

PGRID: Power Grid Reconstruction in Informal Developments Using High-Resolution Aerial Imagery

Simone Fobi Nsutezo^{1*}, Amrita Gupta¹, Duncan Kebut², Seema Iyer³,
 Luana Marotti¹, Rahul Dodhia¹, Juan M. Lavista Ferres¹, Anthony Ortiz¹

¹Microsoft AI for Good Research Lab, ² HOTOSM,
³USA for UNHCR

Abstract

As of 2023, a record 117 million people have been displaced worldwide, more than double the number from a decade ago [22]. Of these, 32 million are refugees under the UNHCR's mandate, with 8.7 million residing in refugee camps. A critical issue faced by these populations is the lack of access to electricity, with 80% of the 8.7 million refugees and displaced persons in camps globally relying on traditional biomass for cooking and lacking reliable power for essential tasks such as cooking and charging phones. Often, the burden of collecting firewood falls on women and children, who frequently travel up to 20 kilometers into dangerous areas, increasing their vulnerability. [7]

Electricity access could significantly alleviate these challenges, but a major obstacle is the lack of accurate power grid infrastructure maps, particularly in resource-constrained environments like refugee camps, needed for energy access planning. Existing power grid maps are often outdated, incomplete, or dependent on costly, complex technologies, limiting their practicality. To address this issue, PGRID is a novel application-based approach, which utilizes high-resolution aerial imagery to detect electrical poles and segment electrical lines, creating precise power grid maps. PGRID was tested in the Turkana region of Kenya, specifically the Kakuma and Kalobeyi Camps, covering 84 km² and housing over 200,000 residents.

Our findings show that PGRID delivers high-fidelity power grid maps especially in unplanned settlements, with F1-scores of 0.71 and 0.82 for pole detection and line segmentation, respectively. This study highlights a practical application for leveraging open data and limited labels to improve power grid mapping in unplanned settlements, where the growing number of displaced persons urgently need sustainable energy infrastructure solutions.

*Corresponding author: sfobinsutezo@microsoft.com

1. Introduction

Reliable electricity access in refugee camps is critical to supporting everyday tasks such as cooking, phone charging and powering small businesses. Electricity consumption and income/earning potential are tightly correlated over time [20], where increased electricity consumption is in lockstep with increased income. Thus electricity access in camps may better support residents' ability to earn a living.

Many refugee camps are often found in remote areas with little to no access to electricity. As a result, about 8.7 million people in refugee camps do not have access to much needed electricity [7], to improve their livelihoods. Thus mapping the intricate layout of power grids is important for quantifying electricity access within camps and for strategic planning to support further electrification of camps.

High resolution information on power grid layout/structure is scarce, partial and often outdated. Advanced sensing devices, such as Advanced Metering Infrastructure (AMIs) and Phasor Measurement Units (PMUs), offer a pathway to grid mapping [16]. These technologies are capable of delivering real-time data on voltage measurements, thereby supporting voltage comparisons across meters to recreate the distribution grid. However, deploying smart meters in refugee camps across the distribution network may come with financial and technical challenges. The cost of these devices across the network and the specialized knowledge required to interpret the data they provide, particularly in terms of power-flow physics to map power grid layouts, makes them infeasible for unplanned settlements.

Mapping power grids from remotely sensed imagery.

As an alternative to mapping power grids from smart meter devices, satellite and aerial imagery has been utilized. Early work by Development Seed [10] demonstrated a human-in-the-loop approach to mapping high voltage (HV) infrastructure. This approach used high resolution 50 cm/pixel optical satellite imagery to first detect HV towers, and then a team of experienced mappers traced HV lines between tow-

ers. Alternatives to this workflow leverage lower resolution Synthetic Aperture Radar (SAR-C) imagery to detect HV steel pylons [9]. This work indicated that grid infrastructure can be identified in satellite imagery, given the right image resolution and power infrastructure material visible in the image. While HV lines provide information about the transmission network, they do not provide a view of the distribution network which is more reflective of where electricity is consumed at the residential and commercial levels. Moreover, the distribution network captures the various electric connection points and the ease and cost of making a power connection [12]. [11] show that identifying the distribution network from satellite imagery remains challenging as distribution networks typically consist of smaller poles and lines, which can be hard to detect even at 50cm/pixel. As an alternative, Visible Infrared Imaging Radiometer Suite (VIIRS) nighttime lights is used in combination with network algorithms (e.g. Dijkstra’s shortest path) to estimate the distribution network given electrified settlements. While this approach approximates the grid, it does not provide a true representation of the network and is sensitive to noise in VIIRS data such as flares, background lighting and temporal changes in luminosity [3]. GridTracer is an approach to mapping power grids from overhead imagery using deep learning [13]. Alongside GridTracer, a benchmark dataset of overhead imagery, towers and lines was released to evaluate power grid mapping algorithms. Our work is most similar to GridTracer in its usage of overhead imagery, thus it serves as a baseline for comparison. Streetview imagery has also been used as an alternative to overhead imagery to map grid layouts [18]. While this provides an interesting application of streetview imagery, this approach is challenging to implement in unplanned settlements as streetview imagery is often not available for these locations.

