GeoAI for Damage Assessment - Gaza

UNDP leverages GeoAI to rapidly map destroyed buildings in the Gaza Strip for effective humanitarian response, with 85% model accuracy.

UNDP
May 1, 2025

Introduction

The ongoing conflict in Gaza has caused massive destruction to public, economic, and social infrastructure, severely affecting the people of Gaza. This situation necessitates effective and efficient response plans to address the evolving humanitarian crisis.

Without near real-time data from the ground, effective response is impossible. To address this, UN agencies—including UNOSAT and FAO—are making significant efforts to map damaged buildings and agricultural areas across the Gaza Strip, publishing the results on a monthly basis. However, these updates take time, as the mapping methods rely heavily on manual, human-driven analysis. Given the need for rapid information, it is critical to develop automated approaches to map damaged structures and deliver results in near real-time.

Consequently, UNDP has explored how rapidly GeoAI can assist in mapping damage as new satellite images become available. In addition, rapid damage mapping can help track changes in the landscape and infrastructure over time as a result of the ongoing conflict.

The main objective is to explore and evaluate GeoAI models for mapping damage—primarily destroyed buildings—in the Gaza Strip. This objective is further divided into the following sub-objectives:

1. Select the highest-performing GeoAI model for detecting destroyed buildings.
2. Map the destroyed buildings using the selected GeoAI model.
3. Compare the GeoAI-derived results with UNOSAT's comprehensive damage assessment.
4. Define the way forward for scaling, refinement, and operational deployment of GeoAI in similar contexts.

Methodology

UNDP utilized the ArcGIS Pro Deep Learning Framework and high-resolution Maxar satellite imagery to map destroyed buildings across the entire Gaza Strip in late January 2025, just a few days after the ceasefire came into effect.

The following section presents the methodology followed.

