

Dataset

Not peer-reviewed version

---

# A Curated RGB Object Detection Dataset of Urban Electrical Distribution Assets with YOLO Annotations

---

[Igor Garcia-Atutxa](#), Hodei Calvo-Soraluze, [Francisca Villanueva-Flores](#)\*

Posted Date: 10 April 2026

doi: 10.20944/preprints202604.0778.v1

Keywords: object detection; urban infrastructure; electricity distribution assets; smart cities; electric vehicles; RGB imagery; YOLO; dataset curation; grid mapping



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Dataset

# A Curated RGB Object Detection Dataset of Urban Electrical Distribution Assets with YOLO Annotations

Igor Garcia-Atutxa <sup>1</sup>, Hodei Calvo-Soraluze <sup>2</sup> and Francisca Villanueva-Flores <sup>3,\*</sup>

<sup>1</sup> Escuela Politécnica Superior. Universidad Católica de Murcia (UCAM), Av. de los Jerónimos, 135, 30107, Murcia, Spain

<sup>2</sup> Universidad de Deusto, Avenida de las Universidades, 24, Bilbao, Biscay, Spain

<sup>3</sup> Centro de Investigación en Ciencia Aplicada y Tecnología Avanzada, Instituto Politécnico Nacional. Boulevard de la Tecnología, 1036 Z-1, P 2/2, 62790 Atlacholoaya, Morelos

\* Correspondence: fvillanuevaf@ipn.mx

## Abstract

Open, well-documented datasets are essential for the reproducible development of vision systems for urban utility management. This Data Descriptor presents a curated RGB object-detection benchmark of four classes associated with electrical distribution and street-level utility assets: Inspection Chamber, Overhead-to-Underground Transition, General Protection Box, and Transformer Substation. The public release contains 997 valid image-label pairs partitioned into 698 training, 150 validation, and 149 test images. Images were acquired during 2019 in multiple localities across Spain, predominantly with a mobile phone and, in occasional cases, using Google Maps as a complementary visual source, and were manually annotated with LabelImg before export to YOLO format. During curation, four invalid image-label pairs were removed because at least one YOLO bounding box exceeded the normalized image domain. The benchmark contains 1,939 object instances, with marked class imbalance: General Protection Box accounts for 50.2% of objects whereas Transformer Substation represents 4.7%. Images are heterogeneous in size and viewpoint, ranging from 90 × 170 to 4160 × 4032 pixels, with a median resolution of 619 × 544 pixels and a median of two annotated objects per image. The public GitHub release is organized into images/, labels/, and metadata/ directories; metadata stores split definitions, classes.txt, data.yaml, inventory information, annotation schema documentation, and diagnostic summary figures. Beyond detector benchmarking, the dataset can support scalable mapping of visible distribution-grid assets, with potential value for smart-city digital twins and data-informed EV charging deployment.

**Keywords:** object detection; urban infrastructure; electricity distribution assets; smart cities; electric vehicles; RGB imagery; YOLO; dataset curation; grid mapping

---

## 1. Summary

Urban asset inventories and condition assessment are foundational tasks for municipal management, electricity distribution planning, and smart-city digitalization (1). Vision-based workflows can reduce manual inspection effort and enable large-scale asset mapping from existing RGB imagery. Generic object detection benchmarks such as COCO (2) have accelerated algorithmic development, but domain-specific datasets remain necessary when classes are operational infrastructure elements rather than everyday objects. Related work has shown the utility of street-level imagery for detecting and localizing urban drainage components such as manhole covers and storm drains (3), while reviews of vision-based power-line inspection have highlighted the growing relevance of computer vision and deep learning for electricity asset management (4,5). Recent studies in specialized imaging also underline the importance of rigorous performance assessment and

transferable deep-learning designs when curating benchmarks intended for real-world deployment (6,7).

