarily focus on natural disasters, which typically require a one-time

[1]Photogrammetry and Remote Sensing, ETH Zurich, Zurich, Switzerland. [2]EcoVision Lab, Department of Mathematical Modeling and Machine Learning, University of Zurich, Zurich, Switzerland. [3]Department of Political Science, University of Zurich, Zurich, Switzerland. [4]International Committee of the Red Cross, Geneva, Switzerland. ✉e-mail: odietrich@ethz.ch

analysis of relatively localized, clustered damages. In contrast, armed conflicts can last for months, years, or even decades, resulting in spatially dispersed and gradually increasing destruction. Consequently, they require continuous monitoring over extended periods so that automated change detection can on the one hand, bring larger benefits, but on the other hand, is more challenging[10].

Recent work has begun to apply research specifically to the task of assessing building destruction induced by armed conflicts. In a seminal study, Mueller et al. trained a CNN with VHR optical imagery and labels from UNOSAT[11] to assess war-induced building destruction in Syria, leveraging the persistent nature of war damage over time due to the absence of reconstruction efforts amidst ongoing fighting[12]. However, accessing VHR optical imagery for conflict zones is costly, as commercial sources typically do not provide free access to their database in these situations[13], unlike after natural disasters[14,15]. Despite being an active field of research[16,17], relying on repeated, large-area coverage with VHR images to monitor armed conflicts at scale is not realistic for most actors. Very recently, Hou et al. proposed a deep learning-based approach that can handle both VHR and moderate-resolution optical imagery[18].

Beyond optical images, Synthetic Aperture Radar (SAR) imagery offers a promising alternative. SAR is an active sensor system that illuminates the Earth's surface with microwave pulses and captures the backscattered signals. Unlike its optical counterpart, SAR can operate at any time of day (night, respectively) and is largely unaffected by clouds. SAR has a well-established history in damage assessment for both natural and anthropogenic events[19,20]. Broadly, SAR-based change detection can be divided into incoherent and coherent methods. Incoherent approaches rely only on backscatter amplitude, typically comparing intensity or correlation differences between pre- and post-event images[21–25]. Coherent methods, on the other hand, use both amplitude and phase information through interferometry (InSAR), generally resulting in more accurate results[26–28]. Additionally, SAR's sub-meter wavelength allows coherent techniques to detect changes far smaller than the pixel resolution, particularly when using permanent scatterers, which maintain coherence over long time series[29]. Coherent change detection is a standard approach in the InSAR community to assess building damage after a disaster rapidly[30–33], c.f. the popular (and patented) *Damage Proxy Maps*. Recent works have combined machine learning and coherence data to produce more accurate maps, either with handcrafted features[34,35] or by using recurrent neural networks trained on coherence time series[36,37]. On the downside, generating high-quality coherence maps is complex and involves a loss of resolution. In this work, we focus on large-scale deployment and accessibility and opted for developing our approach on the cloud-based platform Google Earth Engine (GEE)[38]. Since GEE does not support the Sentinel-1 SLC product[39] necessary for phase-based methods, our approach is purely amplitude-based.

A few studies have already explored satellite-based damage assessment for the Ukraine War, but most focused on specific areas or events. Aimaiti et al. achieved 58% accuracy on UNOSAT labels in the Kyiv region based on handcrafted features from Sentinel-1 and Sentinel-2[40]. Huang et al. used optical and coherence SAR time series to identify burnt areas, respectively urban destruction in Mariupol, reaching an overall accuracy of ≈ 60% across four classes relative to labels retrieved from social media images[41]. Tavakkoliestahbanati et al. leverage multi-temporal InSAR to monitor deformation from a single event, the Kakhovka dam collapse[42]. Ballinger's method is arguably the one most directly comparable to ours. It relies on a pixel-wise statistical *t*-test on Sentinel-1 time series, implemented directly in GEE. UNOSAT labels from various locations were then used to calibrate the decision threshold[43]. Lastly, Scher and Van Den Hoek have generated nationwide coherence maps and performed coherent change detection against a reference pre-invasion image. At the time of writing their maps have not yet been made public[44].

In this work, we introduce an easy-to-use, open-source war impact mapping tool based purely on SAR amplitude data and demonstrate its effectiveness through a comprehensive building damage assessment across Ukraine. Our solution leverages existing, point-wise damage maps from UNOSAT[11] and public, open SAR imagery from the Sentinel-1 mission[39]. Harnessing the parallel computing capabilities of GEE[38], we train a model to estimate the likelihood of war-related changes from paired time series of Sentinel-1 backscatter. We choose a fixed one-year interval in 2020 as our reference period and iteratively generate damage likelihood maps for 3-month periods ranging from February 2021 up to February 2024. Our model, based on the Random Forest algorithm, allows for straightforward deployment on GEE and outputs damage maps with 10×10 m grid spacing, the original ground sampling distance (GSD) of the Sentinel-1 GRD product. In a second processing step, we intersect the per-pixel maps with publicly available building footprints from Overture Maps[45] to generate a damage estimate per building. Our tool produces maps of war damage at a very low cost and enables continuous country-scale monitoring. These maps can serve both as a basis for large-scale, aggregated assessments and as guidance to focus manual verification efforts for individual settlements or buildings. We offer two public online dashboards: In the first one, users can dynamically view and explore the pre-computed damage maps described in the following. The second one allows users to run our method themselves for their desired locations and time periods and create custom maps.

The key advantages of our approach are its scalability, accessibility, and ability to transfer to different geographic contexts. Using commercial VHR imagery for regular screening across entire conflict zones is not viable for humanitarian organizations[46], because it would be too resource-intensive, in terms of both direct data cost as well as infrastructure for data download, storage, and compute. By utilizing moderate-resolution satellite images that are free of charge and acquired with short, regular revisits, we enable a more sustainable monitoring approach. The use of SAR imagery has the additional advantage that the same method can be expected to work well in other conflict contexts, including those with different architectural patterns and/or with even more frequent cloud cover, as it is purely geometrical and does not rely on any location-specific a-priori assumptions. The proposed method is also fast, reproducible, and easy to adapt, even for users who are not experts in remote sensing or image analysis. We believe that our work could be a valuable and accessible resource and could be used to screen large conflict areas beyond our specific application case.