More recently, higher spatial resolution data collection through Unmanned Aerial Vehicles (UAV) such as drones has become more accessible with open source platforms such as OpenAerialImagery. As a result, drone imagery is increasingly utilized to map grid infrastructure. Using the Transmission Towers and Power Lines (TTPLA) dataset [1], [2] propose a transformer based segmentation model to map individual transmission lines and towers. While these works advance the field of grid mapping from drone imagery, the data used in algorithm development consists of high voltage transmission lines at very high resolution, with profile views as opposed to overhead spatial views (See Figure 1). Datasets in profile view, are better suited for mapping individual lines and tower structures as opposed to reconstructing the entire grid structure from the dataset.

Application context. According to the UN High Commission on Refugees (UNHCR), by mid-2023 there were approximately 32 million refugees under UNHCR mandate [22], with about 22% (likely a lower bound) of the refugee



Figure 1. Profile views (left) better support mapping individual lines and tower structures, while overhead imagery (right) better supports mapping the layout of a power grid.

population living in refugee camps [23]. This resulted in the creation of several informal settlements to provide safe shelter and amenities for displaced populations. An important aspect to supporting resettlement and enhancing quality of life for communities within these settlements is access to reliable and sufficient power. Accurate and updated maps of the power distribution infrastructure can inform relief organizations and field operators about power infrastructure needs and better support planning, resource allocation and operations within the camp. Given the nature of informal settlements, power distribution grids more likely appear overhead and are less organized, as the power grid is established to align with housing locations, which may be more unplanned compared to planned settlements. In this paper, in collaboration with USA for UNHCR and HOTOSM, we deploy our approach to mapping overhead power distribution infrastructure in the Kakuma Camp and Kalobeyei Integrated Settlement covering 84 km², and home to over 200,000 residents [21]. PGRID is evaluated using manually tagged labels from well-trained mappers. We also create a geo-referenced power grid layout map and show that our approach can be used to augment existing power grid maps.

Contributions. This work presents an approach to mapping power grid distribution infrastructure from overhead aerial imagery at a spatial resolution of $\sim 6\text{cm/pixel}$. The method is applied in practice to mapping refugee camps thereby enabling humanitarian mappers to rapidly map grid infrastructure in ever growing camps. In this work, we pro-

pose pole detection and line segmentation approaches for mapping low and medium voltage distribution lines. The main contributions of this paper are:

1. Easy-to-deploy algorithms for detecting electrical poles using point supervision and segmenting electrical lines. Our pole detection method simplifies the label acquisition process, as point labels are quicker and easier to obtain compared to bounding boxes.
2. Deployment of our PGRID system in an unplanned settlement where we show high model performance in reconstructing the layout of the existing power grid.
3. We demonstrate that our power distribution grid map enhances existing power grid estimates, and can serve as a viable approach to obtaining high resolution maps.

We share our github repository and demo here: <https://github.com/USAFORUNHCRhive/turkana-grid-mapping>.

2. Problem Formulation

For the task of mapping power distribution grids from overhead high resolution drone imagery, we formulate the problem as follows: for a given area a , the unknown power distribution grid network (\mathcal{G}_a) is defined as an undirected acyclic graph $\mathcal{G}_a(p, l)$, where p is a set of poles ($p = p_1, p_2, p_3, \dots, p_n$) and l are the line segments connecting the poles ($l = l_{12}, l_{13}, l_{34}, \dots, l_{mn} = (p_1, p_2), (p_1, p_3), (p_3, p_4), \dots, (p_m, p_n)$). Under perfect conditions, a power distribution grid can be fully mapped using accurate information about all line segments in the network. However, we observe from drone imagery that electrical lines appear in the image as very thin lines and their visibility within the image is dependent on the time when the imagery is taken, the lens angle and the building density within the image. Moreover, some areas may still be undergoing construction with existing poles but no lines. Because electrical lines may not always be visible, we leverage complementary information about electrical poles and detect both electrical lines and poles. Thus the objective of this paper is to estimate a power distribution grid ($\hat{\mathcal{G}}_a$) in a new area using geospatial machine learning to predict the location of electrical poles and lines that constitute the power distribution grid, given points and lines from a known area.

$$\hat{\mathcal{G}}_a = M(\mathcal{P}(X), \mathcal{L}(X)) \quad (1)$$

$\hat{\mathcal{G}}_a$ is obtained in a two-step process as follows: First, we train a pole detection model (\mathcal{P}_a), given input imagery (\mathcal{X}_a) to accurately detect electrical poles given point pole locations from a known area. Next, we train a line segmentation model (\mathcal{L}_a) also given input imagery (\mathcal{X}_a) to detect connecting lines. Detected poles and lines are merged after post-processing to recreate the power distribution grid $\hat{\mathcal{G}}_a$.

3. PGRID: Reconstruction of grid layouts.

The PGRID method entails three steps: i) a pole detection model to detect electrical pole locations within the image ii) a line segmentation model to segment electrical lines from imagery and iii) a post-processing step to merge detected poles and line segments to obtain a unified layout. We outline each step of our methodology below.

