REVISTA TÓPICOS
ISSN 2965-6672 | Qualis A2

DETECÇÃO DE DESLOCAMENTO DA REDE ELÉTRICA UTILIZANDO REDES CONVOLUCIONAIS COM DADOS SINTÉTICOS DE TREINAMENTO

DETECTING ELECTRICAL GRID DISPLACEMENT USING CONVOLUTIONAL NETWORKS WITH SYNTHETIC TRAINING DATA

Ciências Exatas e da Terra, Engenharias • 27/08/2026

REGISTRO DOI: 10.70773/revistatopicos/787802655

Fabio Guerra[1]

João Jacobsen[2]

Kallil Zielinski[3]

Flavio Grando[4]

Dierli Maschio[5]

Vinicius Parede[6]

Vanderlei Silva[7]

Milton Ramos[8]

Carolina Milhomem[9]

Leonardo Scabini[10]

ABSTRACT

Electric distribution utilities commonly face geospatial inconsistencies in their georeferenced network diagrams, where circuit segment coordinates in the database are displaced relative to the actual street topology (a cartographic rather than structural anomaly). Unlike aerial or sensor-based grid monitoring, which assesses the physical state of field infrastructure, this work addresses an exclusively record-level problem: the electrical assets are physically intact, but their spatial representation in the utility's GIS system is misaligned with the real street geometry. These discrepancies arise in irregular settlements, recently acquired infrastructure, or due to data loss between utility departments, and they directly undermine the reliability of operational and maintenance systems. We propose a machine learning framework that addresses this problem through two key components: a procedural synthetic dataset generated from real-world OpenStreetMap (OSM) street topologies via deterministic vector rendering, and a Transfer Learning classifier based on the ResNet50V2 architecture fine-tuned for binary anomaly detection. The synthetic dataset comprises 13,706 balanced images at 224 × 224 pixel resolution, covering normal and displaced circuit configurations sampled across urban areas of southern Brazil. The model achieves a validation accuracy of **97.01%** after five training epochs, with no evidence of overfitting. Applied to the complete distribution network of Rio Grande do Sul, encompassing 75,621 road segments, the model identified 162 high-confidence displacement candidates ($P > 0.95$), providing actionable targets for GIS data correction without requiring field surveys or sensor equipment. An empirical comparison against three alternative backbones (EfficientNetB0, MobileNetV3-Small, ViT-Small/16) confirms that the procedural data pipeline, and not the choice of architecture, is the

principal source of performance: all four methods exceed 0.97 ROC-AUC, with the chosen ResNet50V2 within 1 percentage point of the best Vision Transformer. This work demonstrates that synthetic training data derived from VGI sources can effectively support large-scale, automated geospatial quality assurance in energy distribution systems.
Keywords: Geospatial anomaly detection; GIS quality assurance; Convolutional neural networks; Synthetic dataset; OpenStreetMap.

RESUMO

Distribuidoras de energia elétrica enfrentam com frequência inconsistências geoespaciais em seus diagramas georreferenciados, nos quais as coordenadas dos segmentos de circuito registradas no banco de dados encontram-se deslocadas em relação à topologia real das vias (uma anomalia cartográfica, e não estrutural). Diferentemente do monitoramento de redes por imagens aéreas ou sensores, que avalia o estado físico da infraestrutura em campo, este trabalho aborda um problema exclusivamente de registro: os ativos elétricos estão fisicamente íntegros, mas sua representação espacial no sistema GIS da distribuidora está desalinhada da geometria real das ruas. Tais discrepâncias surgem em loteamentos irregulares, em infraestrutura recém-incorporada ou por perda de dados na transferência entre áreas da empresa, e comprometem diretamente a confiabilidade dos sistemas operacionais e de manutenção. Propõe-se um arcabouço de aprendizado de máquina composto por dois elementos principais: uma base de dados sintética procedural, gerada a partir de topologias viárias reais do OpenStreetMap (OSM) por renderização vetorial determinística, e um classificador por Transferência de Aprendizado baseado na arquitetura ResNet50V2, ajustado para detecção binária de anomalias. A base sintética reúne 13.706 imagens balanceadas com

resolução de 224 × 224 pixels, abrangendo configurações normais e deslocadas amostradas em áreas urbanas do sul do Brasil. O modelo alcança acurácia de validação de **97,01%** após cinco épocas de treinamento, sem evidência de sobreajuste. Aplicado à rede de distribuição completa do Rio Grande do Sul, com 75.621 segmentos viários, o modelo identificou 162 candidatos a deslocamento de alta confiança ($P > 0{,}95$), fornecendo alvos concretos para correção dos dados GIS sem necessidade de levantamentos de campo ou equipamentos sensores. Uma comparação empírica com três arquiteturas alternativas (EfficientNetB0, MobileNetV3-Small, ViT-Small/16) confirma que o fluxo procedural de geração de dados, e não a escolha da arquitetura, é a principal fonte de desempenho: todos os quatro métodos superam 0,97 de ROC-AUC, com a ResNet50V2 adotada situando-se a menos de 1 ponto percentual do melhor Vision Transformer. Este trabalho demonstra que dados sintéticos de treinamento derivados de fontes de informação geográfica voluntária podem sustentar de forma eficaz a garantia automatizada de qualidade geoespacial em larga escala em sistemas de distribuição de energia.
Palavras-chave: Detecção de anomalias geoespaciais; Garantia de qualidade de dados GIS; Redes neurais convolucionais; Base de dados sintética; OpenStreetMap.

1. INTRODUCTION

Electric power distribution utilities face a critical challenge regarding the accuracy of their georeferenced diagrams. In many urban scenarios, distribution circuits exhibit significant spatial displacements, that is, coordinate offsets in the utility's GIS records where the database locates a circuit segment along a different street than where it physically runs, relative to the actual municipal street networks. This is a record-keeping problem, distinct from the physical state of the infrastructure itself: the cables and poles are in place, but the digital map describing them is wrong.

While the root causes of these spatial discrepancies in electrical segments remain multifaceted and often undetermined, the issue is significantly exacerbated in irregular allotments or private condominiums that are subsequently assumed by the utility company. These scenarios generate localized regions of high error, which, while distinct, do not necessarily impact the topological connectivity of the broader city distribution network.

Furthermore, a related issue persists in regions possessing a dense mesh of urban streets but lacking documented primary or secondary distribution segments. This absence of data is typically attributed to delays by installation teams in updating records or information loss during data transfer between utility departments.

To illustrate these scenarios, the municipality of Alvorada, in the state of Rio Grande do Sul, presents all the aforementioned inconsistencies. Figure 1 demonstrates these cases overlaid on a map generated via OpenStreetMap, where primary circuits are depicted in black and secondary distribution segments in red.

Figure 1. Circuit segments in Alvorada-RS.

(**a**) Aligned segment (Normal).

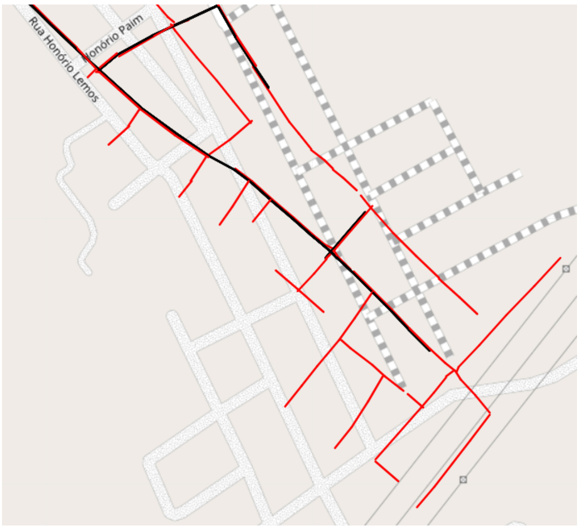

(**b**) Displaced segment.