Results

Evaluation

We first evaluate our model's predictions against additional damage assessments excluded from the training set (see Fig. 1 for the geographic distribution of our test set). The UNOSAT-labeled locations are compared to the corresponding raw pixel predictions of the model. True positive and false negative rates are counted in the 3-month periods from 2022 (after the invasion), and true negative and false positive rates are counted in the periods from 2021 (before the invasion). To guard against potential shifts between the point labels and the satellite tiles, e.g., due to high incidence angles or inaccurate annotation, we compare to the maximum predicted damage probability in a 3×3 pixel window centered at the reference label. We report our results in Table 1. We obtain an F1-score of 74.9% for the damaged class and an AUC of 81.3%. Using the same evaluation protocol, we compare our result to the pixel-wise *t*-test (PWTT) method[43]. To our knowledge, this is the only method directly comparable to ours, since others either require coherence maps or assume that the precise date of the destruction event is known, which is not the case for our UNOSAT labels. Our method consistently outperforms that baseline, see Table 2. For details about the comparison see Supplementary Note 1.

The evaluation of held-out validation data also helps establish an appropriate threshold for distinguishing between damaged and undamaged pixels. Indeed, while the model has been trained on a balanced dataset, the actual distribution of damaged versus undamaged pixels is highly skewed, and the prevalence of undamaged pixels in the test set means that thresholding at damage probability 0.5 will cause a considerable number of false alarms. To address this, we set a target precision of 90%, aligning with the reality of operational deployment, where it is crucial to keep the false positive rate at a manageable level. With this choice, we find an optimal threshold of

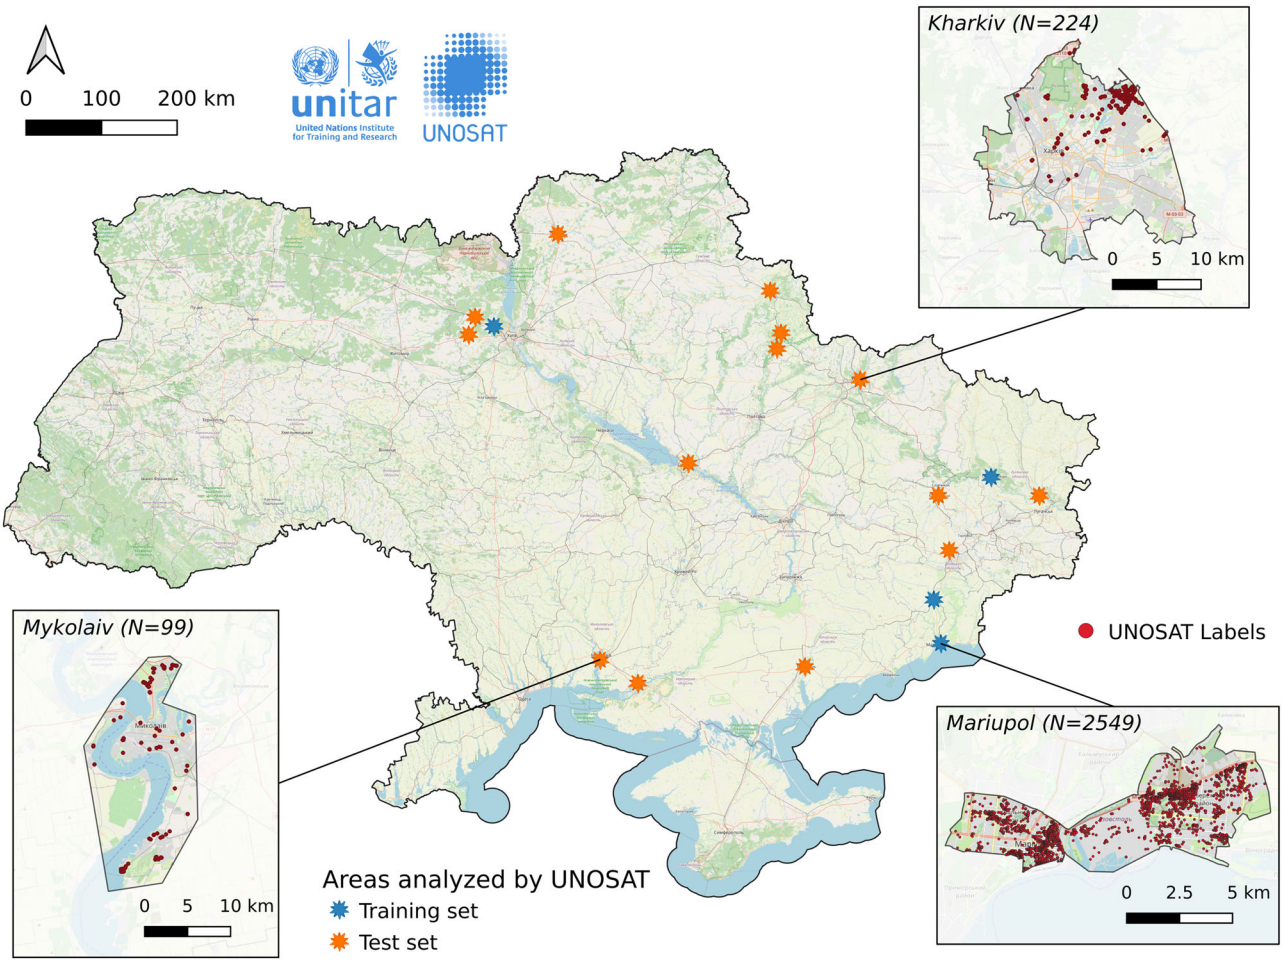


Fig. 1 | Regions analyzed by UNOSAT. Overview of the 18 regions assessed by UNOSAT in Ukraine. Blue and orange dots indicate cities used for training and testing, respectively. The three insets highlight the variation in damage density across different regions. Background map: OpenStreetMap.

Table 1 | Evaluation metrics

Label	Precision	Recall	F_1	Accuracy	AUC
Damaged	67.1 (90.3)	84.6 (46.5)	74.9 (61.4)	80.3 (79.7)	81.3 (71.9)
Undamaged	90.5 (77.4)	78.0 (97.3)	83.8 (86.2)		

Quantitative detection performance relative to UNOSAT labels (at previously unseen locations), using a confidence threshold of 0.5. Values in parentheses indicate performance at the threshold of 0.655, reaching a target precision of 90%. The *undamaged* class is evaluated using the same locations, but acquisition windows before the full-scale invasion. All metrics are given as percentages.

Table 2 | Comparison with PWTT

Method	Recall	F_1	AUC
PWTT[43]	76.3	68.4	75.7
Ours	84.6	74.9	81.3

Comparison of our method with the PWTT[43], with the latter calibrated for Ukraine. Metrics were computed on our test set, using the same settings for both approaches. All metrics are given in percentages.

0.655. For details about the threshold optimization see Supplementary Note 2. With that value, the recall reaches 46.5%. We point out that there is an inherent trade-off between missed detections and false alarms, and the right threshold depends on the application task. Our dashboard allows the user to adjust it with a slider and explore how different thresholds affect the estimates. Results with the confidence threshold of 0.655 are also reported in Table 1.

