Detecting War Destruction in Ukraine using Sentinel-1 Time-Series

Olivier Dietrich[1,*], Torben Peters[1], Vivien Sainte Fare Garnot[2], Konrad Schindler[1], Jan Dirk Wegner[2]

[1]ETH Zurich, [2]University of Zurich
[*]odietrich@geod.baug.ethz.ch

The devastating impacts of armed conflicts on populations and the ensuing humanitarian crisis call for effective conflict monitoring. Satellite imagery has emerged as a powerful tool in this endeavor, its non-invasive capabilities offering a window into remote or high-risk areas, synergizing well with on-ground activities. Traditional methods for detecting conflict-induced events, notably building destruction, hinge on manually analyzing very high-resolution (VHR) satellite images. This approach is time-consuming, labor-intensive, and typically dependent on irregularly acquired and costly commercial satellite imagery, thus limiting the scalability and affordability of such monitoring efforts.[1]

Modern computer vision techniques have shown high potential in identifying damaged buildings from satellite images. However, most research relies on the use of commercial VHR images and predominantly focuses on damage stemming from natural disasters, which are usually spatially concentrated [2, 3]. In stark contrast, destruction from armed conflicts is typically scattered, both spatially and temporally, with VHR images often being unavailable. In this work, we aim to demonstrate the viability of using open satellite imagery for war-induced destruction detection, despite their coarser spatial resolution. We use human-annotated labels from the ongoing conflict in Ukraine to train a machine learning model on time-series of Sentinel-1 images.

We chose Sentinel-1 imagery due to its high revisit frequency, the inherent resilience of SAR images against atmospheric disruptions, and the characteristically strong and stable scatterers that buildings typically exhibit, which allow the extraction of reliable time-series. Within these series, we anticipate that persistent changes over time, such as war-induced destruction, should be easier to detect. Yet, unlike with optical imagery, these signals can look dramatically different based on the satellite's incidence angle and orbit direction, and each satellite pass must therefore be assessed separately. As a result, a single destruction event might appear differently in various time-series. To counteract this, we aggregate the predictions for each site post-inference. We show that this post-processing step is crucial to reach trustworthy results.

We perform extensive analysis with different models, features, and aggregation methods, and show that a simple Random Forest on pixel-wise time-series with a mean aggregation can achieve an f1-score up to 0.84, outperforming pure statistical baselines. We deploy our model in Google Earth Engine to make it readily available to the humanitarian community.

References

[1] Valerie Sticher, Jan D Wegner, and Birke Pfeifle. "Toward the remote monitoring of armed conflicts". In: *PNAS Nexus* 2.6 (May 2023). Ed. by Erik Kimbrough. DOI: `10.1093/pnasnexus/pgad181`. URL: `https://doi.org/10.1093/pnasnexus/pgad181`.

[2] Jihyeon Lee et al. *Assessing Post-Disaster Damage from Satellite Imagery using Semi-Supervised Learning Techniques*. 2020. DOI: `10.48550/ARXIV.2011.14004`. URL: `https://arxiv.org/abs/2011.14004`.

[3] Hannes Mueller et al. "Monitoring war destruction from space using machine learning". In: *Proceedings of the National Academy of Sciences* 118.23 (June 2021). DOI: `10.1073/pnas.2025400118`. URL: `https://doi.org/10.1073/pnas.2025400118`.