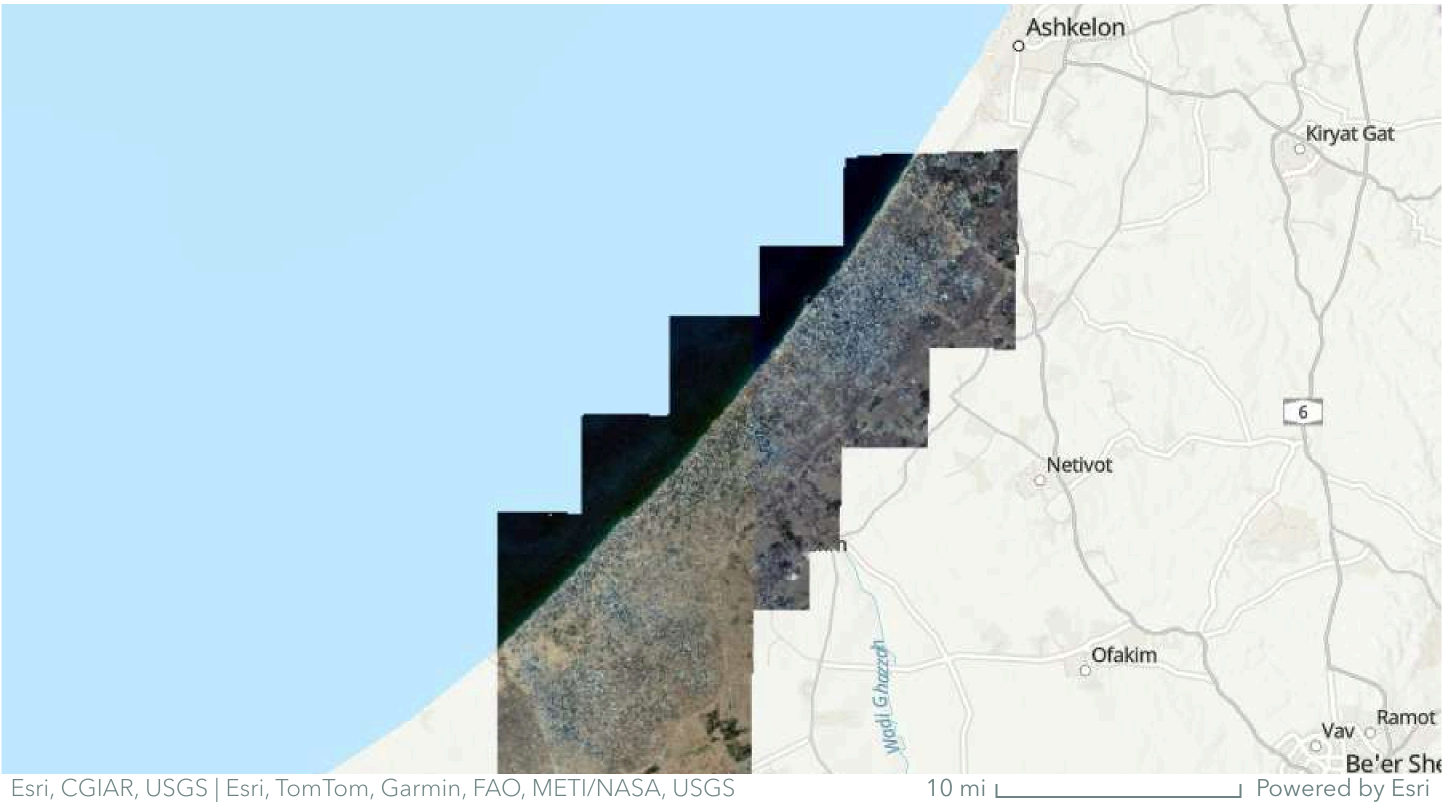

1. Acquire the Image

A satellite image dated October 12, 2024, was obtained from Maxar at a resolution of 50 cm, covering the entire Gaza Strip.

The image was used to create training samples for fine-tuning the GeoAI model.

Zoom in or out on the map to navigate the image.

Esri, NASA, NGA, USGS | Esri, TomTom, Garmin, METI/NASA, USGS | - UN… 5 mi Powered by Esri

2. Create the Samples

- To create well-representative training samples, **13 sampling blocks** across the Gaza Strip were randomly selected. Within each block, training samples were digitized and labeled.

- To create the training samples, **2,239 polygon features** representing destroyed and undestroyed areas were digitized. These samples were classified into two classes: Destroyed (1) and Undestroyed (2), representing destroyed and undestroyed buildings, respectively. The model treats all other background features—such as roads and green areas—as a third class with a value of 0.

- **Map Legend**

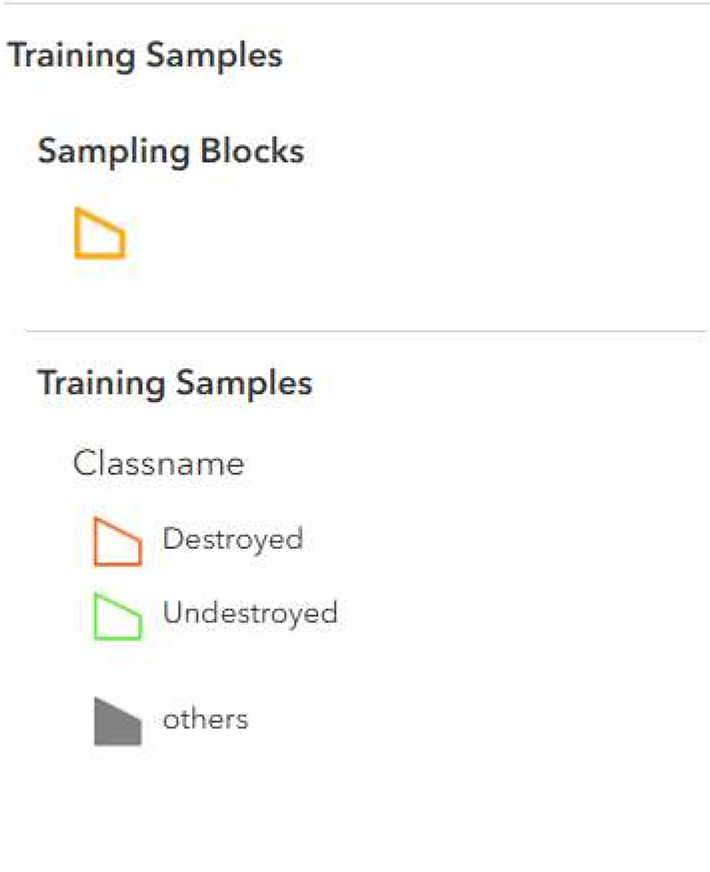
Training Samples
Sampling Blocks
Training Samples
Classname
Destroyed
Undestroyed
others

Map legend

3. Export the Samples

- For each sampling block, the corresponding portion of the input satellite image (captured on October 12, 2024) was segmented and exported as 256×256-pixel image tiles to meet the GPU memory constraints of the training environment.

- Matching labeled sample polygons were simultaneously exported as mask tiles of the same size, ensuring spatial alignment with the image tiles for deep learning model training.