Close-ups

We use Chernihiv, one of the cities in our test set, to qualitatively illustrate the performance and limitations of our approach. Figure 2 shows both heatmaps of damage probability and per-building maps after post-processing with building footprints. Inset A highlights one limitation of using building footprints, as the leftmost damage label has been correctly predicted in the heatmap but does not have a corresponding footprint. Inset B illustrates how the model can detect correct patterns of destruction, even for small structures. However, some labels (e.g., center-top) are completely missed, likely because the affected building volume was too small to be visible at Sentinel-1 resolution. Inset C shows both how multiple UNOSAT labels can fall within a single footprint, and how taking the mean prediction over a partially affected building footprint can discard some correct predictions (third point from the bottom). In the main figure, outside the boundaries of UNOSAT analysis, we observe agricultural fields wrongly identified as damaged in the heatmap. This is expected, since our model has not been trained for patterns that occur outside of settlement areas, and illustrates the need to post-process the raw heatmaps with settlement masks or building footprints.

Web-based tool

We have developed a web application based on Google Earth Engine (GEE) to facilitate rapid reproduction of our results and enable the generation of new maps without any local software installation. The application features a user-friendly dashboard (Fig. 3) that provides the functionality to perform damage assessment for a user-specified region of interest and using user-defined pre- and post-event observation periods. The predicted maps can also be exported as raster files directly from the interface (subject to GEE

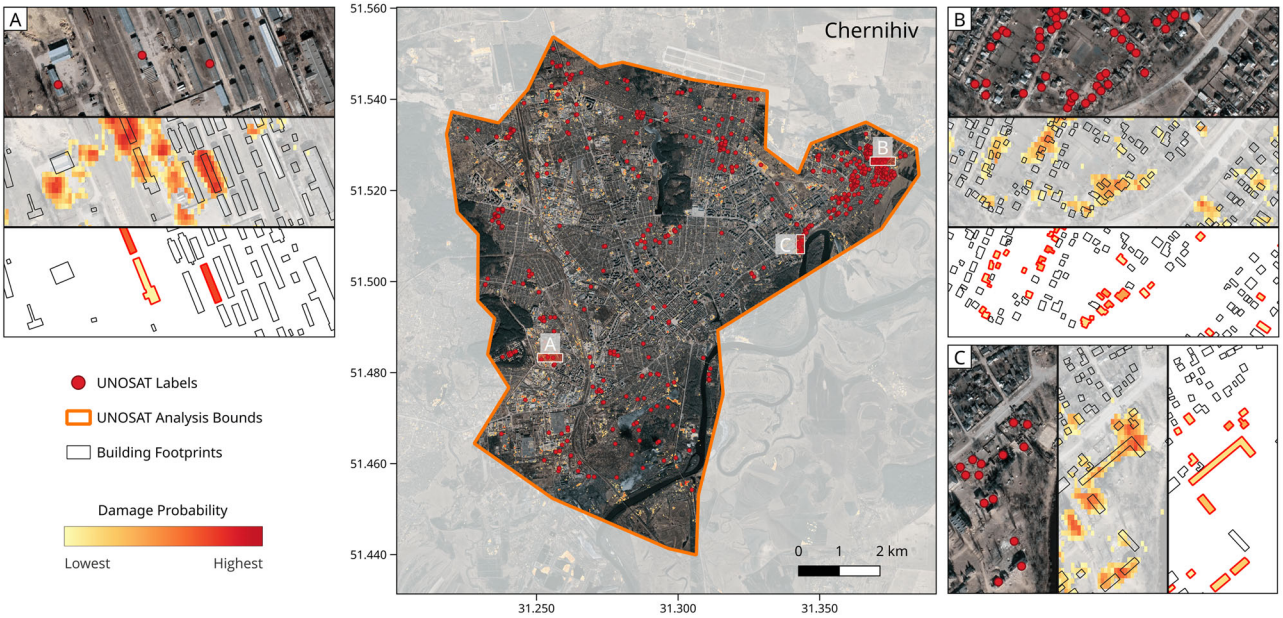


Fig. 2 | Damage estimates and UNOSAT labels for Chernihiv. Building damage estimate for Chernihiv thresholded with the standard confidence threshold of 0.655, aggregated over the first two years of conflict. For clarity, we only present the damage heatmap in the main figure, while three zoomed insets show building-level predictions and highlight some known limitations, such as missing building footprints (Inset A), issues with detecting all small structures (Inset B), and disagreements between UNOSAT points and building footprints (Inset C). Red dots represent UNOSAT annotations indicating buildings marked as either *destroyed* or *severely damaged*. The VHR satellite layer is displayed solely for visualization, all results are derived from 10 m-resolution Sentinel-1 images. Readers are encouraged to explore the maps on the interactive dashboard. Sources: Google Earth/Maxar Technologies, Overture Maps building footprints, UNITAR/UNOSAT damage annotations.



Fig. 3 | The rapid damage mapping tool. Screenshot from the dashboard interface, displaying damage estimates for a region in Mariupol, alongside the Sentinel-1 time series visualizer.

restrictions on download size). Furthermore, the dashboard includes a visualization tool that enables pixel-by-pixel examination of Sentinel-1 time series data at any desired location.

Country-wide analysis

To illustrate the scalability of our approach, we have run damage assessments over the entire country. Our analysis indicates that over 400,000 buildings in Ukraine, or approximately 2.7% of all buildings, have likely sustained damage during the first two years of conflict. This estimate only accounts for buildings larger than 50 m^2. Figure 4 shows the percentage of buildings likely damaged within the first two years of the war, aggregated by *hromadas*, the finest administrative division in Ukraine covering the entire territory. For an analysis of how damage has evolved over time, see Supplementary Note 3. For the most impacted *hromadas* with over 10,000 buildings—Bakhmutska, Marinska, and Popasnianska—damage rates reach 42.6%, 29.0%, and 25.0%, respectively. We observe that this pattern of destruction correlates with the course of the war, the eastern side of the country having been a battlefield for months, while the western part has

