# Using High Resolution Imagery and Neural Networks to Measure Destruction and Reconstruction after a Disaster

Elizabeth Frankenberg
Keenan Karrigan
Peter Katz
Eric Peshkin
Cecep Sumantri
Duncan Thomas

Climate change in combination with sea level rise is intensifying the force and frequency of extreme events that damage the natural environment and the health and wealth of populations across the globe. In 2017 three major hurricanes hit the U.S. in short succession, upending the lives of millions of Americans. Hurricane Florence and Typhoon Mangkhut have done the same, on opposite sides of the globe. The cleanup and rebuilding efforts after these sorts of disasters cost billions and typically  lasts for several years.

Developing tools for assessing and measuring both the extent of damage and the pace of recovery and linking these measurements to data on populations is a critical frontier at the intersection of population-environment linkages. Because satellite imagery is routinely collected on an on-going basis it  is ideally-suited for documenting change before and after an unanticipated disaster, even in the most remote locations. The objective of this paper is to combine high resolution satellite imagery, machine learning tools, and survey and census data in the context of a large-scale natural disaster to demonstrate proof-of-concept for a set of techniques that will enhance social scientists' ability to cost-effectively and efficiently evaluate the long-term impacts of extreme events on the environment and on population well-being.

## The Event

The event we study is the 2004 Indian Ocean Tsunami. The tsunami was triggered by a massive megathrust earthquake. The 1,500 km rupture generated tsunami waves that impacted shorelines throughout the region and propagated wave energy into all other world oceans *(1-3)*. Worldwide casualties totaled over 250,000. The western coastline of the Indonesian province of Aceh was the area hit hardest. Tsunami waves averaging 25-30m struck Aceh's shore some 15 minutes after the quake *(4-5)*. In Aceh, 160,000 people, roughly 5% of the population, perished. The height, force, and inland reach of water from the tsunami depended on topographical features of the ocean floor and the shoreline *(6-7)*. Low-lying coastal communities were largely destroyed in some areas. Along riverways the water encroached as much as 9 kilometers, versus 3-4 kilometers in other nearby locations *(7)*. Inland areas that were flooded sustained substantial damage to the environment but few deaths.

In the tsunami's aftermath a $7 billion recovery effort was launched, at the time the largest, most ambitious rebuilding effort ever undertaken in a developing country, and one that was quite successful by many metrics. As one example, 140,000 homes were built in the five years after the event, to replace the estimated 120,000 that were damaged or destroyed. USAID rebuilt a major transportation artery along Aceh's challenging west coast, and many other groups contributed to the reconstruction of schools, health facilities, village halls, and mosques.

The triptych of high resolution imagery below illustrates the transformation of one community, first by the tsunami waves, where green has been replaced by bare earth, and then by the rebuilding effort, in which the road, bridge, housing, and cropped fields have reappeared. The same process was repeated in hundreds of communities along Aceh's west coast. Harnessing the information contained in these images and linking it to population data offers a wealth of opportunities for understanding how populations and the environment around them rebound in the aftermath of a disaster.