3.1. Electrical Pole Detection

Accurate identification of electrical poles is performed using our proposed pole detection model. We deploy a Fully Convolutional Network (FCN8) [15] semantic segmentation model to detect poles within the image. We build on the network architecture and multi-component loss function proposed by [14] to count and localize cars and adapt it to detect electrical poles. A labeled dataset of point locations is used to indicate electrical pole locations. This labelled dataset, coupled with corresponding images is used as a training signal to learn pole representations within the imagery (see supplementary material for illustration of pole detection network). The model is trained with a four component loss: i) *image-level loss*, is a negative log-likelihood loss which encourages the model at the image-level to detect electrical poles. For images that have poles, the loss increases the likelihood that at least one pixel is predicted as a pole, while decreasing the likelihood of predicting pixels as poles when there are no poles in the image. Here all image pixels are considered, ii) *point-level loss*, which serves as a localization loss, helping the model to accurately learn pole structures within the image. A negative log-likelihood loss is applied only to the annotated pole pixels to encourage the model to correctly identify poles. Both the image and point level losses evaluate the likelihood that a pixel is in the pole class. The next two losses, evaluate the likelihood that a pixel belongs to the background class. The iii) *split-level loss*, encourages the model to predict unique blobs for all annotated poles. First, boundaries between ground truth poles are found using a watershed segmentation algorithm [6]. These boundaries are annotated as background while the area within the boundaries are annotated as foreground. The model learns to predict the probability that a pixel belongs to the background class. Here we also use a negative log-likelihood loss, weighted by the number of poles within the blob. This encourages the model to predict boundaries such that only one pole is in each blob. iv) *false positive loss*, which discourages the model from making false positive predictions. A negative log-likelihood loss is applied to pixels with no ground truth pole annotations, thereby minimizing the false positive errors by the model.

A small learning rate of 1e-6 was used to train the pole detection model. Image augmentation techniques such as vertical and horizontal flips, color jittering and random rotations were used as a form of regularization during training.

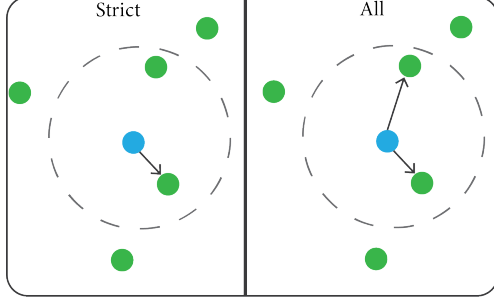


Figure 2. Illustration of *strict* (left) and *all* (right) matching variants for evaluating performance of pole predictions. Ground truth poles are shown at the center in blue, with predicted poles in green surrounding the ground truth pole. The dashed circle represents the threshold (th). In the *strict* match variant, the ground truth pole is matched to its closest prediction, so long as the prediction is within the threshold. In the *all* variant, the ground truth pole is matched to all predicted poles within the threshold. The *all* match is a better approach to evaluate performance for cases where mask predictions for poles is a non-continuous blob.

Pole Detection Metrics. Electrical poles are represented as points in the label dataset, and predictions obtained from the model are vectorized and their centroids taken as predicted pole locations. Matching ground truths with predictions cannot be performed using exact matches given the high precision nature of points. Thus, we set an acceptable distance thresholds th and utilize two matching variants known as *strict* and *all* matches, as illustrated in Figure 2. In the *strict* variant, we perform a one-to-one match between the ground truth pole and a predicted pole, given both poles are within the acceptable distance. When multiple predictions occur within the acceptable distance, we select the closest prediction to the ground truth pole. In the *all* variant, we perform a one-to-many match between the ground truth and predicted poles, given th . The pole matching variants are used to compute the following metrics:

Distance-based Precision. Distance-based precision represents the proportion of predicted poles that have a corresponding ground truth pole, given th . Distance-based precision is computed given the modified precision formula:

$$Precision_{th} = \frac{TP_{th}}{(TP_{th} + FP_{th})} \quad (2)$$

where TP_{th} represents true positives at th , and FP_{th} represents false positives. The *strict* match represents a one-to-one match between predicted poles and ground truth poles, where the closest predicted pole is matched to a ground truth pole, given th . Additional predicted poles within the designated threshold are considered false positives. In the *all* match, all predicted poles within the designated threshold th are matched to the ground truth pole.

Distance-based Recall. Distance-based recall represents the proportion of poles that are correctly predicted as poles,

given th . Thus the modified recall formula is:

$$Recall_{th} = \frac{TP_{th}}{(TP_{th} + FN_{th})} \quad (3)$$

where TP_{th} represents true positives at th , and FN_{th} represents false negatives. TP_{th} is defined as the number of ground truth poles that are correctly predicted as poles, given a distance threshold (th).

Distance-based F1-Score. Distance-based F1-score represents the harmonic mean between precision and recall, given specific matching distance threshold (th).

3.2. Electrical Line Segmentation

We propose patch-level segmentation using an asymmetric DeepLabV3 [8] architecture (see supplementary materials) to segment electrical lines. We apply a scaling factor (sf) to the ground truth label raster to obtain patch-wise labels, indicating line presence or absence. This approach is selected over pixel-level segmentation because of the very thin nature of electrical lines. As an alternative, the scaling factor allows the model to classify each patch as having an electrical line or not. sf of 1 is equivalent to pixel-level segmentation, while sf with a value the same as the image size approximates image classification. Through experimentation, we observe that an sf of 4 yields good detection and localization of line segments. At this sf the ground truth mask is reduced to 1/4 of its original size e.g. 512 x 512 groundtruth mask becomes a 128 x 128 groundtruth after scaling, and the model classifies every 4x4 image patch as containing a line or not. DeepLabV3 is selected due to its atrous convolutions which support rich contextual feature extraction, a desirable attribute for line segmentation.