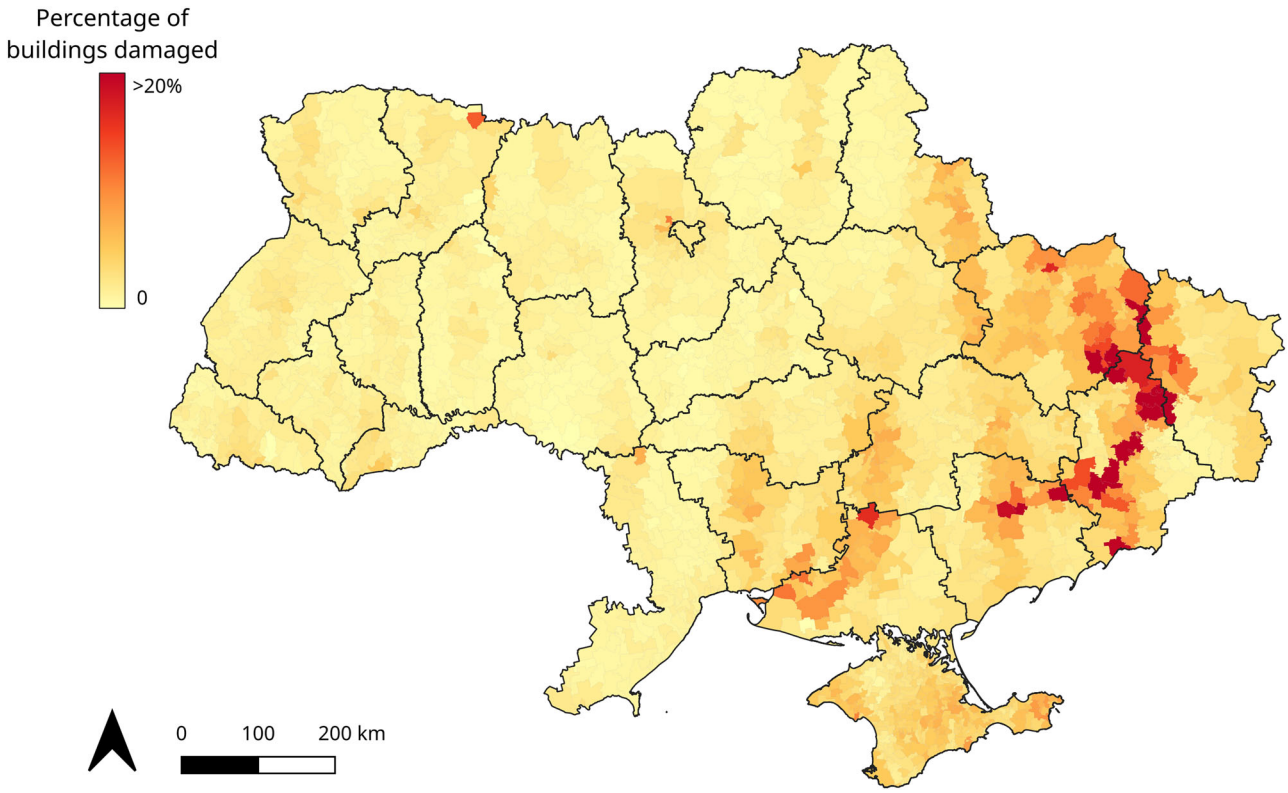


Fig. 4 | Spatial distribution of building damage in Ukraine after two years of conflict. Percentage of buildings likely damaged within the first two years of the war, aggregated by *hromadas*. The predictions were thresholded at 0.655, and only buildings larger than 50 m^2 were considered.

been relatively spared. For a subset of ≈ 1.8M buildings in our dataset, we were able to retrieve meta-data about the building function (such as "residential" or "medical") from the OpenStreetMap catalog. Table 3 summarizes the damaged fractions for different building functions. In this preliminary analysis, we did not find any country-wide patterns that would suggest that certain building types are more prone to war damage than others. Note a more detailed study is needed to confirm or reverse this finding, as biases in the definition and distribution of the available OpenStreetMap building function labels could influence the result.

Discussion

Knowing where and when war-related building damage has occurred is essential for humanitarian organizations to assist populations affected by armed conflicts. The open-access remote monitoring tool that we introduce in this article allows users to rapidly gain an overview of war-related impact. It can also supplement interactive damage mapping and serve as a pre-filtering step to guide reliable, but resource-intensive photo-interpretation by human operators. We have demonstrated the effectiveness of our tool by conducting a comprehensive building damage assessment across Ukraine, showing how war-related destruction has spread throughout the country. The retrieved distribution of damaged buildings broadly reflects the frontlines of the conflict, with the most severe destruction concentrated in areas that experienced prolonged and intense fighting.

While our method has so far been tested only in Ukraine, we plan to extend its application to other regions. Given the characteristics of both the Sentinel-1 SAR data and the Random Forest classifier, we anticipate that the methodology will adapt well to new areas, as long as we have adequate data for fine-tuning. However, the final performance will ultimately be influenced by factors such as vegetation cover and building density.

Limitations

While our tool for screening conflict-related building damage proved effective, it is important to acknowledge its limitations and interpret its outputs accordingly. In this context, we again highlight that the purpose of our method is to complement existing damage mapping solutions with a scalable screening component, not to replace human photo-interpretation or ground-based damage reports.

Table 3 | Number of buildings damaged per OSM class across Ukraine

OSM Class	# Damaged	% Total
Residential	57,089	3.6
Industrial	3177	5.0
Outbuilding	599	1.6
Commercial	1309	3.9
Education	658	3.5
Agricultural	775	4.1
Service	284	3.6
Religious	177	2.1
Medical	126	3.6
Civic	117	4.6
Transportation	106	4.9
Entertainment	8	2.6
Military	6	4.1

The footprints with OSM class information represent a subset of 11.9% of all building footprints analyzed.

Classification threshold. One key limitation is the sensitivity of the quantitative estimates to the selected confidence threshold. As a pragmatic solution, our tool provides a slider that lets the user adjust the threshold according to the precision/recall trade-off needed for a specific use case. For example, a higher threshold may be appropriate to gather

spatially coarse aggregate statistics about destruction patterns as in 4, whereas a lower threshold might better support the search for individual, hitherto overlooked damages.

Google Earth engine. Our study focuses on large-scale assessment and accessibility, leveraging GEE's computational capabilities. However, this reliance also imposes significant constraints. First, it restricts the choice of machine learning methods, as only a limited set of classification algorithms are available at the time of writing. Second, while Sentinel-1 data processed by GEE is orthorectified, it is not radiometrically terrain-flattened, which can introduce distortions in areas with pronounced slopes. Although most cities in Ukraine are situated on relatively flat terrain, our method may be less reliable in towns and cities located in mountainous regions. Additionally, as previously mentioned, GEE currently does not support the ingestion of complex SAR data, preventing the use of phase coherence for change detection. Lastly, while the proposed method is computationally efficient in principle, the web tool may take several minutes to generate results due to GEE's technical overhead. As an alternative, users can download the relevant image tiles and process them locally using the provided code. In the future, we hope to further accelerate the online tool.

Dependence on external data. Our model has specifically been trained on building data and therefore presupposes that settlement or building boundaries are available to mask out areas without human settlements (e.g., agriculture, forest, etc.), which would otherwise cause many false positives. We found that it is sufficient to utilize standard, global land cover maps[47], settlement layers[48], or building footprints. Still, the reliance on such external sources is a potential source of error (see Supplementary Note 4 for possible issues specific to building footprints).