- The screenshot on the right showcases a sample batch of exported image tiles, each overlaid with its corresponding classification mask. Red areas represent buildings classified as Destroyed, while green areas indicate those identified as Undestroyed.

- This step resulted in the generation of **1,561 image tiles**, along with an equal number of corresponding mask tiles, ensuring a one-to-one correspondence for model training.

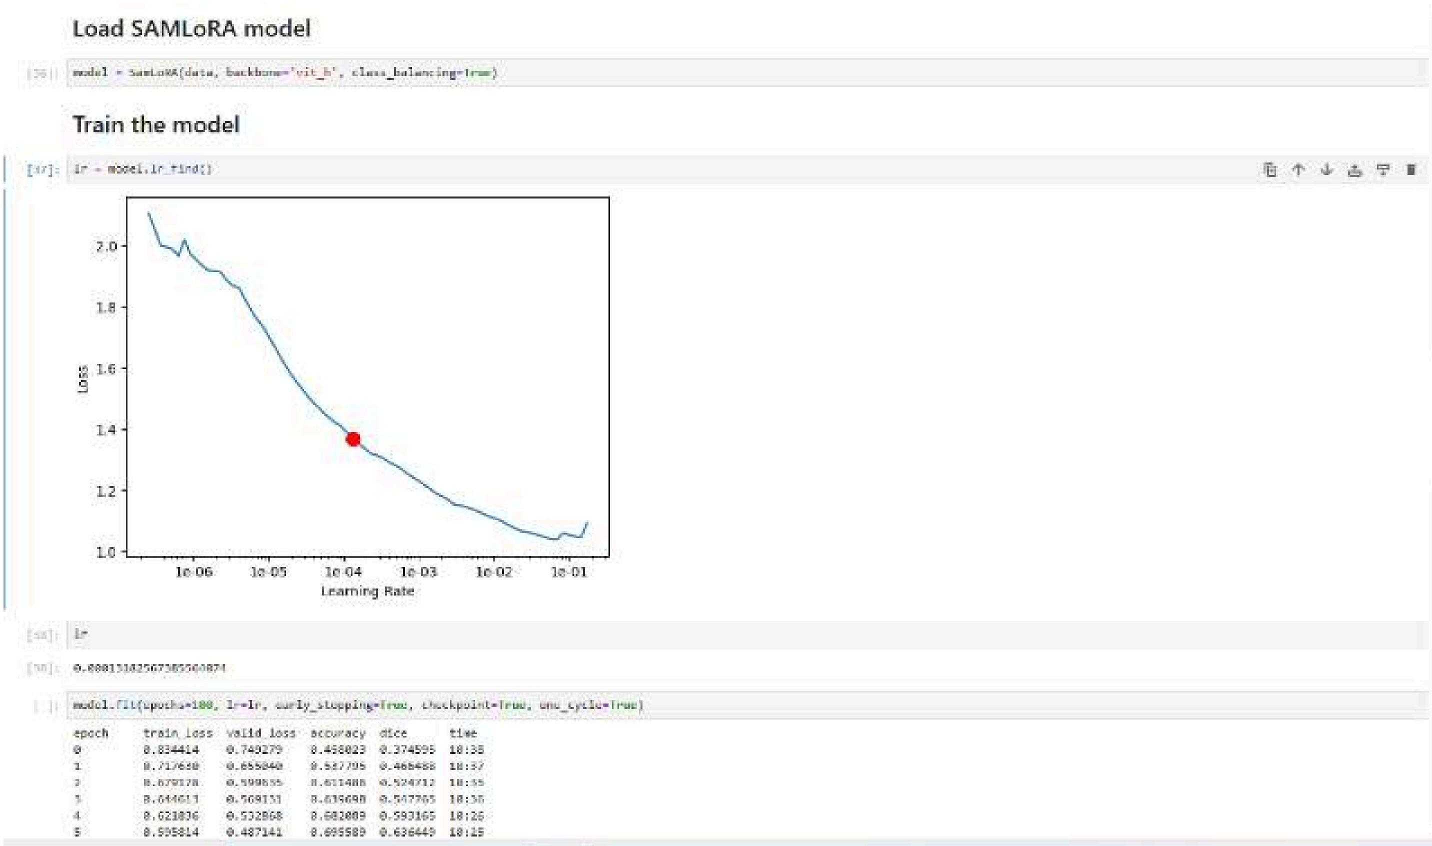

4. Fine-tune the GeoAI model

- The Esri pretrained **SAMLoRA model** was selected to segment the satellite image into destroyed and undestroyed building areas.