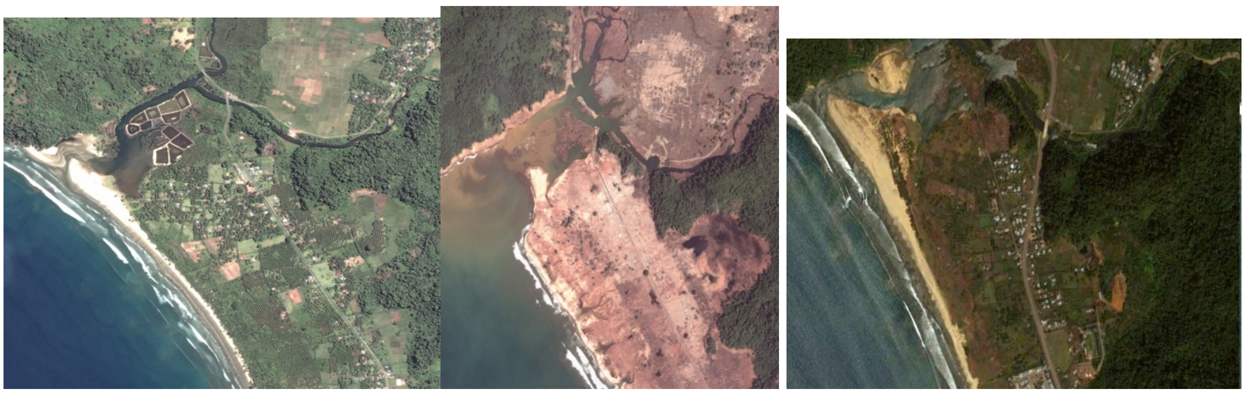

2004, pre-tsunami 2005 two weeks post-tsunami 2010, five years later

**The Population Data**
We draw on two primary sources of population data. The first is the Study of the Tsunami Aftermath and Recovery (STAR). STAR is a longitudinal survey of individuals, households and communities. The baseline, conducted 10 months before the tsunami, interviewed over 33,000 respondents of all ages in nearly 500 communities (for more information see stardata.org). The STAR sample includes 13 *kabupaten* that border the coast in Aceh and North Sumatra. Tsunami survivors and their children born after the tsunami were tracked and interviewed annually for 5 years after the disaster and again in a 10 year follow-up in 2015.

Attrition poses special challenges in a study of a major disruptive event. Nearly two-thirds of the population in areas that were most affected by the 2004 tsunami was displaced; some of those people returned to their pre-tsunami locations, others moved multiple times. It is critical to measure the impact of the disaster on a sample that is representative of the entire post-tsunami population, including those who moved. We have, therefore, worked extremely hard to track every surviving baseline respondent including movers to other provinces in Indonesia and to Malaysia. By designing, testing and implementing sophisticated follow-up procedures, we have successfully re-interviewed over 96% of the survivors and, therefore, *our study sample is representative of the population of survivors immediately after the tsunami.*

The second data source is census data of the two provinces conducted in 2005 (a special post-tsunami census) and 2010 (the regular Indonesian decennial census). The questionnaires are short but the extensive geographic coverage provides a useful counterpart to the STAR survey when examining how populations respond to changes in their environment, and how they contribute to environmental "reshaping" after a disaster.

**Harnessing Satellite Imagery**

The images displayed above provide a compelling snapshot of the extent of change to landcover over a six year period. Harnessing the full potential of high resolution imagery to characterize how events that change the environment affect landcover requires methods to examine and code visually observable changes in a way that scales efficiently. We have developed a machine learning approach to classification of landcover that allows us to quantify features of landscapes before, just after, and five to ten years after the tsunami.

The tool we designed uses a principled image preprocessing pipeline and a version of the SegNet convolutional neural network architecture described by Badrinarayanan and others (8). We settled on this method after testing and rejecting several others (boosted random forest, linear discriminant classifier, quadratic discriminant classifier, and K nearest neighbor). Our preprocessing pipeline uses mathematical tools to augment the available information in the imagery data (Gram-Schmidt pansharpening, normalization, and the application of affine transformations). The SegNet approach involves teaching the network to recognize particular patterns by providing it with a training data set and then applying the trained network to new images, eliminating the need for manual labelling. Variations of this approach have been used successfully for other building recognition tasks (Ghaffarian and Ghaffarian 2014; Yuan 2016).

As a first step we obtained images from Digital Globe via their grants program for the years 2004, 2005, 2007, and 2009. We developed a training dataset that includes over 9 million manually labelled building pixels, based on images from the subdistricts of Banda Aceh (the capital of Aceh province) and Aceh Besar (the surrounding more rural subdistrict) (Peshkin 2018). It is important to note that the images we are working with depict buildings with far less regular patterns than the images other research has relied on. This is particularly true for images from the aftermath of the tsunami (see image below), making our task considerably more complex than previous applications. On the other hand, reconstruction housing is typically far more recognizable (third panel).

| | | |
|---|---|---|
| Building images from prior work (Ghaffarian and Ghaffarian 2014) | Aceh, 2005 | Aceh ~2009 (reconstruction) |

**Preliminary Results**

To this point we have used our tool to capture buildings using images from 2005 (just after the tsunami) and 2007, when some rebuilding was underway. The images below provide an example of how a high-resolution image on the left appears after the network has labelled the buildings (right image).