Choice of reference period. To have a common pre-event reference for both negative and positive samples, we used the year 2020 as a fixed reference period throughout this work. This means, however, that the temporal gap between the pre-event and post-event periods can be rather large; giving rise to misclassifications due to surface changes other than war damage, e.g., construction. Ideally, a future version of the tool would integrate context knowledge about the conflict to minimize the temporal gap between the pre- and post-event observation periods.

Reference data. Our supervised learning approach necessitates a sufficient volume of reference data. Despite existing calls for large-scale datasets and open data programs dedicated to armed conflicts[13], such resources remain limited. Data availability varies considerably between different contexts: for some conflicts, e.g. the wars in Ukraine and Gaza, quite a lot of damage data is available; whereas there may be very little information about building damages in conflicts that receive less public attention. It is precisely in these under-reported conflicts that automated analysis could have the biggest impact. By offering an accessible, open-access tool based on public satellite data, we hope to take one step toward a global monitoring solution. But to get there it will likely be necessary to collect at least some training data that represent the building patterns and imaging conditions of regions with under-reported conflicts.

Policy implications

The combination of free Earth observation data and machine learning offers a promising route toward automated war damage monitoring, complementing interactive mapping and on-the-ground surveys.

But there is inherent uncertainty in predicting damage probabilities from SAR image patterns at 10 m GSD, and users should understand this and align their expectations with what such methods can deliver. Our approach is perhaps best suited to obtain a high-level overview of how damage spreads across time and space. Beyond scientific goals like tracing the development of the conflict, this should often be enough to understand where populations are affected. At the same time, our method cannot confirm damage to an individual building, meaning that it is not directly comparable to assessments based on VHR imagery, ground-based photography, witness reports, etc. Importantly, the automatic large-scale approach and reliable local evidence are complementary: satellite damage maps can identify regions that warrant the effort to gather detailed, reliable evidence.

Post-processing the model outputs with additional data sources, as demonstrated for Ukraine with building footprints, adds an important dimension of information. Future research could seek to assimilate further sources of information, for instance, news reports or social media. Despite the inherent biases of such sources[49], they contain a wealth of complementary information that could be harvested, adding contextual knowledge and narrative depth to maps derived purely from satellite images. For example, attacks on residential buildings carry different policy implications if these buildings are inhabited than if civilians have fled an area in anticipation of strikes. Of course, even when data sources are combined to enhance the information derived from satellite images, contextual knowledge offered by observers and humanitarian organizations remains crucial for decision-making.

Beyond immediate humanitarian action, our large-area war impact maps are potentially useful also to support other efforts to remedy the consequences of armed conflicts, such as planning, prioritization, and resource allocation during reconstruction efforts. In terms of scientific inquiry, wall-to-wall, spatially explicit data allows for fine-grained geospatial analysis of conflict dynamics, which at present relies almost entirely on events derived from news articles. Building damage information obtained through our approach can serve as an additional data source that is complementary to the more commonly used metric of battle-related fatalities.

Methods

Data

All the data used in this work can be found freely available online.

Reference damage assessments. To the best of our knowledge, the only public, georeferenced datasets that provide building-level damage information for Ukraine are the UNOSAT maps[11]. They are based on commercial VHR imagery, where analysts manually compare pre- and post-event images using a standardized damage scale. To maximize the chances that these damages produce discernible changes in the SAR imagery, we only keep the two most severe damage levels, i.e., *destroyed* or *severely damaged*, and discard the others. We aggregate all available UNOSAT maps for Ukraine into a dataset consisting of 10,934 unique entries, distributed across 18 different areas of interest (AOIs) as depicted in Fig. 1. Importantly, these labels are attached to geo-coordinates, without any relation to concrete events. We, therefore, lack information on both the spatial extent and the precise date of the destruction; all that is known is that it happened between the pre- and post-event dates. For the present work, we further restrict this interval by assuming that no damage occurred before the onset of the invasion on February 24, 2022. We point out that this assumption may not strictly hold in the eastern regions, particularly Luhansk and Donetsk oblasts. A detailed summary of our dataset can be found in Supplementary Note 5.

SAR image time series. We use freely available SAR data from the Copernicus Sentinel-1 mission[39], consisting of two satellites equipped with C-band SAR instruments and moving in the same orbital plane, phased 180° apart. This setup originally allowed for global coverage with a revisit period below 6 days anywhere on Earth. Unfortunately, this period has doubled to up to 12 days since the technical failure of Sentinel-1B in December 2021. For our study, we only use data from Sentinel-1A to maintain consistent temporal resolution. We utilize the Ground Range Detected (GRD) product available on GEE, consisting of log-amplitudes for the VV and VH polarizations, resampled at 10 m. Internally, GEE

automatically preprocesses every tile with precise orbit correction, border and thermal noise removal, radiometric calibration, and terrain correction. We do not add any further preprocessing.

Building footprints. We leverage the building layer from the Overture Maps Foundation's Open Map Data[45], which combines building footprints from diverse sources. In Ukraine, this equates to a total of 26.5M buildings, with 24.7% coming from OpenStreetMap's crowd-sourced database[50] and the remaining 20M from Microsoft's Global Building Footprints[51]. We note here that these building footprints have their own limitations. E.g., OpenStreetMap data may be outdated, and Microsoft footprints, retrieved from VHR optical imagery with a deep learning model, may suffer from missing buildings and geometric inaccuracies, particularly for small buildings. Considering the resolution of Sentinel-1 we only take into account the 15M buildings that exceed 50 m^2 in surface area, i.e., they are larger than half a pixel. Supplementary Note 4 illustrates some of these limitations.

Machine learning framework

An overview of our machine learning framework is shown on Fig. 5. For a detailed ablation study of our key design choices, see Supplementary Note 6.

Definition of time intervals. For clarity, we first define the temporal acquisition windows used in our study. Time interval T_0 covers one year from Feb. 24, 2020, until Feb. 23, 2021. The subsequent intervals T_1 to T_{12} represent consecutive 3-month time windows, with T_1 ranging from Feb. 24, 2021, to May 23, 2021, and the last interval T_{12} spanning Nov. 24, 2023, to Feb. 23, 2024. As such, intervals T_1 to T_4 represent the year preceding the invasion, while T_5 to T_{12} represent the two years following it.

Time series extraction. For each location marked by UNOSAT, we stack the corresponding Sentinel-1 tiles and extract all backscatter values, for both the VV and VH polarizations. Limiting the training to the points annotated by UNOSAT corresponds to a fully supervised learning scheme where the target value is known for every training example. Relying only on pixel-wise time series discards all spatial context. We found that the time series signal carries most of the information relevant to our analysis, presumably because conflict-induced building damages are mostly small and localized. Per-pixel processing also scales particularly well on parallel processing architectures like the one underlying GEE.

