

Mapping Electricity Infrastructure with Deep Learning



BASS CONNECTIONS

Ben Alexander^{1,2}, Wendell Cathcart³, Atsushi Hu⁴, Varun Nair^{2,4}, Lin Zuo^{1,2}

Faculty Advisors: Kyle Bradbury, Leslie Collins

¹Department of Statistical Science, ²Department of Computer Science, ³Nicholas School of the Environment, ⁴Department of Mathematics

Bass Connections in Energy & Environment

Introduction

Over 14% of the global population lacks access to electricity. This lack of access is largely concentrated in the developing nations of Sub-Saharan Africa and Southeast Asia [1] (Figure 1). Since electricity access positively correlates with improved economic, educational, and health outcomes [2], identifying cost-effective pathways to electrification is important. However, this process requires data on existing grid infrastructure, which is often unavailable or expensive.

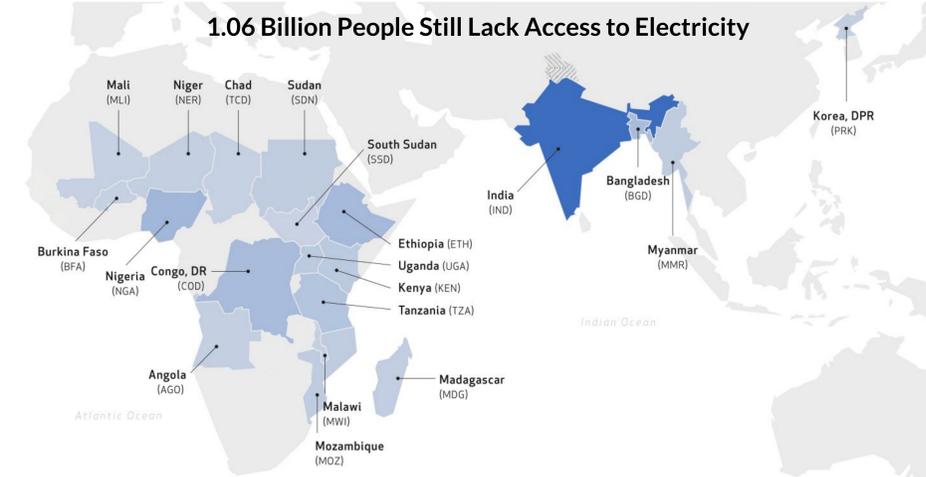


Figure 1: Countries with the greatest need for electricity access span two continents and represent a wide range of geographies [3]. While India has the single largest unelectrified population, countries in Sub-Saharan Africa tend to have larger unelectrified populations as a percent of their total populations.

Our goal is to fill this information gap in the electricity access space by generating maps of electric power grid networks that enable policymakers to identify optimal strategies for providing electricity access to communities without such access. To accomplish this, we use deep learning techniques to detect electricity transmission and distribution infrastructure in satellite and aerial images. Past research on aerial imagery has focused on extraction and classification of land use [4][5], identification of solar PV arrays [6], building detection [7], or segmentation of road networks [8], but the detection of electricity infrastructure is relatively unexplored.

