Creating Artificial Worlds with AI to Improve Access to Energy Data

Caroline Tang
Alex Kumar
Frankie Willard

Introduction

Worldwide, about 1.2 billion people still don't have access to electricity in their homes. This map from the World Bank in 2017 highlights the places that are still greatly lacking in electricity. Access to electricity is becoming increasingly critical, especially for promoting economic development, social equity, and improving quality of life. For example, it has been shown that electricity access is correlated with improvements in income, education, maternal mortality, and gender equality. (1)

One of the first steps in improving energy access is acquiring comprehensive data on the existing energy infrastructure in an area. This includes information on the type, quality, and location.

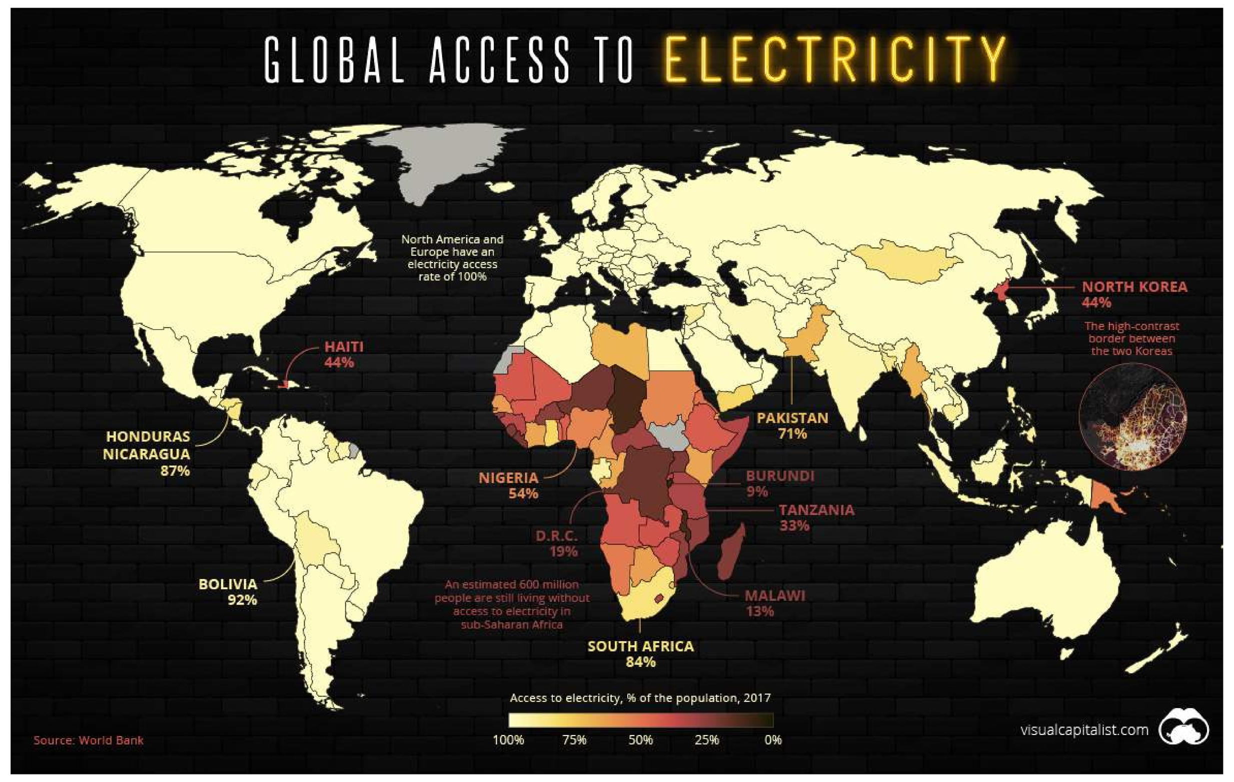

Global heat map showing access to electricity (2)

This information is key for helping make decisions about where to prioritize development and more. However, this critical information for expanding energy accessibility is often unavailable or low quality.

Methods & Data

A solution to efficiently filling this data gap is to automate the process of mapping energy infreastructure in satellite imagery. Using a deep learning object detection model, we can input satellite imagery to the model and make predictions about the characteristics and contents of the energy structures in the image.

