Title: Predicting Levels of Household Electricity Consumption in Low-Access Settings

Stakeholder: Kenya Power and Lighting Company(KPLC), World Resources Institute(WRI), Nithio Energy, TFE Energy, VIDA

Theory of change:

In low-income settings, accurately predicting a customer's anticipated electricity consumption is crucial for electricity providers. This task is particularly challenging when many households lack electricity connections, leading to a wide variability in expected consumption. Connecting low consumers instead of higher consumers can result in the inefficient use of resources. The consumption estimations from this work were made available through an API that has been adopted by electrification planners to design appropriate electricity infrastructure for unelectrified communities.

Technical Abstract:

This study is the first of its kind to predict electricity consumption at the building level, rather than for an entire administrative area, in low-income settings. A Convolutional Neural Network (CNN) was trained using pre-electrification daytime satellite imagery and utility bills from 20,000 geo-referenced electricity customers in Kenya. This innovative two-stage approach leverages no-cost satellite imagery to maximize the use of limited customer data, achieving competitive accuracies in predicting building-level consumption. The study highlights the importance of both the building's characteristics and its surrounding context. Additionally, incorporating lower resolution geospatial datasets, such as nighttime lights and census data, into the training process enhances predictions. The results are already aiding in site selection and distribution planning in Kenya, with potential applicability to other countries.

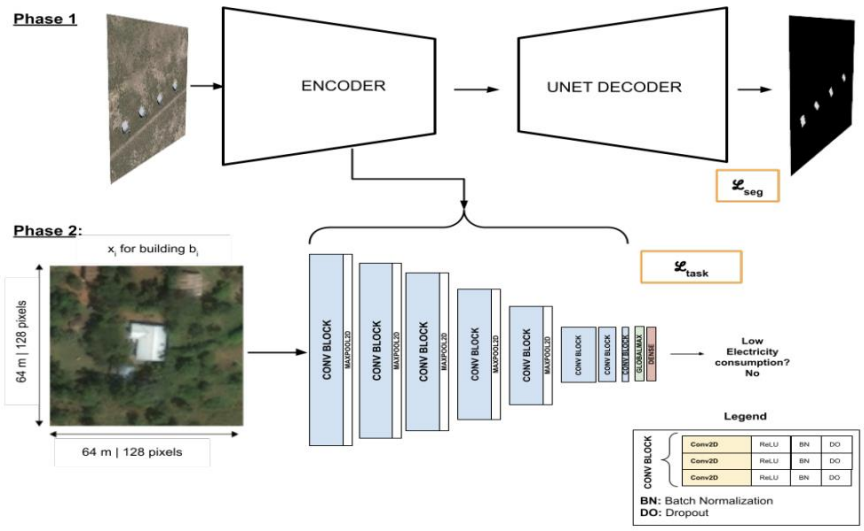

In the figure above, Phase 1 involves training a building segmentation model using an encoder from Phase 2 and a UNET decoder. In Phase 2, the pretrained encoder from Phase 1 is repurposed to learn the electricity prediction task. An image (x_i) of a household's building is fed into the pretrained encoder, which is then trained with a negative log-likelihood loss function to predict electricity consumption levels upon electrification.

Link to paper