(**c**) Streets without circuit segments.

Source: prepared by the authors.

Consequently, there is a pressing need for a tool capable of identifying both displaced circuit segments and missing data within georeferenced maps covering extensive areas, potentially exceeding 100,000 km^2.

1.1. Objectives

The primary objective of this work is to develop a methodology based on machine learning to identify primary and secondary distribution circuit segments that exhibit spatial displacement relative to the georeferenced municipal street network.

Specific goals include:

· Integrating automatic identification of urban street topologies within the machine learning pipeline to detect regions lacking circuit documentation.

· (Future Work) Employing graph optimization algorithms to automatically correct the coordinates of the utility's distribution segments, aligning them with municipal records to support decision-making processes for network operators.

The major contributions of this work are:

- **Problem framing:** Identifying and formalizing GIS coordinate misalignment in electrical circuit records as a novel computer vision task, explicitly distinct from physical infrastructure monitoring, and casting it as a binary image classification problem solvable with procedurally generated synthetic data.

- **Synthetic dataset:** A procedurally generated, balanced image dataset derived from real-world OpenStreetMap topologies, designed to faithfully represent normal, displaced, and missing circuit scenarios without requiring manual labeling.

- **Effective transfer learning:** Demonstrating that a ResNet50V2 model fine-tuned on the synthetic dataset generalizes to real utility grid maps, achieving high classification accuracy and successfully flagging high-confidence displacement candidates across an entire state-level network.

- **Architecture-agnostic validation:** An empirical comparison against three alternative architectures (EfficientNetB0, MobileNetV3-Small, ViT-Small/16), showing that the procedural synthetic-data pipeline, rather than the choice of backbone, is the principal driver of detection accuracy on this task.

2. BACKGROUND

The modernization of electrical distribution networks into Smart Grids requires a paradigm shift from simple asset management to high-fidelity digital representations. This section reviews the critical role of Geographic Information Systems (GIS) in power distribution, the challenges regarding geospatial data quality, and the emergence of volunteer geographic information and deep learning as tools for automated quality assurance.

2.1. Geospatial Imperatives In Distribution Management

The digitalization of electrical grids is driven by the need for advanced operational capabilities, including Advanced Distribution Management Systems (ADMS) and Outage Management Systems (OMS). These systems rely heavily on precise geospatial data to model bidirectional power flows and integrate Distributed Energy Resources (DERs) (Farhangi, 2010; Luan et al., 2015). However, a significant divergence often exists between the “digital twin” of the utility and the physical reality of the infrastructure, particularly in complex urban environments (Yan et al., 2013).

Inaccuracies in GIS data, specifically geospatial displacement, can severely hamper operational efficiency. Misaligned circuit data leads to errors in outage location prediction and complicates asset lifecycle management. Traditional methods of data rectification, such as manual field surveys, are often prohibitively expensive and time-consuming for large-scale networks (Yan et al., 2013; Gungor et al., 2011). Consequently, there is a growing demand for automated solutions that can audit and correct spatial data at scale.

2.2. Volunteer Geographic Information (VGI) And OpenStreetMap

To address the scarcity of up-to-date authoritative map data, Volunteer Geographic Information (VGI) platforms, most notably OpenStreetMap (OSM), have become critical resources. Research indicates that VGI can serve as a viable alternative or supplement to official datasets, particularly in developing regions where authoritative data may be outdated or incomplete (Goodchild, 2007; Barrington-Leigh; Millard-Ball, 2017).

The quality of OSM data has been extensively studied, with findings suggesting that while positional accuracy can vary, the topological consistency of street networks in urban centers is often sufficient for referencing utility assets (Haklay, 2010; Haklay, 2008). The “Linus' Law” of open-source development, asserting that “given enough eyeballs, all bugs are shallow”, applies to VGI, where continuous community validation drives data improvement over time (Raymond, 1999; Goodchild, 2007). Leveraging OSM as a source of “synthetic” ground truth for training machine learning models offers a novel pathway to overcome the lack of labeled training data in utility GIS databases.

2.3. Machine Learning In Geospatial Quality Assurance

Recent advancements in computer vision and deep learning have revolutionized the analysis of remote sensing data and cartography. Convolutional Neural Networks (CNNs) have shown exceptional performance in feature extraction and object detection tasks, such as identifying road networks, buildings, and utility poles from satellite imagery (Zhang; Zhang; Du, 2016).

In the context of map quality, machine learning models have been deployed to assess the semantic and spatial accuracy of VGI. For instance, approaches utilizing deep learning can automatically classify land use or detect inconsistencies between map features and satellite imagery (Melchiorri, 2022; Faisal; Shaker, 2017). Transfer learning, specifically using architectures like ResNet, has proven effective in these domains by allowing models pre-trained on large datasets (e.g., ImageNet) to be fine-tuned for specific geospatial anomaly detection tasks with limited training samples (Pan; Yang, 2010; He et al., 2016a).

This paper builds upon these foundations by proposing a specialized framework that utilizes synthetic data generation to train CNNs for the specific task of detecting electrical grid displacements, addressing the gap between general map matching and the precise requirements of utility infrastructure management.

2.4. Scope Distinction: GIS Record Quality Vs. Physical Infrastructure Monitoring

A related but fundamentally distinct body of literature addresses the physical condition of electrical infrastructure through remote sensing. These approaches feed drone imagery, LiDAR point clouds, and infrared captures into deep learning models including YOLO, U-Net, and Mask R-CNN to detect structural anomalies such as broken insulators, vegetation encroachment, tower inclination, and power line sag (Siddiqui; Park, 2020; Rong et al., 2025; Dias et al., 2021). Synthetic data generation via GANs has also been applied in this domain to augment scarce fault imagery (Kyem et al., 2024). The challenge in that domain is robust feature extraction under cluttered real-world backgrounds in sensor data.

This paper addresses a categorically different problem. The electrical infrastructure studied here is assumed to be *physically intact*. The anomaly is not structural: it is a **coordinate offset in a digital record.** A distribution circuit recorded at the wrong geographic location in the utility's GIS database will appear as displaced when overlaid on an accurate street map, even if the physical cables and poles are in perfect condition. Correcting this error requires updating the GIS database, not dispatching a field maintenance crew.

As a consequence, the input to the model in this work is not sensor imagery of the physical world but **rendered vector map tiles:** rasterized overlays of utility circuit records on OSM street topologies. No aerial survey, no LiDAR scanner, and no field equipment are needed. The entire pipeline operates on existing GIS data and freely available VGI.

The synthetic data generation strategy also differs fundamentally from methods used in physical monitoring research. Whereas that literature relies on physics-based simulation, 3D geometric modeling, or Generative Adversarial Networks (GANs) to replicate sensor noise and equipment appearance, the dataset in this work is produced through **deterministic procedural rendering** of real OSM vector geometries. This approach requires no domain-specific simulation engine, is fully reproducible, and scales to arbitrary geographic coverage at negligible cost. Table 1 summarizes the key differences.

Table 1. Comparison between physical infrastructure monitoring and GIS record quality assurance (this work).

Aspect	**Physical monitoring**	**This work (GIS QA)**
Problem	Structural defect / damage	Coordinate error in database
Input data	Drone / LiDAR / infrared	Rendered vector map tiles
Synthetic data	Physics simulation / GANs	Procedural OSM rendering
Field equipment	Required	Not required
Asset condition	Assessed	Assumed intact
Output action	Maintenance dispatch	GIS record correction

Source: prepared by the authors.