- The **SAMLoRA** model combines the capabilities of **Segment Anything (SAM)** and **Low-Rank Adaptation (LoRA)**, enabling precise object segmentation with minimal labeled data. It is optimized for high-resolution imagery and supports rapid fine-tuning for specialized tasks such as damage detection.

- As shown on the right, **a Jupyter Notebook** was developed to manage the entire process of model fine-tuning and evaluation.

Esri, … 500 ft Powered by Esri

Esri, … 500 ft Powered by Esri

5. Evaluate and use the model

- To evaluate the performance of the model, 20% of the training samples were held out as testing samples.

- As shown below, the model achieved strong performance in identifying both destroyed and undestroyed buildings. For the Destroyed class, it achieved a precision of 85%, meaning that when the model identifies a building as destroyed, it is correct 85% of the time. It also achieved a recall of 88%, indicating that it successfully detects 88% of all buildings that are truly destroyed. This suggests the model is slightly better at finding all destroyed buildings—an especially important capability in crisis response.

	Undestroyed	**Destroyed**
Precision	0.84	0.85
Recall	0.84	0.88
F1 Score	0.84	0.87

Model Performance on testing samples

- The **F1 Score** combines both precision and recall into a single measure of overall performance. A higher F1 Score reflects a better balance between accuracy and completeness. Overall, the model performs very well, particularly in detecting destroyed buildings, which is the primary focus of this damage assessment.
- The model was then used to classify a Maxar satellite image dated on **January 25, 2025**, a few days after the ceasefire came into effect (January 19, 2025).

- The map on the right shows the final layer of destroyed areas found by the model. Swipe left or right to show the destroyed areas as detected by the model.

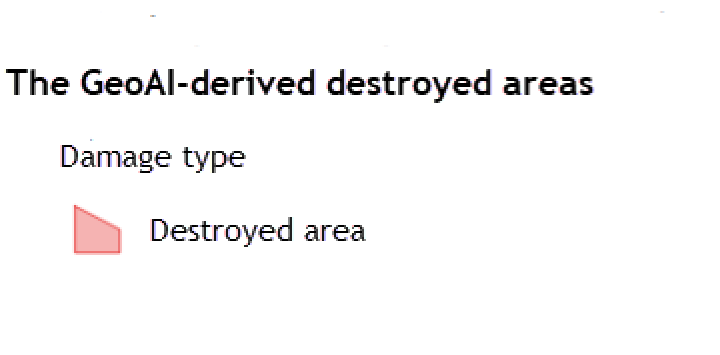

Map legend

Esri, NASA, NGA, USGS | © OpenStreetMap, Microsoft, Esri, TomTom, Gar… 500 ft Powered by Esri

6. Postprocessing the output layer

In this final step, two postprocessing procedures were implemented to produce the final layer of destroyed buildings

- Very small areas, represented as isolated pixels and generated as a side effect of the segmentation process, were removed during postprocessing, as they do not represent actual damaged structures.

- The segmentation layer has been then overlaid by the Microsoft Buildings Layer to classify each building into destroyed or undestroyed building.

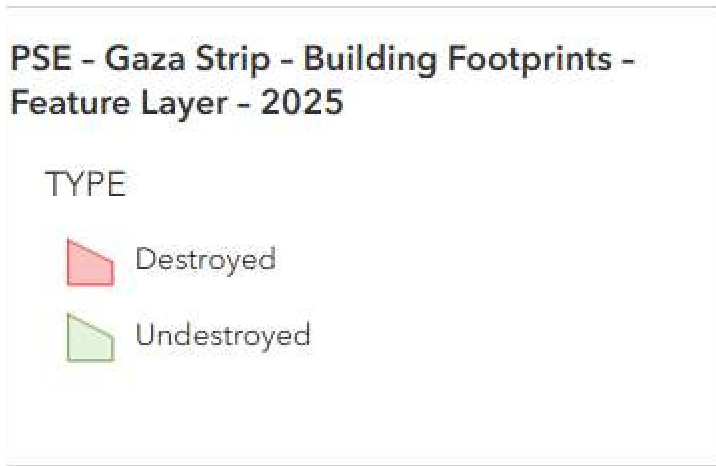

Map legend

Results

Geoal Assessment Dashboard

The below dashboard presents summary statistics about the destroyed buildings in the Gaza Strip.