Importantly, Sentinel-1 has the typical side-looking viewing geometry of SAR sensors, where the same location may look very different depending on the incidence angle and orbit direction (see Supplementary Note 7 for examples). We, therefore extract the backscatter values independently for each orbit, resulting in 2 to 4 time series per UNOSAT reference point. Overall, we obtain 33,304 distinct time series, denoted as x_i^o, where i denotes the ID (respectively, location) and o is the orbit index.

Classification. We formulate damage mapping as a supervised, binary classification. From each time series, we extract two segments: the fixed reference period $x_{i,\text{ref}}^o$ and the assessment period $x_{i,\text{new}}^o$. The task of the classifier is to estimate the probability that a war-related change has occurred in the time window between the two segments. For each segment and each band, we extract a set of N statistical features and combine them into a final feature vector denoted as $\phi(x_i) \in \mathbb{R}^{4N}$ that forms the input to the classifier. This strategy has several advantages. First, by fixing the reference period before the start of the war one can ensure that it represents the state without any war-induced damages. Second, extracting fixed-length segments makes the method independent of the overall conflict duration. Third, using multiple different assessment periods greatly increases the number of training examples and also helps to make the model robust against seasonal variations.

As a classification algorithm, we use a Random Forest as implemented in the SMILE library[52], since that algorithm is available off-the-shelf in GEE, facilitating large-scale deployment. For every input feature vector the model outputs a score $\hat{y}_i^o \in [0, 1]$, which can be interpreted as the likelihood that war-induced damage has happened between the associated reference and assessment periods. At inference time, we compute the overall damage probability map by averaging the estimates from different orbits, $\hat{y}_i = \frac{1}{N_o}\sum \hat{y}_i^o$.

Training details. We perform all computations directly in GEE. We train the model on the four AOIs with the highest numbers of reference labels, Mariupol, NW Kyiv, Rubizhne, and Volnovakha. Together they represent 75.0% of all UNOSAT annotations and 82.6% of all backscatter time series. The remaining fourteen AOIs are reserved for evaluation, ensuring geographic diversity (see Fig. 1). We choose T_0 as a fixed reference period and randomly use periods from T_1 to T_8 (e.g., spanning one year before and after the beginning of the war) as assessment periods.

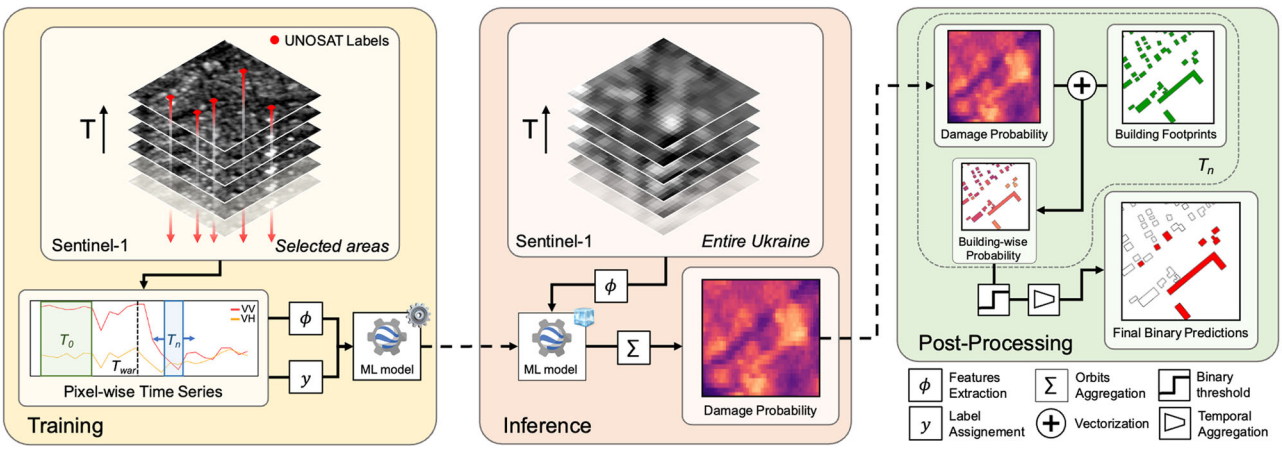


Fig. 5 | Overview of the machine learning framework. For training (left), we use per-pixel Sentinel-1 time series extracted at the location of UNOSAT point annotations. The model is fed with a pair of time series from the same location. The first one spans a fixed 12-month time interval T_0 from 2020, and the second one spans one of the 3-month time intervals T_n between 2021 and 2023. Both time series are encoded with a custom features extractor, and damage labels are dynamically assigned according to T_n. At inference time (center), the model generates a damage probability heatmap valid at T_n and spanning the entire country, aggregating the predictions of different Sentinel-1 orbits. The raw damage probabilities are intersected with building footprints. For the final map the estimates for different time intervals T_n are thresholded and aggregated.

The label y_i is dynamically assigned to each x_i based on the end date $t_{\max}$ of the assessment period:

$$y_i = \begin{cases} 0 & \text{if } t_{max} \le t_{\text{invasion}}, \\ 1 & \text{if } t_{max} > t_{\text{unosat}}, \\ -1 & \text{otherwise} \end{cases} \quad (1)$$

where t_{invasion} is 2022-02-24, the day the invasion was launched; and t_{unosat} is the acquisition date of the post-event image used for annotation by UNOSAT. All time series with $y_i = -1$ are discarded because we cannot determine whether they contain the damage event. We transform x_i into $\phi(x_i)$ using the following seven summary statistics: minimum, maximum, mean, median, standard deviation, kurtosis, and skewness. Finally, we configure the Random Forest to use 50 decision trees, with *minLeafPopulation* set to 3 and *maxNodes* set to 10,000. These values ensure sufficient robustness while remaining within the computational budget permitted by GEE.

Country-wide inference

After model training, we leverage the parallel computing capabilities of GEE to generate damage probability maps for the entire country. We compute one map for each period seen during training, as well as four additional maps for T_9 to T_{12}, the assessment periods between 2023 and February 2024. Each map, in EPSG:4326 projection and stored in UInt8 format, has an extent of 88,867 × 201,284 pixels and a file size of 17.88GB.

Post-processing

To quantify the impact of the war on the Ukrainian building stock, we cross-reference our maps with the building footprints sourced from Overture Maps[45]. For each building b_j and each period T_n, we assign a damage likelihood $\hat{y}_{j,T_n}$ by averaging the likelihood values of all pixels that fully or partially overlap the footprint, weighted by the overlap fraction:

$$\hat{y}_{j,T_n} = \sum_i w_{ij} \hat{y}_{i,T_n} \quad (2)$$

Here $\hat{y}_{i,T_n}$ represents the model output at pixel i for the period T_n, and w_{ij} is the proportion of that pixel that falls within b_j. By construction, this process also discards all damage estimates that fall outside of buildings, e.g. in agricultural areas or forests.

