Watching Armed Conflicts from Space

Leveraging open-access satellite images with deep learning can help humanitarian and human rights organizations address violent conflicts more rapidly. To ensure the effectiveness of remote monitoring, organizations need to align their strategies with the technical constraints of deep learning-based monitoring systems.

By Valerie Sticher, Olivier Dietrich, Birke Pfeifle and Jan Dirk Wegner

From the buildup of Russian tanks along the Ukrainian border to images of destruction raging through cities and villages, we all watched the war in Ukraine unfold with images from the sky. When Russia bombed the Mariupol theater in March 2022, it was not primarily the photographs of the victims that garnered global attention and outrage. Instead, it was a satellite image of the theater captured shortly before the attack. The image revealed the word "children" in Russian, written twice outside the theater and clearly visible even from space. Russian tactical fighter aircraft attacked the theater anyway. For many, the attack came to symbolize Russia's blatant disregard for civilian lives in the war in Ukraine, a powerful demonstration of how satellite images can profoundly affect public discourse.

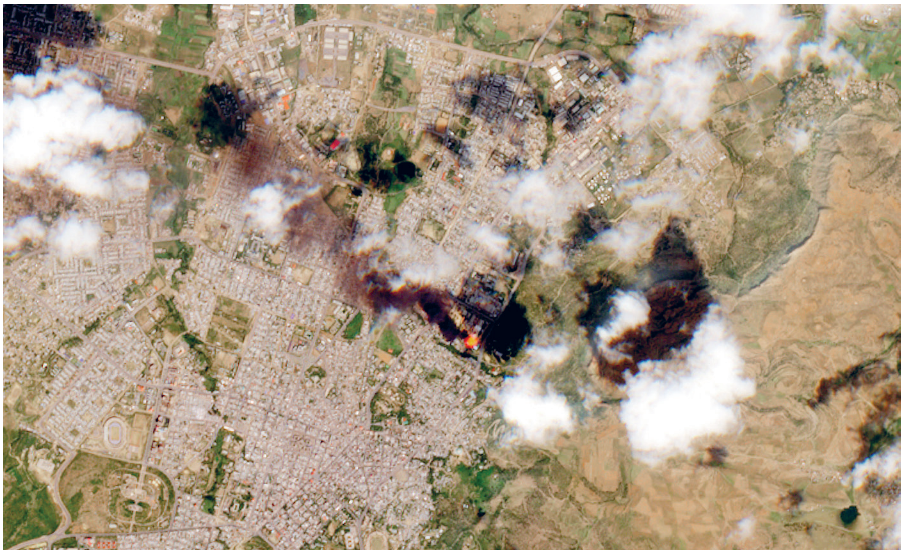

Smoke rising from Mekelle, Ethiopia on 20 October 2021. *Copernicus Sentinel 2 imagery, processed by the EU DG for Defense Industry and Space, access via Reuters.*

Throughout their history, satellite images have influenced the course of armed conflicts. However, with the proliferation of commercial satellite providers, satellite images are increasingly used not only by those who fight wars but also those who report on them, and those who seek to mitigate the damage they inflict. In contrast to other remote sensing instruments, such as drones, satellites are much less intrusive, as they are operated from space. Human rights and humanitarian organizations like the International Committee of the Red Cross (ICRC) or Amnesty International have come to rely on satellite images as an essential tool to understand what is happening in areas that are otherwise inaccessible. They use satellite images for human rights investigations, for example, to determine which conflict actor was in control of an area when an atrocity took place. More generally, satellite images help these organizations understand where fighting is happening, prioritize humanitarian assistance, and monitor humanitarian corridors for the safe passage of refugees and aid.

To date, most organizations rely on spatial analysts to meticulously examine high spatial resolution images for changes on the ground, such as damage to buildings or the destruction of infrastructure. However, the labor-intensive nature of manual image annotation limits the capacity of these organizations to comprehensively screen vast conflict areas. In view of these constraints, and as the world's attention focuses on a few select conflicts, numerous conflicts go largely unmonitored. This is precisely where progress in deep learning could make an important impact.

This CSS Analysis explores the potential of automating satellite imagery analysis through deep learning, a development that could enable organizations to monitor conflicts more systematically and on a larger scale. The article highlights current challenges facing the use of automated analysis. These challenges include the costs involved in screening with high spatial resolution images and how verifying conflict events identified by automatic systems can quickly become overwhelming. The article suggests that actors should prioritize the development of applications that rely primarily on open-access satellite images and align their expectations with the technical constraints of deep learning-based systems. This is to enable them to reap the benefits of automatic analysis without creating new bottlenecks. While remote monitoring can support the activities of human rights and humanitarian actors, other forms of observation will remain crucial.