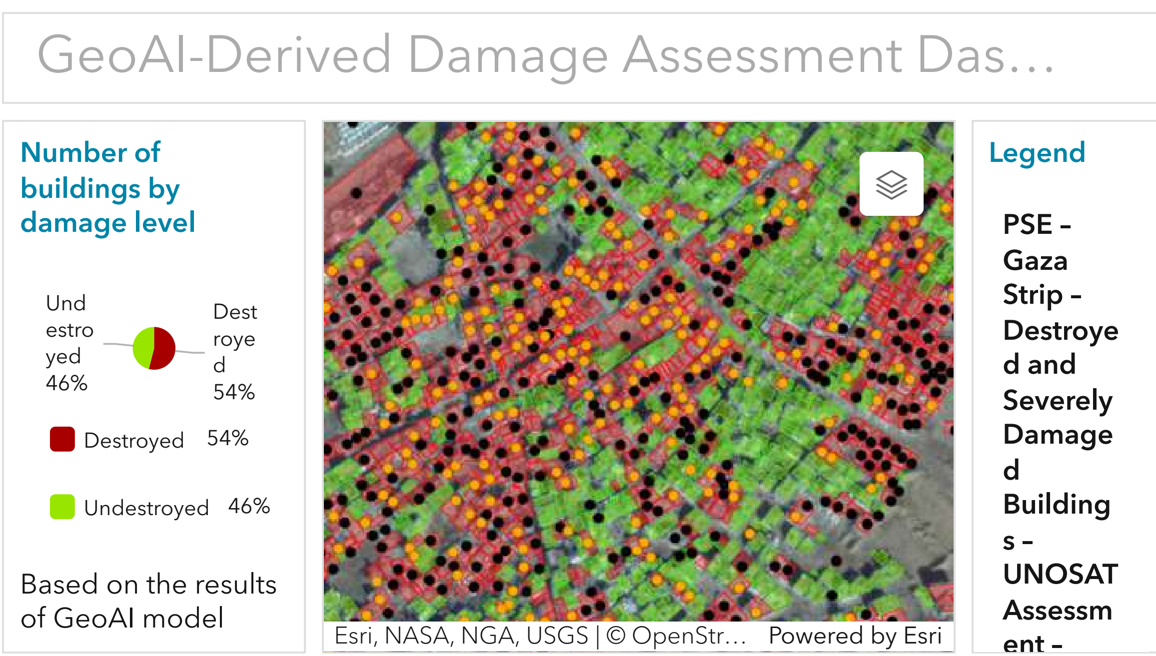

GeoAI vs. UNOSAT assessment

One effective way to evaluate the performance of the GeoAI model is by comparing its results with those of the UNOSAT damage assessment, which relies on expert visual interpretation of high-resolution satellite imagery. The GeoAI-derived damage layer was generated using satellite imagery captured on January 25, 2025, during the ceasefire period. To ensure temporal alignment for comparison, the comprehensive damage assessment conducted by UNOSAT on February 25, 2025—using imagery from a similar period—was selected as the reference benchmark, where only the destroyed and severely damaged buildings were used for comparison.

The map below shows the destroyed (red) and undestroyed (green) building footprints detected by the GeoAI model, alongside the UNOSAT destroyed (black) and severely damaged (orange) building centroids. It can be observed that most of the destroyed and severely damaged centroids fall within or near the destroyed building footprints, indicating a high level of agreement between the GeoAI and UNOSAT assessments.

According to UNOSAT, approximately **87,700** buildings were destroyed or severely damaged, while the GeoAI model identified **89,892** buildings as destroyed. This difference may be attributed to variations in the building datasets used by each method, as well as false positives introduced by the GeoAI model.


Esri, … 500 ft Powered by Esri | Esri, … 500 ft Powered by Esri

The way Forward

This experiment in using GeoAI for damage assessment has demonstrated that GeoAI is a promising approach. It can be both effective and efficient in crisis contexts, as it allows for automation and near real-time results to support humanitarian actors and assist in response and early recovery interventions.

UNDP plans to upgrade the model to include additional damage classes aligned with UNOSAT's damage classification system. Accordingly, UNOSAT data can be considered as ground truth for future experiments.

UNDP also plans to generalize the model and fine-tune it for application in other regions, as is currently being done for damage mapping in Haiti.

Credits

GeoAI Model	UNDP BMS/ITM
Satellite Images	Maxar©2024, Maxar©2025

Damage Assessment data	UNITAR/UNOSAT
Language review	OpenAI (ChatGPT)