The line segmentation model is trained as a binary model, with a cross-entropy loss and an Adam optimizer at a learning rate of 1e-5. Other loss functions, such as Dice and Focal loss, yielded comparable performance to cross-entropy so no further optimization was done. Similar to the pole detection model, the line segmentation model is trained with augmented data samples to prevent overfitting. We evaluate the line segmentation model at the pixel level (buffered by 2 meters on each side) and compare it to the ground truth labels. Buffering is done because the ground truth test labels weakly align with visible electrical lines in the imagery as opposed to training labels that are highly aligned with visible electrical lines in the image. As a result, the model is trained to predict visible electrical lines in the imagery rather than infer connections between electrical poles. We observe that ground truth labels in the test set have an offset of $\sim 2m$ from visible electrical lines within the imagery. In addition to reporting the mean intersection-over-union (mIOU), precision, recall and F1-score are also reported. For grid layout reconstruction, identifying grid presence/absence is more informative than having high fi-

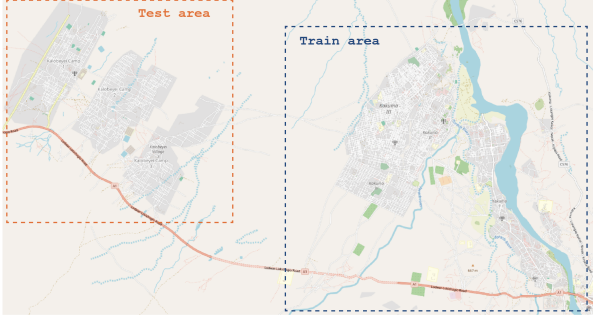


Figure 3. Spatially distinct geographic train and test splits for model training. Models were trained on the Kakuma Camp (right) and tested on the Kalobeyei Camp (left).

delity alignment between the predictions and ground truth, especially given weak ground truth line labels. Thus we measure model accuracy in addition to alignment.

3.3. Hard Negative Mining

Learning to detect electrical poles from imagery can be a difficult task as objects such as fences, street lights, billboard poles and even trees, would be closer in representation to electrical poles, compared to the background. Thus, these objects will be falsely classified as electrical poles thereby making reconstruction of the grid layout difficult. To address this issue, we implement a hard negative mining strategy for pole detection, where difficult negative samples such as fences are labelled and included in the dataset for training. By showing the model hard negative examples, the model learns distinct representations for electrical poles thereby further reducing false positive predictions.

3.4. Unified Power Grid Layout

After training both the pole detection and line segmentation models, the predicted pole raster files are vectorized to obtain polygons, and the polygon centroids are taken as pole locations. For the final grid map, we filter out small polygons prior to generating centroids. Predicted power lines are also vectorized, skeletonized and the same 2 m buffer is applied around the line skeletons. The buffered line predictions are then polygonized to obtain electrical line predictions. This provides a unified power grid ($\hat{\mathcal{G}}_a$) in an easily accessible and lightweight vectorized format for humanitarian mappers to use as visual guides when mapping grids.

4. Datasets

Mapping power grid infrastructure from imagery is contingent on having high resolution imagery with visible power infrastructure and corresponding ground truth annotations for model training. We utilize multiple datasets across diverse geographies to demonstrate the value of the contributions made in this paper.

Benchmark Dataset. We use a benchmark dataset of overhead satellite imagery and corresponding pole tower and line annotations [13] to evaluate the performance of our grid mapping algorithms. This dataset consists of overhead imagery collected from Arizona, USA (AZ); Kansas, USA (KS) and New Zealand (NZ). The images are provided at 0.3m, resampled from native resolutions ranging from 0.15m-0.3m. The first 20% of imagery from each city is used for testing, while the 80% is used for training and validation, consistent with the dataset authors’ format.

Turkana Integrated Settlement. In November 2022, the United Nations High Commissioner for Refugees (UNHCR), Humanitarian OpenStreetMap Team (HOT), and Kenya Red Cross Society (KRCS) partnered to conduct drone mapping of Kakuma-Kalobeyei. Drones successfully covered 84 km², capturing imagery at ~6 cm/pixel. The collected drone imagery was inspected by HOT mappers, who manually labeled electrical poles and lines for ~20% of the area. This paired dataset of manually annotated electrical poles and lines with drone imagery is used. Annotations for the Turkana Integrated Settlement camp were downloaded from OpenStreetMaps (OSM). The “*power*” and “*line*” tags were used to search OSM, where the following features were selected and downloaded: “*pole*”, “*line*” and “*minorline*”. To evaluate the generalizability of PGRID, the Turkana camps were geographically split as shown in Figure 3, where the model was trained with data from the Kakuma Camp (right) and tested in the Kalobeyei Camp (left). For the training data region, pole annotations from OSM were manually realigned to corresponding imagery prior to training the model. The line annotations from OSM were used as a guide to annotate actual lines in imagery as the lines obtained from OSM showed the general grid trajectory and not the actual electrical lines within the image. The models are trained and validated given the Kakuma train split and performance metrics are reported for the unseen Kalobeyei test split. 863 poles were used to train the pole detection model while 636 poles were used for evaluation. 171 line segments with a mean line length of 161 m were used to train the line segmentation model while 118 line segments with a mean line length of 240 m were used for evaluation.