To obtain a final estimate of the number of buildings impacted over the first two years of the war, we aggregate the maps for the relevant assessment periods with the following rule:

$$\hat{y}_j = \begin{cases} 1 & \text{if } \max(T_{[5,12]}) \ge t \text{ and } \max(T_{[1,4]}) < t \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Data availability

All data used in this study is publicly accessible. Links to all results, including damage heatmaps and building footprints with corresponding damage estimates, can be found in the repository hosting our code and/or in Zenodo (https://zenodo.org/records/14811504/).

Code availability

All the code necessary to reproduce our results can be found in the following repository: https://github.com/prs-eth/ukraine-damage-mapping-tool/.

Received: 24 June 2024; Accepted: 3 March 2025;
Published online: 19 March 2025

References

1. Kyiv School of Economics. Russia Will Pay Project https://kse.ua/about-the-school/news/155-billion-the-total-amount-of-damages-caused-to-ukraine-s-infrastructure-due-to-the-war-as-of-january-2024/ (2024).
2. World Bank. Ukraine - Third Rapid Damage and Needs Assessment (RDNA3): February 2022 - December 2023. World Bank Group http://documents.worldbank.org/curated/en/099021324115085807/P1801741bea12c012189ca16d95d8c2556a (2023).
3. Gueguen, L. & Hamid, R. Large-scale damage detection using satellite imagery. In IEEE (ed.) *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1321–1328 https://doi.org/10.1109/cvpr.2015.7298737 (2015).
4. Xu, J. Z., Lu, W., Li, Z., Khaitan, P. & Zaytseva, V. Building damage detection in satellite imagery using convolutional neural networks https://arxiv.org/abs/1910.06444 (2019).
5. Lee, J. et al. Assessing post-disaster damage from satellite imagery using semi-supervised learning techniques https://arxiv.org/abs/2011.14004 (2020).
6. Zheng, Z., Zhong, Y., Wang, J., Ma, A. & Zhang, L. Building damage assessment for rapid disaster response with a deep object-based semantic change detection framework: From natural disasters to man-made disasters. *Remote Sens. Environ.* **265**, 112636 (2021).
7. Gholami, S. et al. On the deployment of post-disaster building damage assessment tools using satellite imagery: A deep learning approach. In IEEE (ed.) *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*, 1029–1036 https://doi.org/10.1109/ICDMW58026.2022.00134 (IEEE, 2022).
8. Kaur, N., Lee, C., Mostafavi, A. & Mahdavi-Amiri, A. Large-scale building damage assessment using a novel hierarchical transformer architecture on satellite images. *Computer-Aided Civ. Infrastruct. Eng.* **38**, 2072–2091 (2023).
9. Gupta, R. et al. xBD: A dataset for assessing building damage from satellite imagery https://arxiv.org/abs/1911.09296 (2019).
10. Sticher, V., Dietrich, O., Pfeifle, B. & Wegner, J. D. Watching armed conflicts from space. *CSS Analyses in Security Policy*. https://css.ethz.ch/content/dam/ethz/special-interest/gess/cis/center-for-securities-studies/pdfs/CSSAnalyse336-EN.pdf (2024)
11. UNOSAT. United Nations Satellite Center. https://unosat.org/about-us.
12. Mueller, H., Groeger, A., Hersh, J., Matranga, A. & Serrat, J. Monitoring war destruction from space using machine learning. *Proceedings of the National Academy of Sciences* **118** https://doi.org/10.1073/pnas.2025400118 (2021).
13. Bennett, M. M., Van Den Hoek, J., Zhao, B. & Prishchepov, A. V. Improving satellite monitoring of armed conflicts. *Earth's Future* **10** https://doi.org/10.1029/2022EF002904 (2022).
14. MAXAR. MAXAR Open Data Program. https://www.maxar.com/open-data.
15. Planet. Planet Disaster Data Program. https://www.planet.com/disasterdata.
16. Nabiee, S., Harding, M., Hersh, J. & Bagherzadeh, N. Hybrid U-Net: Semantic segmentation of high-resolution satellite images to detect war destruction. *Mach. Learn. Appl.* **9**, 100381 (2022).
17. Singh, D. K. & Hoskere, V. Post disaster damage assessment using ultra-high-resolution aerial imagery with semi-supervised transformers. *Sensors* **23**, 8235 (2023).
18. Hou, Z. et al. War city profiles drawn from satellite images. *Nat. Cities* **1**, 359–369 (2024).
19. Plank, S. Rapid damage assessment by means of multi-temporal SAR -" a comprehensive review and outlook to Sentinel-1. *Remote Sens.* **6**, 4870–4906 (2014).
20. Ge, P., Gokon, H. & Meguro, K. A review on synthetic aperture radar-based building damage assessment in disasters. *Remote Sens. Environ.* **240**, 111693 (2020).
21. Aoki, H., Matsuoka, M. & Yamazaki, F. Characteristics of satellite SAR images in the damaged areas due to the Hyogoken-Nanbu earthquake. In ACRS (ed.) *Proc. 19th Asian Conf. Remote Sens*, vol. 7, 1–6 (1998).