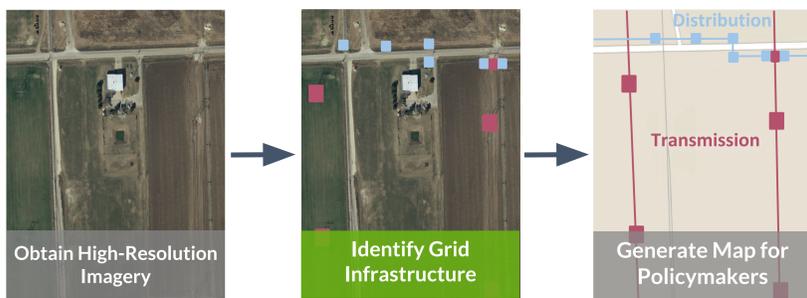


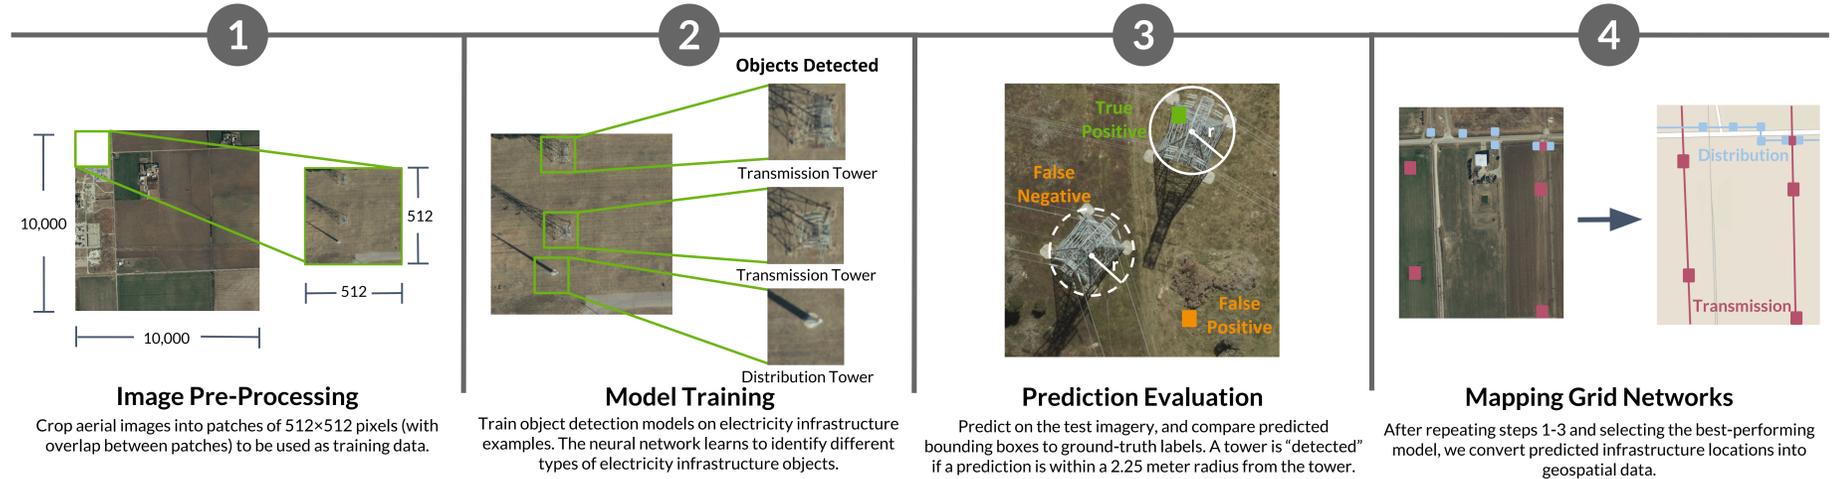
Figure 2: Project Overview. Overview of the automated process of mapping electric power grid networks by detecting electricity infrastructure in high-resolution satellite and aerial imagery with deep learning.

This project focuses on three main research questions:

- Q1** How accurately can we detect transmission and distribution infrastructure in satellite and aerial imagery using state-of-the-art object detection frameworks?
- Q2** How generalizable is the performance of electricity infrastructure object detection models across different geographic locations?
- Q3** How does the resolution of aerial or satellite imagery impact the electricity infrastructure detection model's performance?

We answered the first question (Q1) by applying three state-of-the-art object detection frameworks, Faster R-CNN [9], RetinaNet [10], and YOLOv2 [11], to an annotated dataset containing labeled electricity infrastructure in aerial imagery from regions in Arizona, Connecticut, North Carolina, and Kansas [12]. We then chose the best performing model to help answer the second (Q2) and third (Q3) questions.

Methods



Q1 Accuracy of Grid Infrastructure Detection

Using the object detection architectures RetinaNet, Faster R-CNN, and YOLOv2, we trained three models on aerial imagery from Arizona, Connecticut, Kansas, and North Carolina. Each model's parameters were fine-tuned before comparing results. RetinaNet had the best overall performance, with a precision and recall both around 0.65 (at the point with the best F1 score). RetinaNet also achieved a maximum recall of 0.72 (at a precision of 0.52). Together, these results indicate that the model detected the majority of transmission and distribution towers, while not producing an excessive number of false positives.