Table 1. Distance based mean Average Precision (DmAP) for the tower detection task on the benchmark dataset. Comparison of GridTracer method with ours (PGRID). On average, PGRID tower detection outperforms GridTracer on the benchmark dataset.

Methods	Backbone	AZ	KS	NZ	Mean
GridTracer	ResNet50	0.45	0.48	0.54	0.49
PGRID	ResNet50	0.57	0.75	0.55	0.62
GridTracer	ResNet101	0.73	0.53	0.59	0.62
PGRID	ResNet101	0.70	0.72	0.55	0.66

Table 2. Mean Intersection-Over-Union (mIOU) for electrical line segmentation on the benchmark dataset.

Methods	AZ	KS	NZ	Mean
GridTracer-UNet	0.50	0.34	0.38	0.41
GridTracer-StackNetMTL	0.54	0.40	0.47	0.46
PGRID	0.51	0.34	0.37	0.41

5. Experiments & Results

5.1. Performance on benchmark dataset

We evaluate the performance of our PGRID models on the benchmark dataset. Tower detection and line segmentation models are evaluated on the test data from all three regions. The data split is consistent with GridTracer’s data split for evaluation. Table 1 shows Distance-based mean Average Precision (DmAP) on the tower detection task for the benchmark dataset. We compare results from PGRID to the GridTracer algorithm. PGRID pole detection method using point labels outperforms the FasterRCNN object detection method utilized by the GridTracer algorithm. Using a ResNet50 architecture, PGRID outperforms the GridTracer tower detection method on all three locations. When the more expressive ResNet101 model is used, further performance gains are observed with PGRID outperforming by 19% in Kansas and under-performing by only 3-4 % in Arizona and New Zealand. In addition, PGRID utilizes point labels, as opposed to bounding boxes, thereby making the label acquisition easier. [5] show that acquiring point annotations is faster and requires less effort than acquiring bounding boxes, thereby making point supervision more practical for real-world deployment. Line segmentation model performance on the benchmark dataset is presented in Table 2. PGRID performs comparably to the GridTracer UNet model without the added complexity of predicting orientation with StackNetMTL [4].

5.2. Performance in Turkana Integrated Settlement

PGRID pole detection and line segmentation models are run on the Turkana drone imagery and results are reported.

Pole detection performance. Figure 4 shows detected poles using the PGRID pole detection method for a small sample within the test set. Predicted poles are shown as red blobs while ground truth poles are visible within the images. This sample shows that point supervision results in detection of the full pole, sometimes with predicted blobs occurring on both edges of the pole (row 3, column 3).

Table 3 presents pole detection model performance for each of the three camps in the Kalobeyi test area, given ResNet50 and ResNet101 encoder backbones. Distance based Precision, Recall and F1-score for the *strict* and *all* match variants, across three trials are shown. Higher F1-scores are observed when using a ResNet101 backbone



Figure 4. Test set sample images with predicted poles from the trained pole detection model shown as red blobs.

compared to a ResNet50 backbone. This aligns with expectation as the more expressive ResNet101 backbone learns better pole representations within the dataset. Predicted blobs many times encompass the whole pole, such that the centroid of the blob (used to compute precision) falls some distance away from the ground truth point (occurring at a pole edge). Thus, we evaluate performance at 3 distance thresholds (5m, 7m, 10m), as we observe that poles cast varying shadow lengths in the imagery (see supplementary material). Improved detection is also observed as the distance thresholds are increased. When pole shadow lengths for a random sample of 100 poles in the test set are measured, less than 30% of randomly sampled poles have a pole length under 5m. At this distance, viable detections are rejected for poles casting shadows larger than 5m. Thus at a distance threshold of 10m, even using the *strict* one-to-one match, we observe improved performance for both ResNet backbones over a 5m and 7m. Note that 90% of sampled poles cast shadows less than or equal to 10m. We also observe small variances between the *strict* and *all* match variants, indicating that the model is mostly making single detections for each ground truth pole. We

Table 3. Pole detection results for 3 camps in the test region, given ResNet50 and ResNet101 architectures. Metrics are reported for a strict (one-to-one) match between the ground truth and predictions, and the all (many-to-one) match, given 3 distance thresholds.