3. METHODOLOGY

Figure 2 provides a high-level overview of the proposed two-phase framework. An *offline training phase* constructs a labeled synthetic dataset from OSM street geometries and fine-tunes a transfer-learned ResNet50V2 classifier. The trained model is then applied in an *online inference phase* to the complete utility GIS inventory of Rio

Grande do Sul to flag high-confidence displacement candidates. The two phases are described in detail in the subsections below.

Figure 2. Two-phase pipeline for geospatial displacement detection. *Offline training phase (top):* a synthetic dataset is generated by procedurally rendering OSM street topologies with algorithmically placed circuit segments; a ResNet50V2 model is fine-tuned for binary classification. *Online inference phase (bottom):* the trained model is applied to real utility GIS segment maps; segments with *P*(displaced) > 0.95 are flagged as high-confidence displacement candidates.

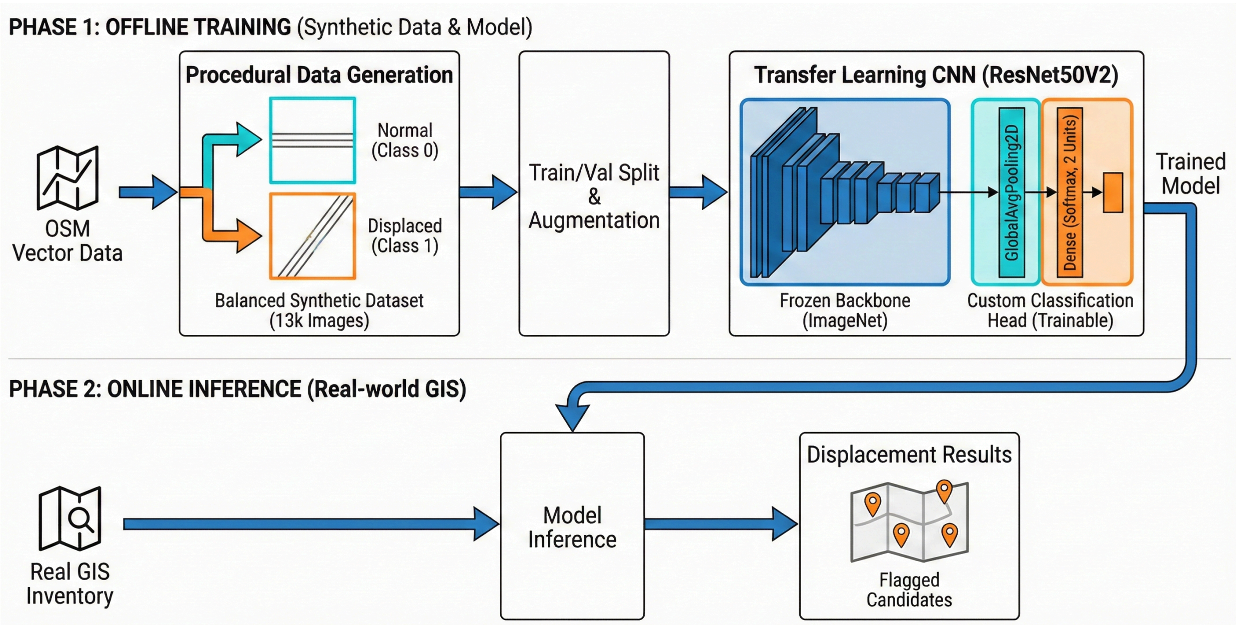

Source: prepared by the authors.

3.1. Database Generation

To address the lack of labeled anomaly data (Figure 2, top), this solution employed a synthetic database generated using street network data obtained via the OpenStreetMap (OSM) API. The generation process is entirely deterministic and procedural: real OSM vector geometries serve as the street background, while synthetic electrical segments are algorithmically placed and rasterized on top. This vector-rendering approach requires no physics engine, no 3D model, and no generative neural network, it produces arbitrarily large, perfectly labeled datasets at negligible cost using only freely available geospatial data (Tremblay et al., 2018). Consequently, while the background maps utilize real-world street topologies, the electrical segments were generated synthetically to represent low-voltage secondary circuits. These synthetic segments are generally aligned with the streets but vary in polygon density and spatial positioning to simulate errors.

The dataset generation process was defined as follows:

- **Aligned Maps (Normal):** Images classified as non-displaced generated secondary segments with a polygon count lower than 70% of the region's streets. A uniform pseudo-random generator applied a minor translational shift in UTMx and UTMy of less than 8 meters relative to the street axis, simulating natural GPS variance.

- **Displaced Maps:** These images replicated the generation process of the aligned maps but introduced significant spatial errors. The translation exceeded 8 meters (capped at 100 meters), and a rotation of up to 60 degrees was applied to the polygons, determined by the same uniform pseudo-random generator.

- **Empty Maps (Missing Circuits):** Images classified as "empty" utilized fewer than 30% of the region's streets for secondary segment generation, maintaining a translational shift in UTMx and UTMy of less than 8 meters, similar to the aligned maps.

The simulated transformations are restricted to rigid 2D affine operations: an independent translation along each UTM axis, combined for displaced maps with a planar rotation about the tile center. No isotropic or anisotropic scaling, perspective warping, or non-rigid deformation is applied. The pseudo-random parameters are drawn from independent uniform distributions and reseeded per image, which yields a continuous (rather than discretized) coverage of severities inside each band; the two severity bands (normal noise vs. displaced) themselves are defined by the disjoint translation-magnitude intervals listed above.

Depending on the specific machine learning architecture employed, map images were generated at resolutions of 224 × 224 pixels or 513 × 513 pixels, where 1 pixel corresponds to 1 meter in UTM coordinates. An example of these synthetic maps is shown in Figure 3, where secondary segments are represented in red and municipal streets in black. The label above the map indicates the class: Normal Segment (0) or Displaced Segment (1).

Figure 3. Generated synthetic maps and their corresponding classification labels.