22. Matsuoka, M. & Yamazaki, F. Use of satellite SAR intensity imagery for detecting building areas damaged due to earthquakes. *Earthq. Spectra* **20**, 975–994 (2004).
23. Matsuoka, M. & Yamazaki, F. Building damage mapping of the 2003 Bam, Iran, earthquake using Envisat/ASAR intensity imagery. *Earthq. Spectra* **21**, 285–294 (2005).
24. Liu, W., Yamazaki, F., Gokon, H. & Koshimura, S.-i Extraction of tsunami-flooded areas and damaged buildings in the 2011 Tohoku-oki earthquake from TerraSAR-X intensity images. *Earthq. Spectra* **29**, 183–200 (2013).
25. Uprety, P., Yamazaki, F. & Dell'Acqua, F. Damage detection using high-resolution SAR imagery in the 2009 L'Aquila, Italy, earthquake. *Earthq. Spectra* **29**, 1521–1535 (2013).
26. Matsuoka, M. & Yamazaki, F. Use of interferometric satellite SAR for earthquake damage detection. *Sat* **2**, z1 (2000).
27. Fielding, E. J. et al. Surface ruptures and building damage of the 2003 Bam, Iran, earthquake mapped by satellite synthetic aperture radar interferometric correlation. *J. Geophys. Res.: Solid Earth* **110** (2005).
28. Gamba, P., Dell'Acqua, F. & Trianni, G. Rapid damage detection in the Bam area using multitemporal SAR and exploiting ancillary data. *IEEE Trans. Geosci. Remote Sens.* **45**, 1582–1589 (2007).
29. Ferretti, A., Prati, C. & Rocca, F. Permanent scatterers in SAR interferometry. *IEEE Trans. Geosci. Remote Sens.* **39**, 8–20 (2001).
30. Yun, S.-H. et al. Rapid damage mapping for the 2015 M w 7.8 Gorkha earthquake using synthetic aperture radar data from COSMO–SkyMed and ALOS-2 satellites. *Seismol. Res. Lett.* **86**, 1549–1556 (2015).
31. Watanabe, M. et al. Detection of damaged urban areas using interferometric SAR coherence change with PALSAR-2. *Earth Planets Space* **68** https://doi.org/10.1186/s40623-016-0513-2 (2016).
32. Washaya, P., Balz, T. & Mohamadi, B. Coherence change-detection with Sentinel-1 for natural and anthropogenic disaster monitoring in urban areas. *Remote Sens.* **10**, 1026 (2018).
33. Tay, C. W. J. et al. Rapid flood and damage mapping using synthetic aperture radar in response to typhoon Hagibis, Japan. *Sci. Data* **7** https://doi.org/10.1038/s41597-020-0443-5 (2020).
34. Mastro, P., Masiello, G., Serio, C. & Pepe, A. Change detection techniques with synthetic aperture radar images: Experiments with random forests and sentinel-1 observations. *Remote Sens.* **14**, 3323 (2022).
35. Akhmadiya, A. et al. Application of GLCM textural based method with Sentinel-1 radar remote sensing data for building damage assessment. In IEEE (ed.) *2022 International Conference on Smart Information Systems and Technologies (SIST)*, 1-5 https://doi.org/10.1109/SIST54437.2022.9945758 (IEEE, 2022).
36. Stephenson, O. L. et al. Deep learning-based damage mapping with InSAR coherence time series. *IEEE Trans. Geosci. Remote Sens.* **60**, 1–17 (2022).
37. Yang, Y. et al. Large-scale building damage assessment based on recurrent neural networks using SAR coherence time series: A case study of 2023 Turkey-Syria earthquake. *Earthq. Spectra* **40**, 2285–2305 (2024).
38. Gorelick, N. et al. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sens. Environ.* **202**, 18–27 (2017).
39. Copernicus Programme, European Space Agency. Sentinel-1. https://sentinels.copernicus.eu/web/sentinel/missions/sentinel-1.
40. Aimaiti, Y., Sanon, C., Koch, M., Baise, L. G. & Moaveni, B. War related building damage assessment in Kyiv, Ukraine, using Sentinel-1 radar and Sentinel-2 optical images. *Remote Sens.* **14**, 6239 (2022).
41. Huang, Q., Jin, G., Xiong, X., Ye, H. & Xie, Y. Monitoring urban change in conflict from the perspective of optical and SAR satellites: The case of Mariupol, a city in the conflict between RUS and UKR. *Remote Sens.* **15**, 3096 (2023).
42. Tavakkoliestahbanati, A., Milillo, P., Kuai, H. & Giardina, G. Pre-collapse spaceborne deformation monitoring of the Kakhovka dam, Ukraine, from 2017 to 2023. *Communications Earth & Environment* **5** (2024).
43. Ballinger, O. Open access battle damage detection via pixel-wise T-Test on Sentinel-1 imagery https://arxiv.org/abs/2405.06323 (2024).
44. Scher, C. & Van Den Hoek, J. Nationwide mapping of damage to human settlements across Ukraine using Sentinel-1 InSAR coherence change detection. In AGU (ed.) *AGU Fall Meeting Abstracts*, vol. 2023, GC23A–01 (2023).
45. Overture Maps Foundation. Open Map Data. https://overturemaps.org.
46. Sticher, V., Wegner, J. D. & Pfeifle, B. Toward the remote monitoring of armed conflicts. *PNAS Nexus* **2** https://doi.org/10.1093/pnasnexus/pgad181 (2023).
47. Brown, C. F. et al. Dynamic World, near real-time global 10m land use land cover mapping. *Scientific Data* **9** https://doi.org/10.1038/s41597-022-01307-4 (2022).
48. Schiavina, M., Melchiorri, M. & Pesaresi, M. GHS-SMOD R2023A - GHS settlement layers, application of the degree of urbanisation methodology (stage I) to GHS-POP R2023A and GHS-BUILT-S R2023A, multitemporal (1975-2030) http://data.europa.eu/89h/a0df7a6f-49de-46ea-9bde-563437a6e2ba (2023).
49. Miller, E., Kishi, R., Raleigh, C. & Dowd, C. An agenda for addressing bias in conflict data. *Scientific Data* **9**, https://doi.org/10.1038/s41597-022-01705-8 (2022).
50. OpenStreetMap contributors. OpenStreetMap. https://www.openstreetmap.org.
51. Microsoft. Global ML Building Footprints. https://github.com/microsoft/GlobalMLBuildingFootprints.
52. Li, H. Smile - statistical machine intelligence and learning engine https://haifengl.github.io (2014).

Acknowledgements

This work was supported in part by ETH4D's Engineering Humanitarian Action initiative through the Remote Monitoring of Armed Conflict project in collaboration with the ICRC, as well as by ESA's Discovery Program under the Humanitarian Monitoring with Copernicus Satellites project (Contract No. 4000138343).

Author contributions

Olivier Dietrich: conceptualization (equal), data curation (lead), investigation (lead), methodology (lead), software (lead), visualization (lead), writing - original draft preparation (lead). Torben Peters: conceptualization (support), investigation (support), methodology (equal), software (support), supervision (equal), writing - original draft (equal). Vivien Sainte Fare Garnot: conceptualization (support), supervision (support), writing - review & editing (support). Valerie Sticher: conceptualization (lead), data curation (support), funding acquisition (lead), project administration (equal), writing - original draft (equal). Thao Ton-That Whelan: funding acquisition (equal) project administration (equal), validation (support), writing - review & editing (support). Konrad Schindler: conceptualization (equal), funding acquisition (lead), methodology (equal), project administration (equal), resources (lead), supervision (lead), writing - review & editing (lead). Jan Dirk Wegner: conceptualization (lead), funding acquisition (lead), methodology (equal), project administration (equal), resources (support), supervision (lead), writing - review & editing (lead).

Funding

Open access funding provided by Swiss Federal Institute of Technology Zurich.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s43247-025-02183-7.

Correspondence and requests for materials should be addressed to Olivier Dietrich.

Peer review information *Communications Earth & Environment* thanks Ritwik Gupta, Eric Fielding and David Sandwell for their contribution to the peer review of this work. Primary Handling Editors: Teng Wang and Joe Aslin. A peer review file is available.

Reprints and permissions information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2025