	ResNet50 & 10m					ResNet101 & 10 m				
	P_S	P_A	R	$F1_S$	$F1_A$	P_S	P_A	R	$F1_S$	$F1_A$
K1	0.66 ± 0.06	0.70 ± 0.05	0.66 ± 0.02	0.66 ± 0.03	0.68 ± 0.02	0.78 ± 0.01	0.82 ± 0.01	0.63 ± 0.01	0.70 ± 0.01	0.71 ± 0.01
K2	0.56 ± 0.01	0.62 ± 0.02	0.78 ± 0.05	0.65 ± 0.02	0.69 ± 0.03	0.63 ± 0.05	0.68 ± 0.03	0.72 ± 0.05	0.67 ± 0.05	0.70 ± 0.04
K3	0.63 ± 0.03	0.64 ± 0.02	0.80 ± 0.04	0.70 ± 0.00	0.71 ± 0.00	0.76 ± 0.02	0.78 ± 0.01	0.62 ± 0.06	0.68 ± 0.04	0.69 ± 0.04
	ResNet50 & 7m					ResNet101 & 7 m				
	P_S	P_A	R	$F1_S$	$F1_A$	P_S	P_A	R	$F1_S$	$F1_A$
K1	0.64 ± 0.05	0.65 ± 0.05	0.64 ± 0.02	0.64 ± 0.02	0.65 ± 0.02	0.77 ± 0.01	0.78 ± 0.01	0.67 ± 0.10	0.69 ± 0.01	0.70 ± 0.01
K2	0.54 ± 0.02	0.54 ± 0.02	0.76 ± 0.04	0.63 ± 0.02	0.63 ± 0.02	0.63 ± 0.05	0.64 ± 0.04	0.71 ± 0.05	0.66 ± 0.05	0.67 ± 0.04
K3	0.63 ± 0.03	0.63 ± 0.03	0.80 ± 0.04	0.70 ± 0.00	0.71 ± 0.00	0.75 ± 0.01	0.77 ± 0.01	0.62 ± 0.06	0.68 ± 0.04	0.69 ± 0.04
	ResNet50 & 5m					ResNet101 & 5m				
	P_S	P_A	R	$F1_S$	$F1_A$	P_S	P_A	R	$F1_S$	$F1_A$
K1	0.61 ± 0.04	0.62 ± 0.04	0.62 ± 0.02	0.62 ± 0.02	0.62 ± 0.01	0.76 ± 0.01	0.76 ± 0.01	0.62 ± 0.01	0.68 ± 0.01	0.68 ± 0.01
K2	0.52 ± 0.03	0.53 ± 0.02	0.74 ± 0.06	0.61 ± 0.03	0.61 ± 0.03	0.61 ± 0.04	0.61 ± 0.04	0.69 ± 0.05	0.65 ± 0.05	0.65 ± 0.05
K3	0.63 ± 0.03	0.63 ± 0.03	0.80 ± 0.04	0.70 ± 0.01	0.70 ± 0.01	0.75 ± 0.02	0.75 ± 0.02	0.61 ± 0.06	0.67 ± 0.03	0.67 ± 0.03

considered bounding-box object detection methods such as FasterRCNN [17], and observe that such object detection methods struggle to learn when a fixed sized bounding box is applied for the pole point labels. The uniformly sized label bounding boxes are either too small (leaving out a portion of the pole) or too big (including non-pole information) for a subset of the labels, thereby making it hard for the detection algorithm to consistently learn relevant pole features needed for meaningful detection.

Detecting electrical poles from aerial imagery can be a non-trivial task, as multiple pole-like structures might be predicted as electrical poles, while these structures are in fact not electrical poles. Hard negative mining is applied to suppress false positives, where the model is trained with hard negative examples. Table 4 shows performance across the three camps with and without hard negative mining, at the 10m distance threshold using the *strict* match and ResNet101 encoder. Hard negative mining improved model performance between 3-8 % in K2 and K3.

Line segmentation performance. Table 5 shows line segmentation model performance across the three camps in Kalobeyei. Mean Intersection-Over-Union (mIOU) is reported to measure alignment between the ground truth and predicted electrical lines, while precision, recall, F1-score are used to measure how well the model detects lines. For the purpose of mapping grid layouts, line detection is of higher importance when compared to perfect alignment.

The PGRID line segmentation model is able to detect electrical lines with F1-scores between 0.77-0.82, with the

lowest performance occurring in K2. The lower performance in the K2 camp is because lines are drawn connecting poles while there are no visible lines in the imagery. Intuitively, connecting poles with lines makes sense however, we do not run a post-processing step connecting lines and poles, as our partners also observe that there are areas in the integrated settlement where the grid is still under-construction. Across 3 trials, there is small variance in model performance, indicating the stability of the model in segmenting electrical lines. For our application, model outputs serve as a guide for mappers to accept/reject predictions, rather than as a standalone annotator of the grid. Figure 5, shows sample line predictions (left) and a small portion of the predicted grid (right) given the ground truth labels for poles and electrical lines, in the test region.

5.3. Comparison to a global power systems dataset

Information on power grid layouts is valuable for measuring electricity access and for allocating resources for grid expansion. Open access power system datasets such as the ground truth *Kenya Electricity Network* dataset [19] provided by Kenya Power and Light Corporation (KPLC) or approximations of Low Voltage (LV) coverage such as [3], provide relevant information about distribution-level grid layouts and electricity access. However, these datasets are often incomplete especially in remote areas. Our approach for reconstructing grid layouts from overhead imagery, provides a pathway for augmenting existing power grid datasets, thereby providing more comprehensive estimates of power systems coverage. Figure 6 shows the

Table 4. Comparison of F1-scores using ResNet101 architecture when model is trained with and without hard negative mining. We report results at a distance threshold of 10m.