However, properly training the object detection model requires signficant amounts of already labelled images and in situations where the available training data is from a different geographical region than the region to which the model is applied, the variations in physical characteristics negatively affect the model's accuracy as the model's ability to generalize across these different geographies is poor.

Our proposed solution to this is to generate synthetic images that supplements the original satellite imagery training dataset. We generate these synthetic images by cropping the existing energy infrastructure out of satellite images and placing them on top of a different real image (serves as a background) that originally features no energy infrastructure from one of the target geographic domains. By doing so, we create a larger dataset to train on. To blend the cropped energy infrastructure and the background images together, we use an image harmonization technique, specifically GP-GAN (3), to blend the images together. GP-GAN is a Gaussian-Poisson Generative Adversarial Network that works to blend high-resolution images.

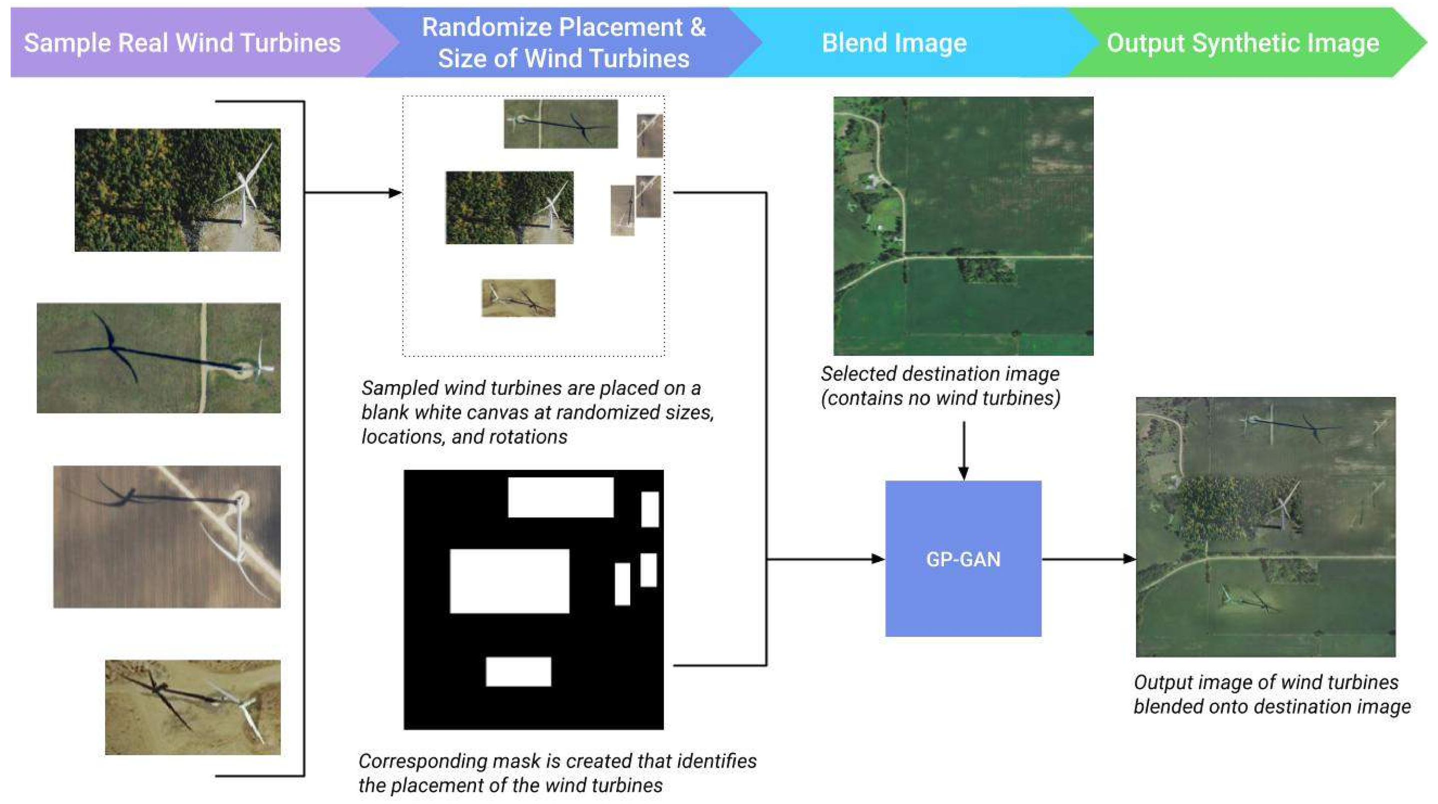

Illustrates process of generating synthetic imagery from real imagery

We created our augmented synthetic dataset by generating 75 images blended by GP-GAN and mixing it with 100 real images. We selected this ratio of real to synthetic imagery based on previous experimentation that aimed to find the ratio of images that optimized performance. (4)