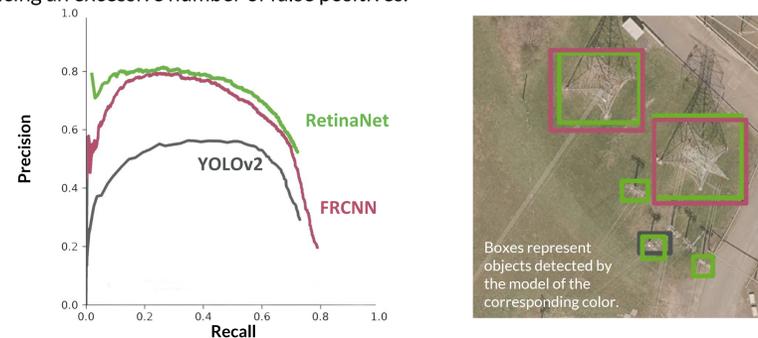


Figure 3: (Left) Comparison of the three object detection models' performance on the same test set. The recall shows that the models (colored above) detect the majority of transmission and distribution towers, and the precision shows that the models don't produce excessively many false positives. (Right) The three models made different object detections in this sample image.

Q2 Generalizability Across Diverse Geographies

To scale up these techniques, we need to be able to train once, and apply in many diverse geographies. To see how well a model trained on imagery from one location can predict on other locations, we evaluated five models that were trained on five different training locations (as labeled below), on each test location. Unsurprisingly, models trained and tested on the same location had higher precision/recall. However, the best performance is from the model trained on the entire USA dataset, illustrating that training on diverse geographies leads to more generalizable models.

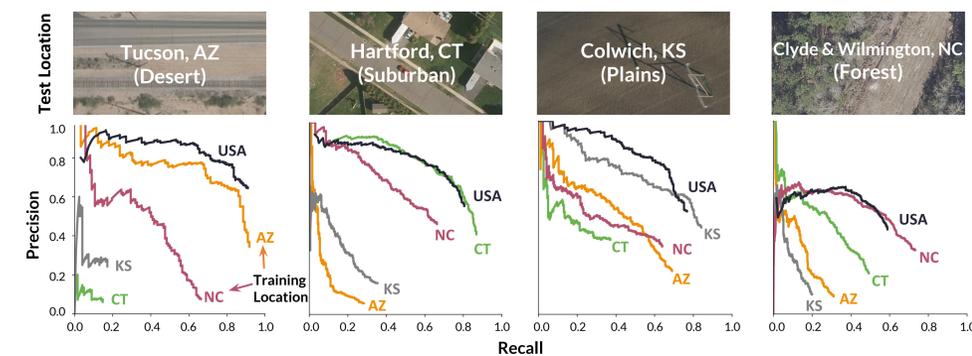


Figure 4: Performance of five models trained on five different training locations, evaluated on test imagery from four different states. Note that some training locations (labeled on the plots) generalize to different test locations better than others, and the model trained on data from across the USA (a combination of the other four locations) achieved the best performance.

Q3 Performance at Different Resolutions

Overhead imagery from different sources exhibit a wide range of resolutions (high: 0.15m to low: >10m) that impact how well infrastructure of different sizes can be seen in imagery. Higher resolution data is generally limited in availability and greater in cost. To explore the relationship between image resolution and model performance, we trained four models on images from Tucson, AZ at different resolutions to represent commonly available image sources. We found that model performance decreases significantly with imagery at 0.5m resolution or coarser.