	$F1_S$			$F1_A$		
	K1	K2	K3	K1	K2	K3
No - HNM	0.70	0.64	0.60	0.72	0.66	0.61
HNM	0.70	0.67	0.68	0.71	0.70	0.69

Table 5. Performance of line segmentation model, reporting alignment with the mIOU and detection accuracy with the F1-score.

	mIOU	Precision	Recall	F1-score
K1	0.70 ± 0.01	0.83 ± 0.00	0.82 ± 0.00	0.82 ± 0.00
K2	0.63 ± 0.01	0.75 ± 0.00	0.81 ± 0.00	0.77 ± 0.00
K3	0.67 ± 0.00	0.77 ± 0.01	0.85 ± 0.01	0.81 ± 0.00

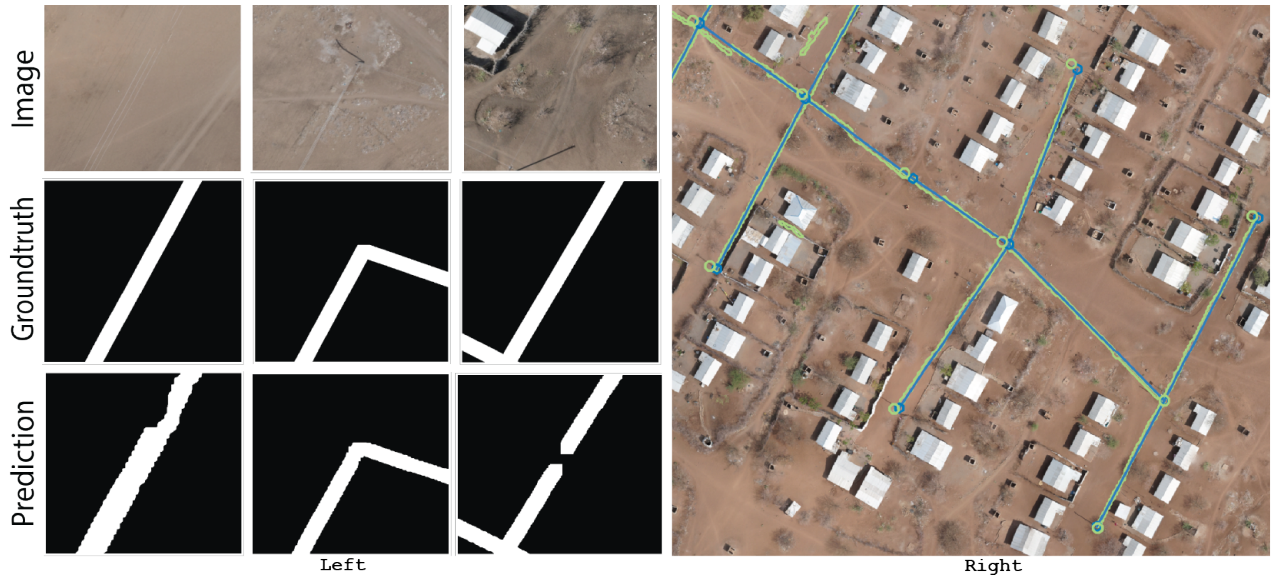


Figure 5. **Left:** Sample images showing electrical lines (top), ground truth lines buffered by 2 m (middle) and model predictions (bottom). **Right:** Sample ground truth power line (blue) alongside on predicted power line (green). Electrical poles are shown as circles.

newly mapped grid coverage after running inference with our model relative to the low voltage estimate in the global power systems dataset by [3]. The figure shows the presence of a power grid (red dots) using PGRID estimates, relative to the global power systems dataset (blue squares) for Kakuma camp at a 250m/grid cell resolution, the native resolution of the global power systems dataset. Only the Kakuma Camp is shown, as no LV data is available from open source datasets for other regions in the Turkana Integrated Settlement. With PGRID, we are able to include previously unmapped regions thereby augmenting existing estimates of grid topology. With growth in drone mapping and open imagery platforms such as OpenAerialImagery, PGRID can be applied across geographies to map grid topology in previously unmapped regions.

6. Conclusion

In this work, we show a real-world application for mapping power grid infrastructure in unplanned settlements using high-resolution aerial imagery. By implementing advanced pole detection and line segmentation models, we show the effectiveness of PGRID in accurately reconstructing power grid layouts in the Turkana region, particularly within the Kakuma and Kalobeyi Camps. Results showed that PGRID provides a detailed and accurate representation of electricity access in these challenging environments. The successful application of this method in resource-constrained settings and the provision of open source code, highlights its potential for broader use in other unplanned or underserved regions. Leveraging open data supports scalable solutions, providing high-precision guidance to map-

pers. This can greatly benefit humanitarian efforts and infrastructure planning, leading to better resource allocation, enhanced service delivery, and ultimately, improved quality of life for communities in these areas.

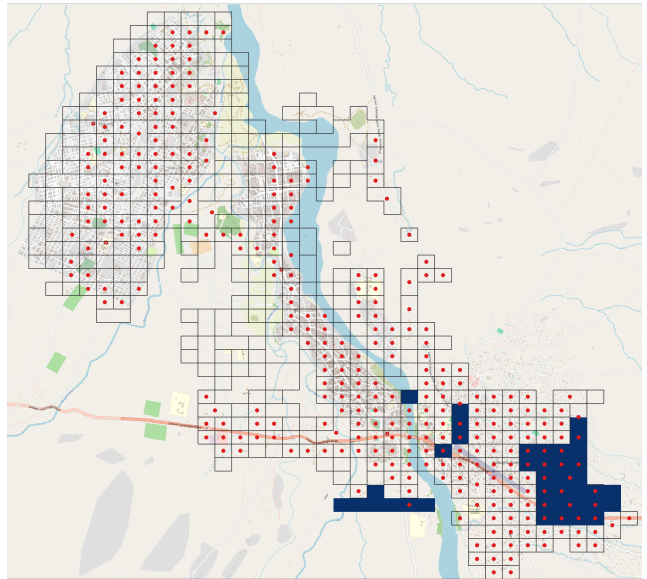


Figure 6. PGRID provides comprehensive mapping of power distribution grids, represented by red points, covering areas that are sparsely mapped in existing datasets, shown in blue. The existing dataset, based on [3], provides limited coverage of low-voltage power distribution grid for the Turkana Integrated Settlement. Grid coverage is depicted at the resolution of 250m/grid cell, demonstrating the increased granularity of our approach.