Comparison of Commercial Satellite Imagery and Open-Access Imagery

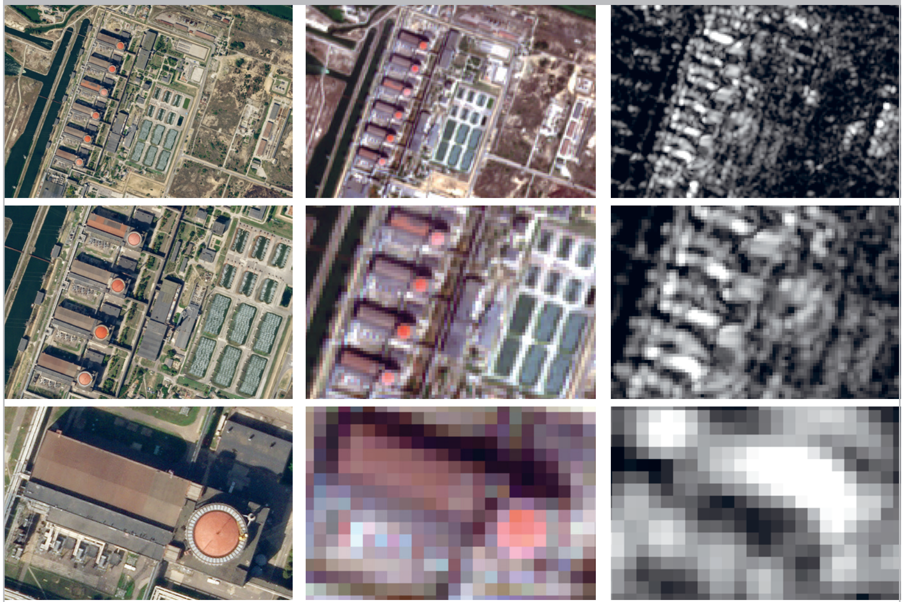

Optical satellite images from a commercial provider *(© Planet Labs PBC/Reuters).*

Optical open-access images *(Copernicus Sentinel-2, processed by ESA).*

SAR open-access images *(Copernicus Sentinel-1, processed by ESA).*

The Promise of Remote Monitoring

The term "remote monitoring of armed conflicts" refers to the continuous screening of entire conflict regions through satellite images. The aim of such monitoring is to detect the impact of conflict events, such as the destruction of buildings, within a few days. Satellite imagery is also used to map the effects of disasters like earthquakes or floods. However, armed conflicts often extend over much longer periods of time than disasters. This means conflicts often require continuous screening rather than one-off mappings. This could make the use of automation in the remote monitoring of armed conflicts more beneficial than in disaster situations. However, it also makes it more challenging.

Advancements in deep learning and the expansion of computing power present new prospects for such automation. Supervised techniques may be particularly interesting for remote monitoring. These techniques learn from reference data annotated by humans and apply learned insights to new data. A key advantage of such systems is that they can be expanded to entire countries and potentially the globe. The training of deep learning models that scale to entire conflict regions and beyond requires large amounts of manually annotated reference data that cover a wide range of possible scenarios. Once set up and carefully validated by experts, systems based on deep learning can be run on a regular basis with little extra effort. This allows for the continuous near-real time screening of conflict areas rather than sporadic mappings.

Despite the promise of deep learning, only a handful of pilot projects have been implemented by human rights and humanitarian actors to date. One example is Amnesty International's exploration of AI in satellite data analysis to detect the destruction of human settlements in Darfur. There are two important reasons for the lack of real-life applications at a larger scale. The first is the focus on high spatial resolution imagery. The second is the misalignment of expectations about what deep learning systems can and cannot deliver.

Conflicts often require continuous screening rather than one-off mappings.

Types of Satellite Images

A notable obstacle to developing real-world remote monitoring applications is that models are usually trained to work with high spatial resolution imagery. The preference for this type of image intuitively makes sense: It is easier to detect objects, such as trenches or damaged buildings, if images display a high level of detail. At the same time, high spatial resolution images bring a host of challenges to automated screening solutions. The most obvious challenge relates to the costs of acquisition. Most commercial satellite providers use a business model that allows customers to choose between tasking a satellite to take images of a specific location or selecting images from its archive. Other providers offer subscription-based models that give users access to continuously acquired images of a specific region. Subscription-based models are costly and tasking a satellite to take images of a location even more so. Archival images are cheaper, but there is no guarantee that a specific location has been covered in the past. Moreover, acquisition costs are only one part of the problem. There are also the more hidden costs of handling high spatial resolution images, which requires massive processing power and high download bandwidth.

Open-access images with a lower spatial resolution present a promising alternative. For example, the European Space Agency's

(ESA) Sentinel satellite constellations continuously and densely map the globe, including remote conflict areas, by capturing images of this type. The images provided by these constellations are publicly accessible, free of charge, and made available at regular intervals for almost the entire globe. The processing of moderate spatial resolution images requires significantly less download bandwidth and computing power compared to high spatial resolution imagery.