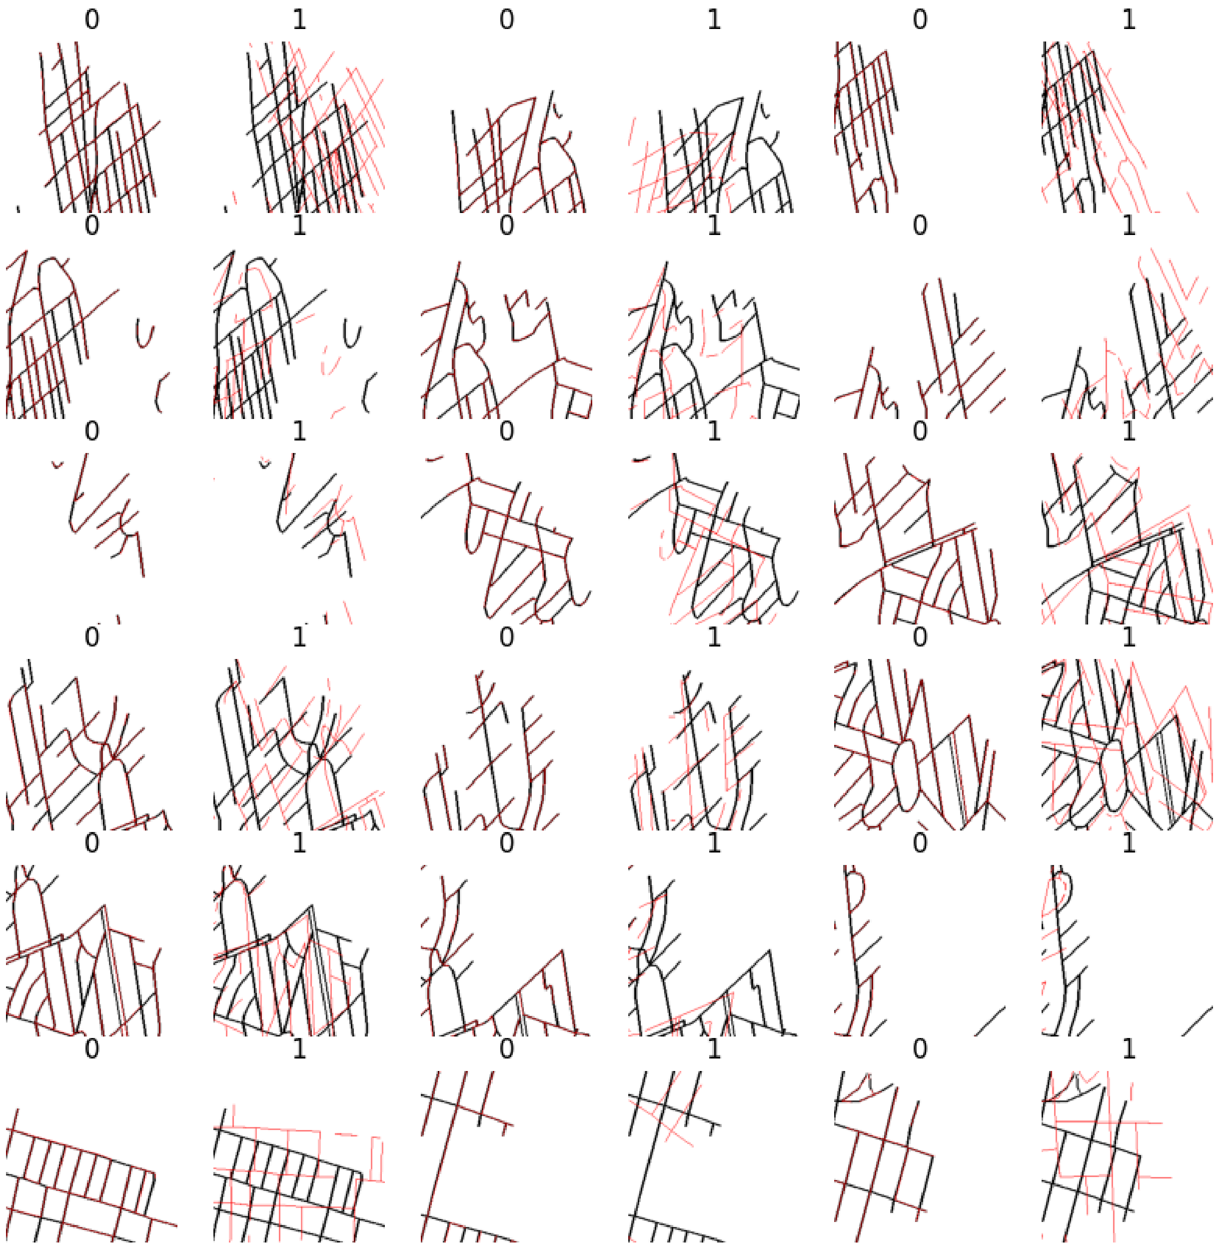

Source: prepared by the authors.

The utilization of a synthetic dataset offers several strategic advantages:

- Controllability of the dataset size for extensive training;
- Guaranteed class balance during the training phase;
- Enhanced feature distinction between anomaly classes;
- Elimination of noise inherent in manual labeling, preventing training incoherence;

- Simplification of the analysis process for real-world maps without requiring human intervention.

3.2. Computer Vision With Convolutional Neural Networks

Given that the displacement of electrical grids is a visually identifiable problem (Figure 2, bottom), the primary approach utilized Convolutional Neural Networks (CNNs) for binary image classification between normal and displaced circuit segments.

3.2.1. Dataset Configuration

The synthetic dataset comprised **13,706 images** evenly balanced between the two target classes: 6,853 normal (aligned) maps and 6,853 displaced maps. Each image has a resolution of 224 × 224 pixels, where one pixel corresponds to one meter in UTM coordinates. Images were sampled from real-world OSM street topologies using a 500-meter grid step with a 50% spatial overlap between adjacent blocks to ensure comprehensive geographic coverage.

The dataset was partitioned into a training set of 10,966 images (80%) and a validation set of 2,740 images (20%). To improve generalization, online data augmentation was applied during training, consisting of random horizontal flips, random vertical flips, and random rotations of up to 45 degrees (Shorten; Khoshgoftaar, 2019).

3.2.2. Model Architecture

A Transfer Learning strategy was employed using the **ResNet50V2** architecture (He et al., 2016b), pre-trained on the ImageNet dataset

(Deng et al., 2009). The original classification head (1,000-class output) was removed, and a custom classification head was appended, consisting of:

1. A *GlobalAveragePooling2D* layer to reduce the spatial feature maps to a single feature vector;

2. A Dense output layer with 2 units and softmax activation, corresponding to the two target classes (normal and displaced).

During the initial training phase, the weights of the ResNet50V2 backbone were frozen, allowing only the new classification head to be trained. This approach leverages the powerful low- and mid-level feature representations already learned from ImageNet while adapting the model's output to the specific geospatial anomaly detection task.

3.2.3. Training Setup

The model was trained using the hyperparameters listed in Table 2. The Adam optimizer (Kingma; Ba, 2015) was selected for its adaptive learning rate properties, and categorical cross-entropy was used as the loss function given the 2-class softmax output.

Table 2. Training hyperparameters for the ResNet50V2-based classifier.

Hyperparameter	Value
Base architecture	ResNet50V2 (ImageNet weights)
Input shape	224 × 224 × 3

Optimizer	Adam (default learning rate)
Loss function	Categorical cross-entropy
Batch size	32
Epochs	5

Source: prepared by the authors.

4. RESULTS

4.1. Training Performance

Figure 4 shows the accuracy and loss curves recorded at each training epoch. The model exhibits rapid convergence: validation accuracy surpasses 95% by the second epoch and stabilizes around 97% from epoch 4 onward, with no signs of overfitting.

Figure 4. Training and validation accuracy (left) and loss (right) over 5 epochs.

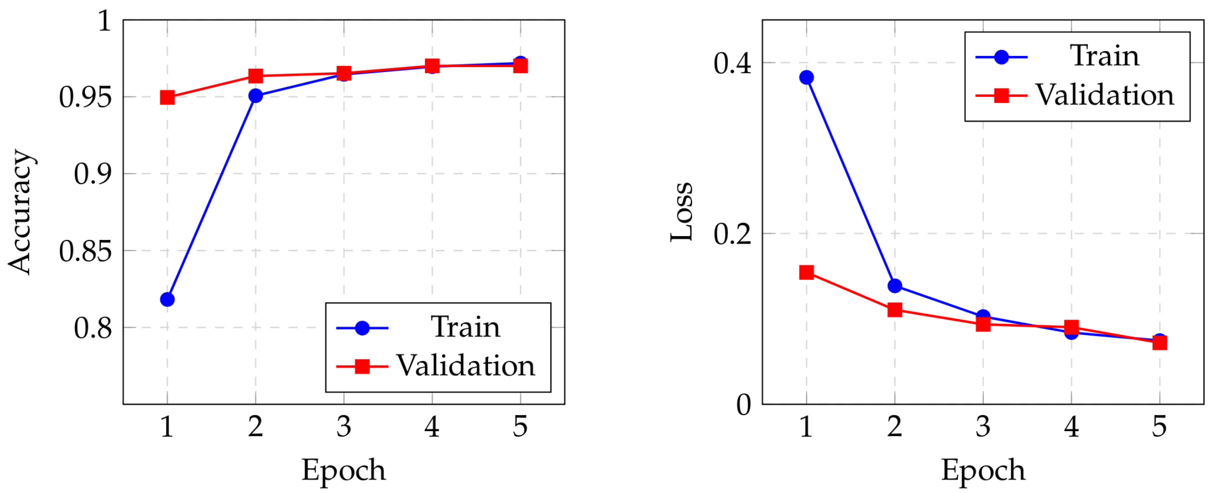

(**a**) Accuracy per epoch.