References

- [1] Rabab Abdelfattah, Xiaofeng Wang, and Song Wang. Ttpla: An aerial-image dataset for detection and segmentation of transmission towers and power lines. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 2
- [2] Deyu An, Qiang Zhang, Jianshu Chao, Ting Li, Feng Qiao, Zhenpeng Bian, and Yong Deng. Duformer: Solving power line detection task in aerial images using semantic segmentation. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, pages 54–66. Springer, 2023. 2
- [3] Christopher Arderne, Conrad Zorn, Claire Nicolas, and EE Koks. Predictive mapping of the global power system using open data. *Scientific data*, 7(1):19, 2020. 2, 7, 8
- [4] Anil Batra, Suriya Singh, Guan Pang, Saikat Basu, CV Jawahar, and Manohar Paluri. Improved road connectivity by joint learning of orientation and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10385–10393, 2019. 6
- [5] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *European conference on computer vision*, pages 549–565. Springer, 2016. 6
- [6] Serge Beucher and Fernand Meyer. The morphological approach to segmentation: the watershed transformation. In *Mathematical morphology in image processing*, pages 433–481. CRC Press, 2018. 3
- [7] Bonyan Organization. Importance Of Green Energy and Solar Panels for Camps. [https://bonyan.ngo/displaced-and-refugees/solar-panels-in-refugee-camps/#:~:text=\(80%25\)%20of%20the%208.7,have%20no%20access%20to%20electricity.,2024.1](https://bonyan.ngo/displaced-and-refugees/solar-panels-in-refugee-camps/#:~:text=(80%25)%20of%20the%208.7,have%20no%20access%20to%20electricity.,2024.1)
- [8] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 4
- [9] Development Seed. Machine learning using satellite images and map features. <http://devseed.com/ml-grid-detection/>, 2017. 2
- [10] Development Seed. Mapping the electric grid: Using ML to augment human tracing of HV infrastructure. <https://devseed.com/ml-grid-docs/>, 2018. 1
- [11] Dimitry Gershenson, Brandon Rohrer and Anna Lerner. A new predictive model for more accurate electrical grid mapping. <https://engineering.fb.com/2019/01/25/connectivity/electrical-grid-mapping/>, 2019. 2
- [12] Simone Fobi, Ayse Selin Kocaman, Jay Taneja, and Vijay Modi. A scalable framework to measure the impact of spatial heterogeneity on electrification. *Energy for Sustainable Development*, 60:67–81, 2021. 2
- [13] Bohao Huang, Jichen Yang, Artem Streltsov, Kyle Bradbury, Leslie M Collins, and Jordan M Malof. Gridtracer: Automatic mapping of power grids using deep learning and overhead imagery. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:4956–4970, 2021. 2, 5
- [14] Issam H Laradji, Negar Rostamzadeh, Pedro O Pinheiro, David Vazquez, and Mark Schmidt. Where are the blobs: Counting by localization with point supervision. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 547–562, 2018. 3
- [15] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 3
- [16] Wenpeng Luan, Joshua Peng, Mirjana Maras, Joyce Lo, and Brian Harapnuk. Smart meter data analytics for distribution network connectivity verification. *IEEE Transactions on Smart Grid*, 6(4):1964–1971, 2015. 1
- [17] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016. 7
- [18] Qinghu Tang, Zhecheng Wang, Arun Majumdar, and Ram Rajagopal. Fine-grained distribution grid mapping using street view imagery. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*, Vancouver, BC, Canada, pages 8–14, 2019. 2
- [19] The World Bank Group. Kenya - Kenya Electricity Network. <https://energydata.info/dataset/kenya-kenya-electricity-network>, 2024. 7
- [20] Moussa Blimpo Lauren Culver Jacob Kincer Meera Mahadavan Vijay Modi Bob Muhwezi Rose Mutiso Varun Sivaram Jay Taneja Mark Thurber Johannes Urpelainen Michael Webber Todd Moss, Morgan Bazalian. The Modern Energy Minimum: The case for a new global electricity consumption threshold. <https://energyforgrowth.org/wp-content/uploads/2021/01/FULL-Modern-Energy-Minimum-final-Jan2021-1.pdf>, 2024. 1
- [21] UNHCR. Refugees and Asylum-seekers in Turkana. <https://data.unhcr.org/en/country/ken/796>, 2024. 2
- [22] UNHCR The UN Refugee Agency. Global Trends Report 2023. <https://www.unhcr.org/global-trends-report-2023>, 2024. 1, 2
- [23] UNHCR The UN Refugee Agency. Refugee Camps Explained. <https://www.unrefugees.org/news/refugee-camps-explained/#Howmanyrefugeesliveinrefugeecamps?>, 2024. 2