A major drawback of open-access images is their larger ground sampling distance (GSD), the distance between the center of two adjacent pixels as measured on the ground. In the case of the ESA's Sentinel-2 images, the GSD is 10 meters, whereas some commercial providers offer satellite images with a GSD as low as 30 centimeters. The larger GSD of open-access images makes it challenging to use them to identify objects and changes occurring on the ground. This is evident in the example of the Zaporizhzhia nuclear power plant in Enerhodar, Ukraine, depicted on p.2. Each of the six power plant buildings, visible to the left of the pictures on the top row, is approximately 190 meters long and 70 meters wide. Each building could be represented by as many as 53,200 pixels in commercial imagery with a 50 centimeters GSD. In contrast, optical images from the Sentinel constellation may depict one building with just 133 pixels. Accordingly, details of these buildings are much more difficult to discern in open-access images compared to commercial high spatial resolution images. Damage to smaller buildings may not even be discernible by the human eye in open-access satellite images.

Deep learning models, however, can identify subtle patterns in satellite data that would not be apparent to the human eye. A particularly promising data source for large-scale mapping efforts is imagery taken by synthetic aperture radar (SAR) satellites. SAR satellites actively emit a microwave signal instead of relying on light visible to the human eye, as is the case with optical satellites. This means SAR data is still usable even when there is cloud coverage. Unlike optical images, which capture reflections of the sun's radiation, SAR images capture back-scattered signals that are emitted by the SAR sensor. Even small objects can lead to significant changes in the recorded signals. Therefore, it should eventually become possible for SAR data to indicate the presence of destruction caused by conflict, even for damage that is much smaller than a pixel in a SAR image. Moreover, SAR data provides additional information, so-called phase information, which enables a technique called interferometric SAR (InSAR). This technique is already widely used in science and practice, such as in measuring the impact of earthquakes. When applied to a specific range of electromagnetic frequencies, InSAR methods allow for the measurement of changes at the centimeter level, even though the ground sampling distance is much larger than this.

SAR data, and its potential combination with existing land cover maps and optical satellite images, makes it possible to detect changes that may not be otherwise observable in open-access images. This potential makes working with open-access images an attractive alternative to the costly and resource intensive work involved with using high spatial resolution images. Nonetheless, high-resolution satellite imagery will continue to play a crucial role in the remote monitoring of armed conflict, especially for organizations aiming to further investigate the nature and impact of damage initially identified in moderate resolution images.

The Challenge of Verification

A second important challenge of using deep learning to analyze satellite images relates to uncertainties inherent in the output it generates. Systems based on algorithms can, at best, only indicate the probable location and time of damage or other impacts resulting from conflict, along with an estimated likelihood that such impacts occurred. In other words, these systems cannot indicate that an event happened with absolute certainty. However, in many situations, organizations will want to be certain that an event took place. In these cases, organizations that use deep learning tools will need to manually verify that an event flagged by the system—such as an airstrike on a civilian building—has indeed taken place. This could involve using high spatial resolution imagery or sources on the ground.

For remote monitoring applications, the need to manually verify each flagged event can quickly become overwhelming if mappings are done on a regular basis. This is because even when war is raging, the prevalence of conflict events—such as the destruction or damage of buildings—is relatively low. For example, according to estimates from The Economist, more than 40 per cent of the built-up area in the Ukrainian city of Mariupol was destroyed in the first four months of the war in Ukraine. Yet even with this enormous level of destruction, only a small fraction of the city's buildings were destroyed each week. In most conflict contexts, the number of unchanged and undamaged structures outweighs that of destroyed buildings to a much greater degree than in Mariupol, especially when violence occurs sporadically. In such scenarios, even a low false positive rate results in a significant number of buildings being flagged as likely having been damaged even though no damage took place. This poses a substantial challenge for organizations aiming to verify each model output and may even increase their workload rather than facilitate their efforts.

Policy Alignment

While there are strategies and techniques to refine the performance of deep learning models, the challenge of verification cannot be addressed solely on a technical level. It also requires a paradigm shift in our understanding of the most effective uses for conflict screening and an acknowledgment of what may not yet be achievable. At least three types of applications are feasible, all focusing on uses that reduce the need for individual verification.

The first types of applications are those that help organizations get a sense of where fighting is happening, and what damage it inflicts. Rather than focusing on individual, verifiable conflict events, the deep learning models would highlight areas where damage accumulates. This can guide manual mapping efforts or, more broadly, increase an organization's situational awareness,

Further Reading

Mia Bennett et al., **"Improving Satellite Monitoring of Armed Conflicts,"** *Earth's Future* 10:9 (2022).