Then, we test our augmented synthetic dataset against the original baseline dataset and against our previous Bass Connections Team's augmented synthetic dataset - which was created by placing cropped 3D models of wind turbines from CityEngine over background imagery. CityEngine is a 3D modeling software that contains assets of various infrastructure, including wind turbines!

Experimental Setup & Results

To evaluate the potential of synthetic imagery in improving the performance of object detection, we set up Within-Domain and Cross-Domain experiments, where a domain is defined as a specific geographic region. The source domain refers to the region that the real training data comes from, while the target domain refers to the region that the object detection model is applied to. These two types of experiments each correspond to a potential real-world situation one might encounter, and help us to evaluate the potential performance of the object detection model in each of these situations with limited real world data.

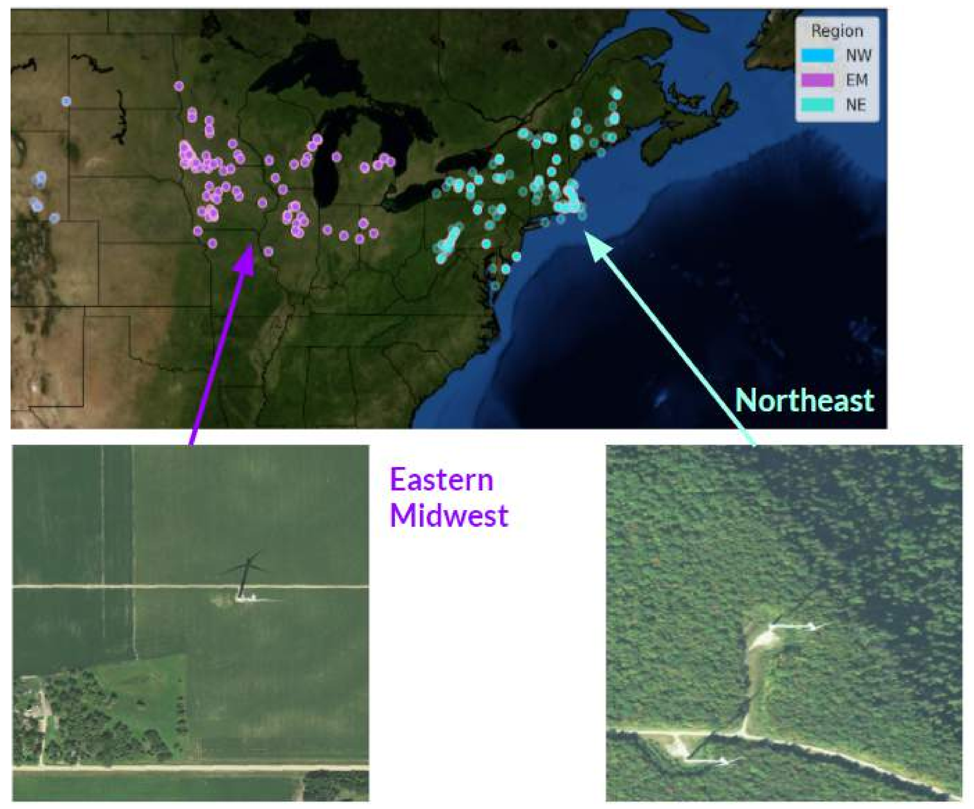

Map displaying the experimental target domains (5)

Within-Domain Experiment: For each of the domains we selected, we ran the baseline and modified experiments, where all of the data came from the same region. This experiment helps to evaluate the overall ability of synthetic imagery (especially using our GP-GAN technique) to improve the object detection performance.