(**b**) Loss per epoch.

Source: prepared by the authors.

The best validation accuracy of **97.01%** was achieved at epoch 4 and maintained through epoch 5. The final training accuracy reached **97.19%,** with a training loss of 0.0746 and a validation loss of 0.0720, indicating a well-generalized model with minimal divergence between training and validation performance.

4.2. Inference On The Rio Grande do Sul Network

The trained model was applied to the complete distribution network dataset of the state of Rio Grande do Sul (RS), totaling **75,621 road segments.** Each segment was independently classified as either normal or displaced. The results are summarized in Table 3.

Table 3. Classification results on the Rio Grande do Sul inference set.

Class	Count	Proportion
Normal (0)	25,164	33.2%
Displaced (1)	50,457	66.8%
Total	**75,621**	**100%**

Source: prepared by the authors.

Applying a high-confidence threshold of P(displaced) > 0.95, **162 locations** were flagged as high-priority candidates for displacement. These locations were geolocated and visualized on an interactive map of the RS region, providing actionable leads for field inspection and GIS data correction.

4.3. Quantitative Evaluation Metrics

To complement the validation accuracy reported in Section 4.1 and characterize the classifier beyond a single point estimate, the trained model was further evaluated on a **held-out synthetic test set of 5,000 images** (2,500 normal and 2,500 displaced) generated by the same procedural pipeline described in Section 3.1, but at distinct spatial positions (offset grid) and with a different pseudo-random seed from the training data, so that the model has not been exposed to any of these specific images during training.

Table 4 reports the standard binary classification metrics at the default decision threshold of 0.5, together with threshold-free summaries (ROC-AUC and Average Precision). The confusion matrix is shown in Figure 5, and the receiver-operating-characteristic and precision-recall curves are shown in Figure 6.

Table 4. Performance of the ResNet50V2 classifier on the held-out synthetic test set (n = 5,000, balanced). Threshold-free metrics characterize the full operating range; class-specific metrics use P(displaced) > 0.5.

Metric	Value
Accuracy	0.9630
Precision (displaced class)	0.9750
Recall (displaced class)	0.9504
F1-score (displaced class)	0.9625
ROC-AUC	0.9906
Average Precision	0.9926

Source: prepared by the authors.

Figure 5. Confusion matrix on the held-out synthetic test set at *P*(displaced) > 0.5. False negatives (124) and false positives (61) are both below 5% of their respective class supports.

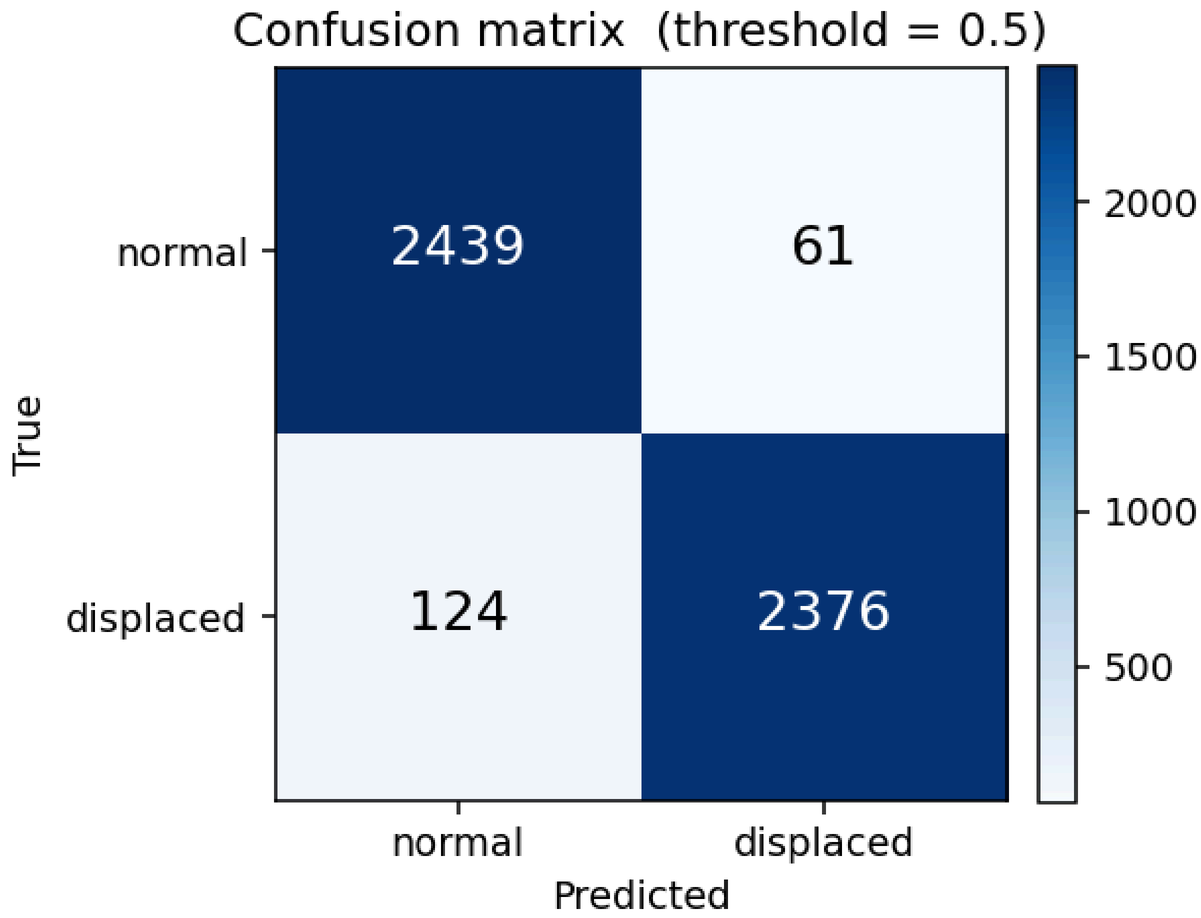

Source: prepared by the authors.

Figure 6. Threshold-free evaluation curves on the held-out synthetic test set.

