Harnessing Artificial Intelligence for the Analysis of Structural Damage After Extreme Weather and Seismic Events

Thomas Y. Chen
Academy for Mathematics, Science, and Engineering

Summary

Natural disasters ravage the world's cities, valleys, and shores on a monthly basis. Having precise and efficient mechanisms for assessing infrastructure damage is essential to channel resources and minimize the loss of life. Using a dataset that includes labeled pre- and post- disaster satellite imagery, we train multiple convolutional neural networks to assess building damage on a per-building basis. In order to investigate how to best classify building damage, we present a highly interpretable deep-learning methodology that seeks to explicitly convey the most useful information required to train an accurate classification model. We also delve into which loss functions best optimize these models. Our findings include that ordinal-cross entropy loss is the most optimal loss function to use and that including the type of disaster that caused the damage in combination with a pre- and post-disaster image best predicts the level of damage caused. Our research seeks to computationally contribute to aiding in this ongoing and growing humanitarian crisis, heightened by climate change.

Theory / Method / Workflow

For this work, we utilize the xBD dataset, which covers a wide range of disasters in fifteen countries around the world, from Guatemala to Portugal to Indonesia (over 850,736 building polygons totaling an area of 45,361 square kilometers). One of xBD's main purposes is to demonstrate changes between pre-disaster and post-disaster satellite imagery to aid in detecting the damage caused. Therefore, each post-disaster building is labeled as one of the following: "unclassified," "no damage," "minor damage," "major damage," or "destroyed." We train a baseline classification model to classify buildings by damage level, as defined by the Joint Damage Scale. The model input is only the post-disaster image. Notably, our baseline model does not use change detection. Because the data is labeled, this is a supervised approach. The model architecture is ResNet18, an 18 layer CNN, and was pre-trained on ImageNet data. This baseline model uses the cross-entropy loss function. The network is trained on 12,800 buildings crops with a batch size of 32. The Adam optimizer with a learning rate of 0.001 is used. The model trained for 100 epochs on NVIDIA Tesla K80 GPUs.

We train other models that improve upon the performance of the baseline model. To do this, we introduce other model inputs, namely the pre-disaster image (in combination with the post-disaster image) and the type of disaster (e.g. volcano, wind, etc.) that caused the building damage. To train a model that takes in both pre-disaster images and their corresponding post-disaster images, we concatenate the RGB channels of the two and use that as input. To train a model that takes in the pre-disaster image, post-disaster image, and disaster type, we do the same, but also concatenate a one-hot encoded representation of the disaster type in one of the later layers of the CNN. Furthermore, we experiment with other loss functions, namely mean squared error loss and ordinal cross-entropy loss to train these models. The other aspects of the training process (optimizer, learning rate, number of epochs, etc.) remain the same. These improved models contribute to our understanding of what information leads to the most accurate prediction results for building damage assessment.

Results, Observations, Conclusions

Table 1: Comparison of Validation Accuracy on Nine Different Models

Model Accuracy on Validation Set with Chosen Loss (100 epochs)			
	Loss Function		
Model Input	Mean Squared Error	Cross-Entropy Loss	Ordinal Cross-Entropy Loss
Post-Disaster Image Only	45.3%	59.5%	64.2%
Pre-Disaster, Post-Disaster Images	50.2%	68.3%	71.2%
Pre-Disaster, Post-Disaster Images, Disaster Type	49.7%	72.7%	74.6%

In Table 1, we present model accuracy on the validation set across nine different models, which are differentiated by three different input combinations and three different loss functions. The baseline model, which is trained with post-disaster data only and the cross-entropy loss function, has an accuracy of 59.5%, as shown. It is important to note that all models were trained and validated on data that is evenly split between building crops of each class (no damage, minor damage, major damage, and destroyed), so a purely blind guessing model would achieve approximately 25% accuracy. However, we note that none of the accuracy numbers are necessarily optimal. This can be explained by the fact that the differences between categories, particularly between minor-damage and major-damage, are often difficult to discern, for both humans and AI. This is certainly a challenge with non-binary classification tasks with building damage that has been acknowledged by many, including Gupta et. al. Additionally, more thorough dataset cleaning may yield marginally more accurate results. These results contribute to the research area of building damage detection by addressing the limited interpretability of current literature in regards to what types of information are most useful to building damage classification models as well as which loss functions are the best criterion. The main insights that can be drawn from our work include using individualized building crops instead of semantic segmentation to train models and performing experiments with various combinations of model inputs and loss functions to explicitly examine their differences. Our work's main contribution to the field is presenting a novel, more interpretable, analysis of how to classify building damage most accurately and effectively in the event of a natural disaster. Practically, our work and others in the field advance methods for more robust emergency responses and more efficient allocation of resources, which saves lives and property. This research is especially important now, when climate change is ramping up the frequency and intensity of these devastating events.

References

Lionel Gueguen and Raffay Hamid. Large-scale damage detection using satellite imagery. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1321–1328, 2015.

Ritwik Gupta, Richard Hosfelt, Sandra Sajeev, Nirav Patel, Bryce Goodman, Jigar Doshi, EricHeim, Howie Choset, and Matthew Gaston. xbd: A dataset for assessing building damage from satellite imagery. arXiv preprint arXiv:1911.09296, 2019.

Hanxiang Hao, Sriram Baireddy, Emily R Bartusiak, Latisha Konz, Kevin LaTourette, MichaelGribbons, Moses Chan, Mary L Comer, and Edward J Delp. An attention-based system fordamage assessment using satellite imagery.arXiv preprint arXiv:2004.06643, 2020.