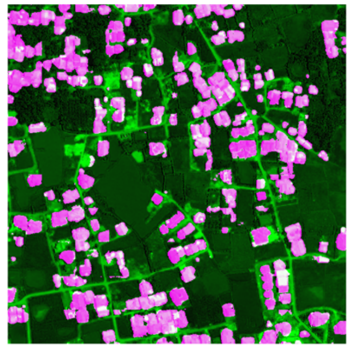

The standard summary measure of performance for neural networks applied to computer vision problems is the F1 statistic, which is the harmonic mean of a measure of recall (the share of relevant items that are detected, which declines as false negative increases) and a measure of precision (the share of detected items that relevant, which declines as false positives increases). This measure is computed by comparing labels created by the network to manually-labeled images that were not used as part of the training data. By these measures our method is quite effective, even for the highly disordered imagery from 2005. In addition our method is highly efficient, classifying 1.1 pixels per minute versus about 200 pixels per minute, on average, by human coders.

| **Segmentation** | **Recall** | **Precision** | **F1** |
|---|---|---|---|
| ISPRS Gold Standard (Kaiser et al 2017) | 0.91 | 0.84 | 0.88 |
| Ghaffarian and Ghaffarian (2014) | 0.88 | 0.72 | 0.78 |
| Yuan (2016) | 0.80 | 0.81 | 0.80 |
| ***Our method*** | 0.85 | 0.77 | 0.81 |
| 2005 | 0.84 | 0.74 | 0.78 |
| 2007 | 0.86 | 0.80 | 0.83 |

**Next Steps**

Having demonstrated the basic feasibility and accuracy of this approach, there are several obvious next steps that are underway and that will be completed in the coming months. Working with our proof-of-concept labelling described above, we will develop measures of surface area covered by buildings at different time points and georeference the measures so that we can compute measures of change over time in share of land cover accounted for by buildings, for known administrative units. Additionally we will link those measures to the communities covered in the STAR survey, which will allow us to link changes over time in building coverage to outcomes and behaviors for STAR respondents. Finally, we are in the process of extending our classification tool so that rather than capturing only buildings, we can capture six separate categories: water, agriculture, rubble, foundations, buildings, and roads. This expansion will substantially enrich the ways we can use the data, for example linking change in land use to changes in population composition, sources of livelihood, and other social and economic outcomes.

1. C. Subarya, M. Chlieh, L. Prawirodirdjo, J-P Avouac, Y. Bock, K. Sieh, A. Meltzner, D. Natawidjaja, and R. McCaffrey. Plate-boundary deformation associated with the great Sumatra-Andaman earthquake. *Nature*. **440**, 46-51 (2006).
2. V. Titov, A. Rabinovich, H. Mofjeld, R. Thomson, F. Gonzalez. The Global Reach of the 26 December 2004 Sumatra Tsunami. *Science*. **309**, 2045-48 (2005).
3. D. Arcas, V. Titov. Sumatra tsunami: lessons from modeling. *Surv Geophys* **27**, 679-705 (2006).
4. K. Monecke, W. Finger, D. Klarer, W. Kongko, B. McAdoo, A. Moore, S. Sudrajat. A 1,000-year sediment record of tsunami recurrence in northern Sumatra. *Nature*.
5. **453,** 1232-1234 (2008).
6. C.M. Rubin, B.P. Horton, K. Sieh, J. Pilarczyk, P. Daly, N. Ismail, A.C. Parnell. Highly variable recurrence of tsunamis in the 7,400 years before the 2004 Indian Ocean tsunami. *Nat Commun*. **8**, 16019 (2017).
7. H. Degueldre, J. Metzger, T. Geisel, R. Fleischmann. Random focusing of tsunami waves. *Nature Physics*. **12**, 259-262 (2016).
8. M. Umitsu, C. Tanavud, B. Patanakanog. Effects of landforms on tsunami flow in the Plains of Banda Aceh, Indonesia, and Nam Khem, Thailand, *Mar Geol*. **242**, 141-153 (2007).
9. Badrinarayanan 2016 Badrinarayanan, V., Kendall, A., Cipolla, R. (2016). SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. https://arxiv.org/pdf/1511.00561.pdfm, Oct 2016
10. Ghaffarian, S., and Ghaffarian, S. (2014). Automatic Building Detection Based on Purposive FastICA (PFICA) Algorithm Using Monocular High Resolution Google Earth Images. ISPRS Journal of Photogrammetry and Remote Sensing, 97, 152–159. DOI:10.1016/j.isprsjprs.2014.08.017.
11. Yuan 2016
12. Peshkin, Eric. 2018. “The Quantification of Markers of economic development from time-series satellite imagery using deep learning.” A thesis submitted to the Duke Department of Mathematics.
13. Kaiser et al 2017. Learning Aerial Image Segmentation From Online Maps. IEEE Transactions on Geosciences and Remote Sensing. 55(11): /ieeexplore.ieee.org/abstract/document/7987710/