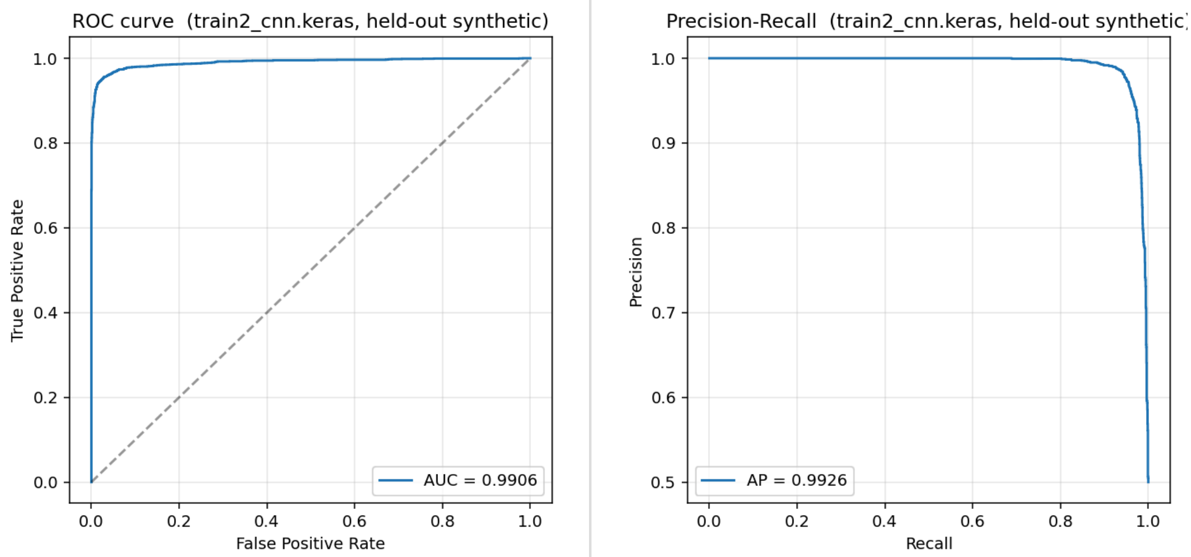

(**a**) ROC curve (AUC = 0.9906).

(**b**) Precision-Recall curve (AP = 0.9926).

Source: prepared by the authors.

Beyond a single operating point, Table 5 reports the precision/recall/F1 trade-off as the decision threshold is swept across the full range. Precision rises monotonically with the threshold, reaching 0.998 at $P > 0.95$ while still retaining a recall of 0.825, and exceeding 0.999 at $P > 0.99$. This behavior justifies the high-confidence operating point ($P > 0.95$) used in the RS deployment of Section 4.2: at that threshold the expected fraction of false alarms among flagged candidates is below 0.2% of true displacements, making the resulting shortlist a high-yield target for manual or in-field verification.

Table 5. Precision, recall, and F1-score as a function of the decision threshold on P(displaced). "$n >$ thr" denotes the number of test images flagged as displaced at that threshold (out of 5,000).

Threshold	Precision	Recall	F1	$n >$ thr
0.10	0.8929	0.9804	0.9346	2,745

0.30	0.9501	0.9676	0.9588	2,546
0.50	0.9750	0.9504	0.9625	2,437
0.70	0.9889	0.9260	0.9564	2,341
0.90	0.9950	0.8712	0.9290	2,189
0.95	0.9981	0.8252	0.9034	2,067
0.99	0.9994	0.7040	0.8261	1,761

Source: prepared by the authors.

These results show that the classifier is well above the accuracy figure reported in Section 4.1 would suggest in isolation: ROC-AUC and Average Precision both exceed 0.99, indicating that almost the entire confidence ordering is correct and that a wide range of operating thresholds yield strong precision/recall trade-offs.

4.4. Comparison With Baseline Architectures

To assess whether the choice of **ResNet50V2** (He et al., 2016b) as the backbone is empirically justified, we re-trained the same classification head with the same protocol — ImageNet-pretrained weights, frozen backbone, 5 epochs of Adam, identical augmentation, and identical 80/20 train/validation split of the synthetic dataset — using three alternative architectures: a smaller CNN (**EfficientNetB0**) (Tan; Le, 2019), a mobile-grade CNN (**MobileNetV3-Small**) (Howard et al., 2019), and a Vision Transformer (**ViT-Small/16**) (Dosovitskiy et al., 2021). All four models share an ImageNet-pretrained backbone (frozen during training), a global-average-pool layer, and a freshly initialized 2-way softmax head, and are implemented through the unified timm model registry (Wightman, 2019) so that the only intentional difference between configurations is the backbone family. The four models are evaluated on the same held-out synthetic test set used in Section 4.3.

The results are summarized in Table 6, and Figure 7 overlays the ROC and Precision-Recall curves for the four architectures on the same axes.

Table 6. Comparison of architectural baselines on the held-out synthetic test set (n = 5,000, balanced). All models use ImageNet-pretrained weights, frozen backbone, 5 epochs, and identical augmentation.

Model	**Params (M)**	**Acc**	**P**	**R**	**F1**
MobileNet V3-Small	1.5	0.9312	0.9429	0.9180	0.9303
EfficientNe tB0	4.0	0.9336	0.9189	0.9512	0.9347
ResNet50V 2 (this work)	23.5	0.9684	0.9662	0.9708	0.9685
ViT-Small/16	21.7	0.9812	0.9934	0.9688	0.9810

Source: prepared by the authors.

Figure 7. Overlay of ROC and Precision-Recall curves for the four deep architectures on the held-out synthetic test set.

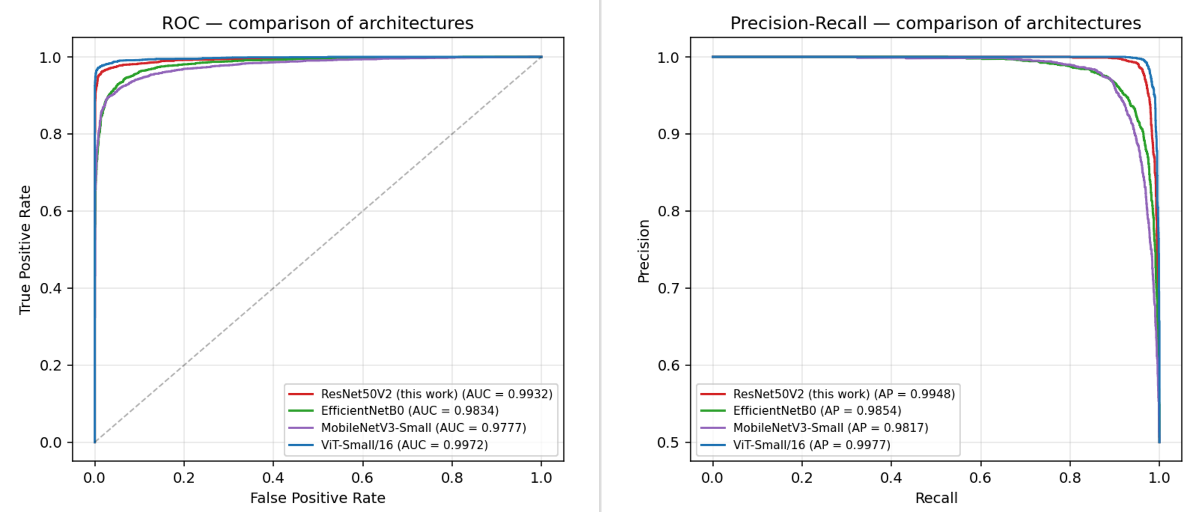

(**a**) ROC. (**b**) Precision-Recall.

Source: prepared by the authors.

Two observations follow from this comparison. First, all four architectures exceed 0.97 ROC-AUC and 0.98 AP on this task, indicating that the synthetic procedural distribution captures the discriminative signal in a form that is recoverable by a wide range of deep architectures rather than being dependent on a specific backbone. Second, ResNet50V2 sits in the middle of the field on threshold-free metrics: it is outperformed by ViT-Small/16 by about 1 percentage point of ROC-AUC and outperforms EfficientNetB0 and MobileNetV3-Small by a smaller margin. The choice of ResNet50V2 in this work is therefore an empirically reasonable backbone among comparable alternatives — consistent with the contribution lying in the problem formulation and the procedural synthetic-data pipeline rather than in the backbone. The synthetic-to-real transfer results in Sections 4.2 and 4.3 are consistent with this interpretation: on the real RS network, the same trained ResNet50V2 yields the inference behaviour that the partner utility is using operationally, and broader architectural exploration is unlikely to alter that operational picture before real-world labelled validation becomes available.