Cross-Domain Experiment: For these experiments, the domains for the real source and target images are different, while the synthetic images used in the modified training dataset are from the target region.

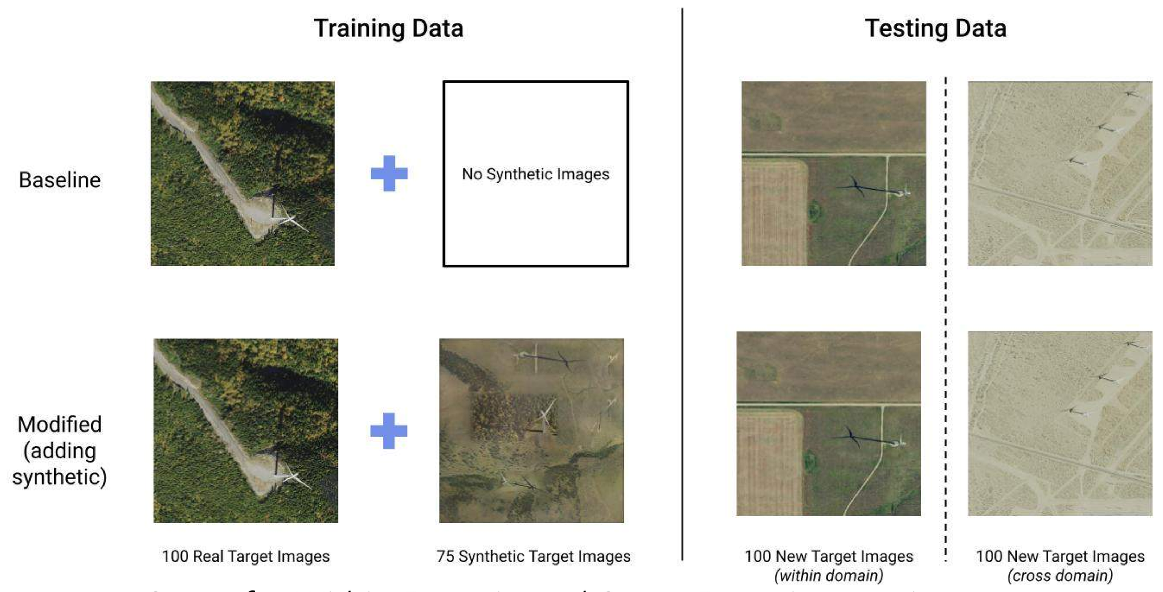

Setup for Within-Domain and Cross-Domain experiments

Thus, with synthetic images more similar to the target region, we hypothesize that the addition of the synthetic images will improve the accuracy of the object detection when the target and source regions are dissimilar in appearance. This means an improvement the model's ability to generalize across different regions despite the limitations of the existing training data.

Overall	Within-Domain	Cross-Domain	Average
Average: Baseline	0.817	0.477	0.647
Average: Added CityEngine	0.881	0.593	0.737
Average: Added GP-GAN	0.896	0.628	0.762

Averaged results for Within-Domain and Cross-Domain experiments

After conducting the Within-Domain and Cross-Domain experiments on the Eastern Midwest (EM) and North Eastern (NE) regions, we found that our augmented synthetic dataset improved Average Precision (AP) by 9.7% Within-Domain and 31.7% Cross-Domain when compared to the baseline performance. When compared to the CityEngine synthetic dataset, Within-Domain AP performance increased by 1.7% and Cross-Domain increased by 5.9%.

Average Precision is the area under a Precision-Recall curve, and is a good general metric for object detection accuracy.

Experiments	Validate EM	Validate NE	Average
Train EM: Baseline	0.822	0.387	0.605
Train EM: Added CityEngine	0.919	0.487	0.703
Train EM: Added GP-GAN	0.900	0.638	0.769
Train NE: Baseline	0.567	0.812	0.690
Train NE: Added CityEngine	0.698	0.842	0.770
Train NE: Added GP-GAN	0.619	0.892	0.755

Full experimental results

Conclusions & Next Steps

From these results, we can conclude that adding synthetic imagery that is blended by GP-GAN to a smaller baseline dataset of satellite imagery yields promising results for boosting Cross-Domain object detection performance. In our overall goal of object detection on energy structures in satellite imagery, it is vital that our model is able to perform well across multiple domains, even when less labelled real data is available for that domain.