Hannes Mueller et al., **"Monitoring War Destruction from Space Using Machine Learning,"** *PNAS* 118:23 (2021).

Valerie Sticher / Aly Verjee, **"Do Eyes in the Sky Ensure Peace on the Ground? The Uncertain Contributions of Remote Sensing to Ceasefire Compliance,"** *International Studies Review* 25:3 (2023).

Valerie Sticher / Jan Dirk Wegner / Birke Pfeifle, **"Toward the Remote Monitoring of Armed Conflicts,"** *PNAS Nexus* 2:6 (2023).

helping it to prioritize humanitarian aid.

The second types of applications are models that detect unusual trends, either temporally or geographically. By alerting users to such trends, these models can serve as an early warning mechanism. Because they focus on trends rather than individual events, they can be set up to screen large conflict areas across the globe. Organizations will still have to verify alerts manually. However, alerts would only be raised if a trend rather than an individual event were identified, guiding human experts where to look for evidence of rising tensions.

The third types of applications are systems that make use of secondary criteria to reduce the need for manual verification. For example, humanitarian actors may be particularly interested in monitoring critical infrastructure, such as hospitals and schools. Applications could be designed to issue alerts only if damage is detected in pre-specified locations. Each alert would have to be verified, but the number of alerts a system issues could be kept at a manageable level by using pre-defined criteria.

Deep learning models are better at detecting some forms of violence, like air strikes or shelling, than others, such as cattle raids.

All these applications are particularly interesting for conflict areas that receive relatively little attention, such as the Central African Republic or Chad. Information about events in these conflict areas may not be picked up by the media, and human rights and humanitarian actors often lack the resources to monitor such contexts systematically. Unfortunately, these are also the types of places where there is no comprehensive "ground truth" data, such as manually annotated images that would allow for the training of deep learning models for these specific conflicts. Acquiring such data is crucial for effectively applying the technical possibilities of deep learning in practice.

Outlook

Humanitarian and human rights organizations can use open-access satellite imagery in conjunction with deep learning models to respond to violent conflicts more swiftly and effectively. To realize this potential, applications should be developed that consider the technical limitations of deep learning systems, the operational limitations of organizations, and the sensitive environments in which they are deployed. Moreover, it is essential for researchers to collaborate closely with human rights and humanitarian organizations. This is to ensure that data and code are shared in a manner that facilitates the development and application of technologies to alleviate the impact of armed conflicts, while minimizing the potential for harm.

Researchers, along with human rights and humanitarian organizations, should also contemplate the long-term impact of remote monitoring solutions on armed conflicts. Deep learning models are better at detecting some forms of violence, like air strikes or shelling that cause heavy damage, than others, such as cattle raids. Some forms, like gender-based violence, will never be detectable through satellite images. As remote monitoring grows more prevalent, the disparities in detection capabilities could influence the actions of human rights and humanitarian actors. This could even potentially prompt a change in tactics by those perpetuating violence. Other forms of observation, including working with actors on the ground, remain crucial. By developing and using remote monitoring thoughtfully, considering both its potential and limitations, we can ensure that it becomes an effective tool to reduce suffering in armed conflicts, without opening up new avenues for instigating violence.

For more on perspectives on Mediation and Peace Promotion, see CSS core theme page.

Valerie Sticher is a visiting scholar at the School of Advanced International Studies at Johns Hopkins University.
Olivier Dietrich is a PhD student at the Photogrammetry and Remote Sensing Group at ETH Zurich.
Birke Pfeifle is a PhD student at the Center for Security Studies (CSS) at ETH Zurich.
Jan Dirk Wegner holds the "Data Science for Sciences" chair as associate professor at the Department of Mathematical Modeling and Machine Learning (DM³L) at the University of Zurich.

This policy brief is based on the following article: Valerie Sticher / Jan Dirk Wegner / Birke Pfeifle, "Toward the Remote Monitoring of Armed Conflicts," PNAS Nexus 2:6 (2023). The underlying research is part of the Remote Monitoring of Armed Conflicts project, a collaborative effort between the EcoVision Lab at ETH Zurich and the University of Zurich, the Center for Security Studies at ETH Zurich, and the ICRC, funded by the Engineering for Humanitarian Action Initiative.

CSS Analyses in Security Policy is published by the Center for Security Studies (CSS) at ETH Zürich. The CSS is a center of competence for Swiss and international security policy. Each month, two analyses are published in German, French, and English.

Editor: Névine Schepers
Language editing: Mike Woods
Layout and graphics: Miriam Dahinden-Ganzoni, Rosa Guggenheim

Feedback and comments: analysen@sipo.gess.ethz.ch
More editions and online subscription: www.css.ethz.ch/cssanalyses

Most recent editions:

ISSN: 2296-0244; DOI: 10.3929/ethz-b-000656792