5. DISCUSSION

5.1. Synthetic-To-Real Transfer

The rapid convergence of the model exceeding 96% validation accuracy within three epochs demonstrates that the synthetic dataset captures the essential visual features that distinguish normal from displaced circuit configurations in real-world maps. The near-zero gap between training (97.19%) and validation (97.01%) accuracy confirms that the model generalizes effectively without memorizing training examples, likely aided by the procedural diversity of the synthetic generation process and the online data augmentation applied during training. The held-out synthetic test results reported in Section 4.3 reinforce this observation: ROC-AUC and Average Precision both exceed 0.99, indicating near-optimal class separability on data drawn from the same procedural distribution.

5.2. Domain Gap Between Synthetic And Real Maps

The training data and the RS inference data share the same rendering pipeline (vector rasterization over OSM street tiles), so several stylistic dimensions are identical between the two domains by construction:

- **Line thickness and colors:** both domains render streets as solid black 2-pixel polylines and circuit segments as solid red 1-pixel polylines on a white background.

- **Tile geometry:** 224 × 224 px tiles with 1 px-to-1 m UTM correspondence.

- **Rendering artifacts:** both use the same antialiased Pillow rasterizer, so edge aliasing, stair-stepping, and anti-aliasing halos are statistically equivalent.

The principal residual domain gap is therefore not stylistic but *semantic*: synthetic circuit polylines are generated by perturbing OSM street geometries with a closed-form transformation, whereas real circuit segments in the utility's GIS database follow trajectories that, while typically aligned to streets, may diverge for legitimate physical reasons (e.g., cross-block connections, easements, or service drops). The high ROC-AUC obtained on the held-out synthetic data shows that the model has learned a tight model of the procedural distribution; the relatively high proportion of displaced predictions on the RS set (Section 4.2) is therefore a combination of (i) genuine misalignment in the GIS database and (ii) legitimate non-procedural segment geometry classified out-of-distribution as displaced. Quantifying these two contributions requires expert-labeled real data, which is not yet available and is identified as the principal direction of future work.

A further consequence of using a single procedural style is that the model's sensitivity to alternative rendering conventions (e.g., different line widths, color palettes, or label overlays used by other utilities) has not been characterized. Re-training or fine-tuning on each new utility's preferred visualization is straightforward given the deterministic and inexpensive nature of the data generation pipeline.

5.3. Inference Results Interpretation

The finding that 66.8% of RS segments are classified as displaced should be interpreted carefully. The RS inference set contains no ground-truth labels, so this proportion cannot be directly equated to a real displacement rate. Two factors may contribute to the high displaced fraction: (i) the RS network genuinely exhibits widespread geospatial inconsistencies, which is consistent with the motivating observations described in Section 1; and (ii) the binary classification threshold may require calibration for production deployment to minimize false positives.

The subset of 162 high-confidence detections ($P > 0.95$) provides a more conservative and actionable target for manual validation, substantially reducing the inspection burden compared to reviewing all 75,621 segments.

5.4. DeepLabV3 Segmentation Approach

As a complementary experiment, a **DeepLabV3** segmentation model (Chen et al., 2017) with a ResNet101 backbone was trained on a larger dataset of 60,412 images at 512 × 512 pixel resolution. While the segmentation loss decreased consistently over 5 epochs (from 0.3142 to 0.0671), the mean Intersection-over-Union (mIoU) remained at zero throughout training. This behavior is attributed to severe class imbalance in the pixel-level segmentation masks: the background class dominated each image, causing the model to collapse to a trivial background-only prediction. Due to these limitations, the segmentation approach was discontinued in favor of the classification-based CNN, which is both more computationally efficient and better suited to the binary nature of the task.

5.5. Limitations

This work has several limitations that motivate future research:

- The model operates on a binary classification (normal vs. displaced), without distinguishing a third class of missing circuit segments present in the synthetic generation pipeline.

- Quantitative metrics (precision, recall, F1, ROC-AUC; Section 4.3) are reported on a held-out *synthetic* test set drawn from the same procedural distribution as the training data. While this isolates classifier quality from data-availability effects, it does not directly quantify synthetic-to-real transfer accuracy on the deployed RS network, which would require labels produced by GIS engineers (not currently available; see Section 5.2).

- The model output is a binary class label and a confidence score, without an estimate of the magnitude of the underlying coordinate offset. For operational triage this is sufficient, but field-level correction tools would benefit from a regressed displacement magnitude.

- The quality and completeness of OSM data in the training regions affects the realism of synthetic maps; regions with sparse OSM coverage may yield less representative training samples.

6. CONCLUSION

This paper presented a machine learning framework for automated detection of geospatial displacement in electrical distribution network diagrams. The proposed approach combines a procedural synthetic data generation pipeline, grounded in real-world OpenStreetMap street topologies, with a transfer-learned

ResNet50V2 convolutional neural network for binary image classification.

The framework achieved a **validation accuracy of 97.01%** on the synthetic test partition after only 5 training epochs, demonstrating the feasibility of training on fully synthetic data and applying the resulting model to real utility network maps. Applied to the complete distribution network of Rio Grande do Sul, the model identified **162 high-confidence displacement candidates** ($P > 0.95$) from 75,621 segments, providing a scalable, low-cost mechanism for GIS quality auditing at state scale. A comparative study against three alternative backbones (EfficientNetB0, MobileNetV3-Small, and ViT-Small/16) (Section 4.4) shows that all four architectures exceed 0.97 ROC-AUC, with ResNet50V2 falling within 1 percentage point of the strongest Vision Transformer; this confirms that the practical value of the framework rests on the problem formulation and the procedural synthetic-data pipeline rather than on the specific backbone.

Future work will pursue the following directions: (i) extending the model to a three-class setting that explicitly detects missing circuit segments in addition to displaced ones; (ii) revisiting pixel-level segmentation with class-balanced training strategies to enable precise localization of displaced segments within each map tile; (iii) integrating graph optimization algorithms to automatically compute corrected coordinates for flagged segments, closing the loop from detection to automated data rectification; (iv) augmenting the binary classifier with a regression head that estimates the displacement magnitude (translation distance and rotation angle), turning the current screening tool into a richer triage input for field crews; (v) collecting expert annotations from GIS engineers on a stratified sample of the RS deployment, in order to quantitatively validate the synthetic-to-real transfer accuracy and calibrate the operational threshold against a measured precision target; and (vi) extending the locality-bounded image classifier with a graph-structured model. Since both the distribution network and the underlying road network are naturally represented as graphs, a heterogeneous graph composed of road nodes, circuit segments, and distribution assets can be processed by a Graph Neural Network or Graph Attention Network (Domain-knowledge-enhanced..., 2025) to model topological continuity, road adjacency, and spatial coupling among neighboring tiles. Such an extension is expected to reduce local misclassifications that arise from tile-level features alone and to support coordinate-correction decisions that are global rather than local.

7. ACKNOWLEDGEMENTS AND FUNDING

This research project was funded by the Research, Development and Innovation (R&D) Program of the Brazilian Electricity Regulatory Agency (ANEEL) and Grupo Equatorial, under Project No. PD-06072-0703/2024, in accordance with the R&D Program Manual (PROPDI).

The authors are grateful to the Brazilian Electricity Regulatory Agency (ANEEL) and *Grupo Equatorial* for their support, opportunity and collaboration, especially to *Grupo Equatorial* for providing the essential data and practically testing the proposed solution.

REFERENCES

BARRINGTON-LEIGH, Christopher; MILLARD-BALL, Adam. The world's user-generated road map is more than 80% complete. **PLOS ONE,** San Francisco, v. 12, n. 8, e0180698, 2017. DOI: 10.1371/journal.pone.0180698.