For the future, we would like to expand our methodology to creating synthetic datasets with all types of energy infrastructure to test model performance in multi-class classification and move closer to the ultimate goal of mapping energy infrastructure in satellite imagery!

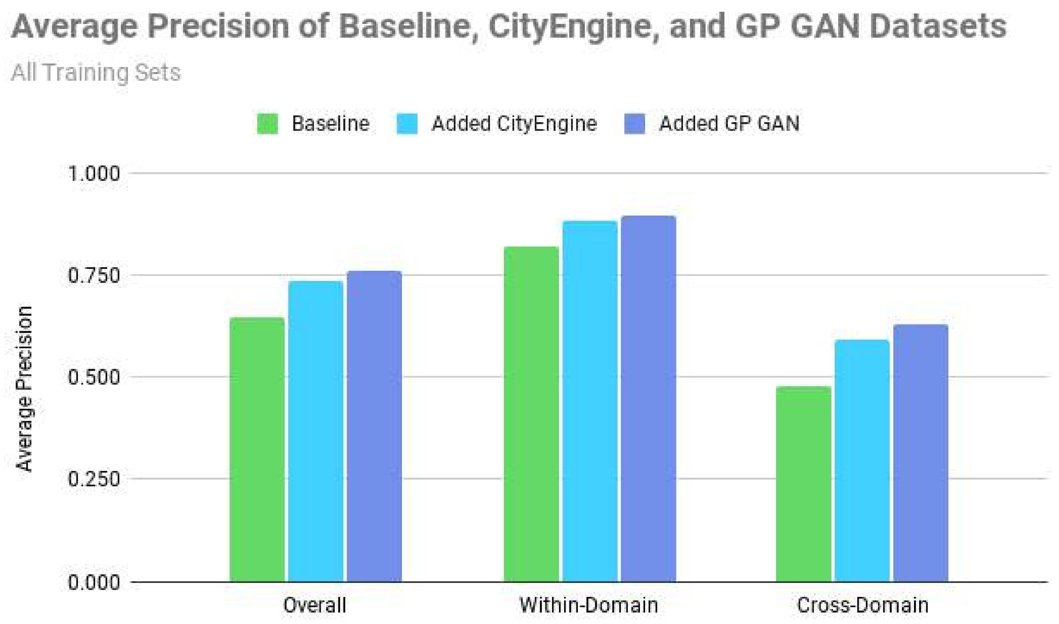

Plotted Average Precision for comparison

Furthermore, we would like to test our methodology on more geographic domains. Additional Cross-Domain improvement in performance would strengthen our results. We would also like to attempt more hyperparameter tuning in the image augmentation process (GP-GAN or other) to ensure higher realisticness of our generated images and increase blending.

References & Acknowledgments

(1) Hanania, J., Stenhouse, K., & Donev, J. (2018, June 25). Access to Electricity. Energy Education. https://energyeducation.ca/encyclopedia/Access_to_electricity.
(2) Routley, N. (2019, November 27). Mapped: The 1.2 Billion People Without Access to Electricity. Visual Capitalist. https://www.visualcapitalist.com/mapped-billion-people-without-access-to-electricity/
(3) Wu, H., Zheng, S., Zhang, J., & Huang, K. (2019, August 5). Gp-gan: Towards realistic high-resolution image blending. arXiv.org. https://arxiv.org/abs/1703.07195.
(4) Hu, W., Feldman, T., Ou, Y. J., Tarn, N., Ye, B., Xu, Y., Malof, J. M., & Bradbury, K. (2021). Wind Turbine Detection with Synthetic Overhead Imagery. International Geoscience and Remote Sensing Symposium. Published.
(5) Ye, A., Lin, E., Moscoso, J. L., Feldman, T., Ou, J., & Zhang, W. (2021). Locating Energy Infrastructure with Deep Learning. Deep Learning for Rare Energy Infrastructure in Satellite Imagery. https://duke-bc-dl-for-energy-infrastructure.github.io/

Thank you to: Wayne Hu, Dr. Kyle Bradbury, Dr. Jordan Malof, and the previous Bass Connections team!