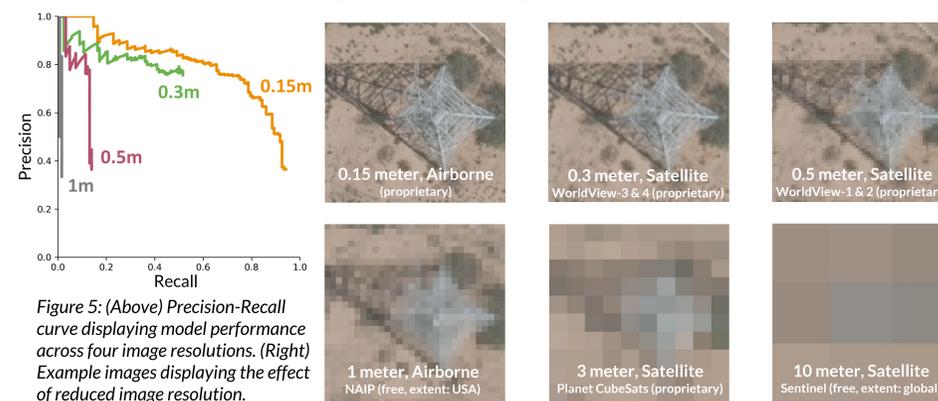


Figure 5: (Above) Precision-Recall curve displaying model performance across four image resolutions. (Right) Example images displaying the effect of reduced image resolution.

Conclusions

- Q1** After training three model architectures to identify grid infrastructure in imagery from four states across the USA, we achieved precision and recall performance both around 0.65, and with the potential to increase recall closer to 0.72 (with lower precision). These results show the promise of using deep learning to identify transmission and distribution towers for grid infrastructure mapping.
- Q2** Our model trained on all four US locations consistently matched or outperformed models only trained and tested on a single location (e.g., USA outperformed KS in column 3). This indicates that the model is able to aggregate knowledge about electricity infrastructure in imagery from multiple locations and apply it to improve performance. This provides a potential solution for making more generalizable models, and supports the prospect of training an even broader "world model" with additional data.
- Q3** Our findings suggest that electricity infrastructure detection works with very high resolution imagery (0.15m to 0.3m resolution), while overhead imagery from most publicly-available sources (of 1m resolution or coarser) is too low-resolution. Model performance declined as image resolution decreased, with resolutions of 0.5m or coarser performing poorly.

References

- "State of Electricity Access Report (SEAR)", World Bank, May, 2017. [Online]. Available: <https://www.worldbank.org/en/topic/energy/publication/sear>. [Accessed: 03-Dec-2018].
- Alstone, P., Gerbenson, D., & Kamen, D. M. (2015). Decentralized energy systems for clean electricity access. *Nature Climate Change*, 5(4), 305.
- "About the Heat Maps: Progress and Challenges on Sustainable Energy | Sustainable Energy for All (SE4ALL)", 2014. [Online]. Available: <https://www.seforall.org/heatmaps>. [Accessed: 03-Dec-2018].
- Yao, X., Han, J., Cheng, G., Qian, X., & Gao, L. (2016). Semantic Annotation of High-Resolution Satellite Images via Weakly Supervised Learning. *IEEE Transactions on Geoscience and Remote Sensing*, 54(6), 3660-3671.
- Scott, G., England, M., Starns, W., Marcum, R. (2017). Training Deep Convolutional Neural Networks for Land-Cover Classification of High-Resolution Imagery. *IEEE Geoscience and Remote Sensing Letters*, 14(4), 549-552.
- J. M. Malof, L. M. Collins and K. Bradbury A deep convolutional neural network, with pre-training, for solar photovoltaic array detection in aerial imagery, 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), pp. 91-99.
- A. Zakharov, A. Tuzhilkin and A. Zhiznyakov, "Automatic building detection from satellite images using spectral graph theory", 2015 International Conference on Mechanical Engineering, Automation and Control Systems (MEACS), Tomsk, 2015, pp. 1-5.
- Panboonyuen, T., Vateekul, P. (2017). Road Segmentation on Remotely-Sensed Images Using Deep Convolutional Neural Networks with Landscape Metrics and Conditional Random Fields. *Remote Sens.* 9(680), 10.20944.
- Rim, S., He, K., Girshick, R., & Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems* (pp. 91-99).
- Lin, T. Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision* (pp. 2980-2988).
- Redman, J., & Farhadi, A. (2017). YOLOv2: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7263-7271).
- K. Bradbury, G. Han, V. Nair, T. Pathirathna, and K. You, "Electric Transmission and Distribution Infrastructure Imagery Dataset", Aug. 2018. [Online serial]. Available: https://figshare.com/articles/Electric_Transmission_and_Distribution_Infrastructure_Imagery_Dataset/6933088. [Accessed: 03-Dec-2018].