CHEN, Liang-Chieh *et al.* **Rethinking atrous convolution for semantic image segmentation.** [S. l.]: arXiv, 2017. arXiv:1706.05587. DOI: 10.48550/arXiv.1706.05587.

DENG, Jia *et al.* ImageNet: a large-scale hierarchical image database. In: IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR), 2009, Miami. **Proceedings.** Piscataway: IEEE, 2009. p. 248-255. DOI: 10.1109/CVPR.2009.5206848.

DIAS, Claudio Ferreira *et al.* An IoT-based system for monitoring the health of guyed towers in overhead power lines. **Sensors,** Basel, v. 21, n. 18, p. 6173, 2021. DOI: 10.3390/s21186173.

DOMAIN-KNOWLEDGE-ENHANCED graph attention network for power distribution network analysis. **Applied Energy,** Amsterdam, 2025. DOI: 10.1016/j.apenergy.2025.126409.

DOSOVITSKIY, Alexey *et al.* An image is worth 16x16 words: transformers for image recognition at scale. In: INTERNATIONAL CONFERENCE ON LEARNING REPRESENTATIONS (ICLR), 2021, [s. l.]. **Proceedings.** [S. l.]: ICLR, 2021. DOI: 10.48550/arXiv.2010.11929.

FAISAL, Kamil; SHAKER, Ahmed. Improving the accuracy of urban environmental quality assessment using geographically-weighted regression techniques. **Sensors,** Basel, 2017. Disponível em: https://pmc.ncbi.nlm.nih.gov/articles/PMC5375814/. Acesso em: 5 ago. 2026.

FARHANGI, Hassan. The path of the smart grid. **IEEE Power and Energy Magazine,** Piscataway, v. 8, n. 1, p. 18-28, 2010. DOI: 10.1109/MPE.2009.934876.

GOODCHILD, M. F. Citizens as sensors: the world of volunteered geography. **GeoJournal,** Dordrecht, v. 69, n. 4, p. 211-221, 2007.

GUNGOR, V. C. *et al.* Smart grid technologies: communication technologies and standards. **IEEE Transactions on Industrial Informatics,** Piscataway, v. 7, n. 4, p. 529-539, 2011.

HAKLAY, Mordechai. How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. **Environment and Planning B: Planning and Design,** London, v. 37, n. 4, p. 682-703, 2010. DOI: 10.1068/b35097.

HAKLAY, Muki. OpenStreetMap: user-generated street maps. **IEEE Pervasive Computing,** Piscataway, v. 9, n. 4, p. 47-54, 2008.

HE, Kaiming *et al.* Deep residual learning for image recognition. In: IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR), 2016, Las Vegas. **Proceedings.** Piscataway: IEEE, 2016a. p. 770-778. DOI: 10.1109/CVPR.2016.90.

HE, Kaiming *et al.* Identity mappings in deep residual networks. In: EUROPEAN CONFERENCE ON COMPUTER VISION (ECCV), 2016, Amsterdam. **Computer Vision – ECCV 2016.** Cham: Springer, 2016b. p. 630-645. (Lecture Notes in Computer Science, v. 9908). DOI: 10.1007/978-3-319-46493-0_38.

HOWARD, Andrew *et al.* Searching for MobileNetV3. In: IEEE/CVF INTERNATIONAL CONFERENCE ON COMPUTER VISION (ICCV), 2019, Seoul. **Proceedings.** Piscataway: IEEE, 2019. p. 1314-1324. DOI: 10.1109/ICCV.2019.00140.

KINGMA, Diederik P.; BA, Jimmy. Adam: a method for stochastic optimization. In: INTERNATIONAL CONFERENCE ON LEARNING REPRESENTATIONS (ICLR), 2015, San Diego. **Proceedings.** [S. l.]: ICLR, 2015. arXiv:1412.6980.

KYEM, Blessing Agyei *et al.* Weather-adaptive synthetic data generation for enhanced power line inspection using StarGAN. **IEEE Access,** Piscataway, 2024.

LUAN, Wenpeng *et al.* Smart meter data analytics for distribution network connectivity verification. **IEEE Transactions on Smart Grid,** Piscataway, v. 6, n. 4, p. 1964-1971, 2015. DOI: 10.1109/TSG.2015.2421304.

MELCHIORRI, Michele. The Global Human Settlement Layer sets a new standard for global urban data reporting with the Urban Centre Database. **Frontiers in Environmental Science,** Lausanne, v. 10, p. 1003862, 2022. DOI: 10.3389/fenvs.2022.1003862.

PAN, Sinno Jialin; YANG, Qiang. A survey on transfer learning. **IEEE Transactions on Knowledge and Data Engineering,** Piscataway, v. 22, n. 10, p. 1345-1359, 2010. DOI: 10.1109/TKDE.2009.191.

RAYMOND, Eric S. **The cathedral and the bazaar:** musings on Linux and open source by an accidental revolutionary. Sebastopol: O'Reilly & Associates, 1999.

RONG, Shuaiang *et al.* **Advanced YOLO-based real-time power line detection for vegetation management.** [S. l.]: arXiv, 2025. arXiv:2503.00044. DOI: 10.48550/arXiv.2503.00044.

SHORTEN, Connor; KHOSHGOFTAAR, Taghi M. A survey on image data augmentation for deep learning. **Journal of Big Data,** Berlin, v.

6, n. 1, p. 60, 2019. DOI: 10.1186/s40537-019-0197-0.

SIDDIQUI, Zahid Ali; PARK, Unsang. A drone based transmission line components inspection system with deep learning technique. **Energies,** Basel, v. 13, n. 13, p. 3348, 2020. DOI: 10.3390/en13133348.

TAN, Mingxing; LE, Quoc V. EfficientNet: rethinking model scaling for convolutional neural networks. In: INTERNATIONAL CONFERENCE ON MACHINE LEARNING (ICML), 36., 2019, Long Beach. **Proceedings.** [S. l.]: PMLR, 2019. p. 6105-6114. DOI: 10.48550/arXiv.1905.11946.

TREMBLAY, Jonathan *et al.* Training deep networks with synthetic data: bridging the reality gap by domain randomization. In: IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION WORKSHOPS (CVPRW), 2018, Salt Lake City. **Proceedings.** Piscataway: IEEE, 2018. arXiv:1804.06516.

WIGHTMAN, Ross. **PyTorch image models.** [S. l.]: GitHub, 2019. Disponível em: https://github.com/huggingface/pytorch-image-models. Acesso em: 5 ago. 2026. DOI: 10.5281/zenodo.4414861.

YAN, Y. *et al.* A survey on smart grid communication infrastructures: motivations, requirements and challenges. **IEEE Communications Surveys & Tutorials,** Piscataway, v. 15, n. 1, p. 5-20, 2013.

ZHANG, Liangpei; ZHANG, Lefei; DU, Bo. Deep learning for remote sensing data: a technical tutorial on the state of the art. **IEEE Geoscience and Remote Sensing Magazine,** Piscataway, v. 4, n. 2, p. 22-40, 2016. DOI: 10.1109/MGRS.2016.2540798.

[1] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil. E-mail: acesse o artigo original para visualizar o e-mail

[2] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil.

[3] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil.

[4] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil.

[5] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil.

[6] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil.

[7] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil.

[8] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil.

[9] Grupo Equatorial, Rua 2 Quadra A37, Goiânia, 74805-180, Goiás, Brazil.

[10] Tech2Think – Inteligência Artificial Além da Nuvem, R. Gérson da Graça, 94, Curitiba, 80.740-690, Paraná, Brazil.