The dataset described here addresses a narrower but operationally relevant task: object detection of four categories of visible electrical distribution assets and associated street-level utility elements. The four class identifiers correspond to Inspection Chamber, Overhead-to-Underground Transition, General Protection Box, and Transformer Substation. Although the images were collected in Spain, these asset types or close functional analogues are common in many distribution networks, so the benchmark can serve as seed data for broader efforts to map visible electricity-network components across cities and countries. Such machine-readable inventories could support grid modernization, EV charging rollout, and smart-city services that depend on more complete knowledge of low-voltage infrastructure (2,8–10).

The curated public release contains 997 valid image-label pairs distributed across training (698), validation (150), and test (149) subsets. The source images were collected during 2019 in different Spanish localities, primarily with a mobile phone and, in occasional cases, through Google Maps, and were subsequently annotated in LabelImg. During quality control, four invalid image-label pairs were removed because at least one YOLO bounding box exceeded the normalized image domain. The resulting benchmark contains 1,939 annotated objects. For distribution, the repository is organized into images/, labels/, and metadata/ folders so that any user can download the benchmark from Huggingface (<https://huggingface.co/datasets/Garcia-Atutxa/Urban-Electrical-Distribution-Dataset/tree/main>) while keeping the data, annotations, and supporting documentation clearly separated.

The archive also makes visible several limitations that are important for downstream users and reviewers. First, the class distribution is imbalanced, with General Protection Box dominating the object count and Transformer Substation being comparatively scarce. Second, image resolutions and viewpoints are heterogeneous, which is desirable for ecological realism but may also introduce domain shift. Third, the geographic scope is currently limited to Spain, so transfer to other regulatory, architectural, or streetscape contexts should be validated empirically before worldwide deployment. Finally, because a small subset of source imagery may come from third-party visual services, redistribution rights and attribution requirements should remain explicit in the final release.

## 2. Data Description


The public release is organized into three top-level folders: images/, labels/, and metadata/ (Table 1). Images are stored as JPEG/JPG RGB files, and each label is stored as a text file using normalized YOLO bounding-box coordinates. The images/ and labels/ folders mirror one another by standardized filename, whereas metadata/ stores the official train/val/test split files, classes.txt, data.yaml, and diagnostic figures. The public release contains 997 valid image-label pairs distributed across the official benchmark partition (Figure 1).

In the official benchmark, images range from  $90 \times 170$  to  $4160 \times 4032$  pixels, with a median resolution of  $619 \times 544$  pixels. The number of annotated objects per image ranges from 1 to 8, with a mean of 1.94 and a median of 2. No empty labels occur in the train, validation, or test subsets.

**Table 1.** Organization of the cleaned public release and recommended use of each folder.

Folder	Image-label pairs	Purpose	Recommended use
train	698	Official training subset	Use for model fitting
val	150	Official validation subset	Use for model selection and early stopping


test	149	Official hold-out subset	Use only for final evaluation
public release	997	Complete benchmark release	Use according to the train/val/test protocol



**Figure 1.** Benchmark composition. Left: number of images containing each class. Right: object distribution across the official train, validation, and test subsets.

### 2.1. Structure of the Dataset

The dataset is organized into four object categories relevant to electrical distribution infrastructure: Inspection Chamber, Overhead-to-Underground Transition, General Protection Box, and Transformer Substation. Representative examples of these classes are shown in Figure 2, where panels A–D correspond to each category, respectively. The benchmark split corresponds approximately to a 70/15/15 partition by image count, implemented through metadata files that define the 698-image training subset, 150-image validation subset, and 149-image test subset.



**Figure 2.** Representative examples of the four object classes included in the dataset: (A) Inspection Chamber, (B) Overhead-to-Underground Transition, (C) General Protection Box, and (D) Transformer Substation.


## 2.2. Class Distribution and Annotation Characteristics

Table 2 summarizes the class distribution in the public release. General Protection Box is the dominant class, accounting for 974 annotated objects (50.2% of the release), whereas Transformer Substation is the minority class, with 91 objects (4.7%). The imbalance should be considered explicitly when benchmarking detectors or reporting aggregate scores.

**Table 2.** Per-class statistics for public release.

Class ID	Class	Images	Objects	Objects (%)	Mean bbox area	Median bbox area
0	Inspection Chamber	291	409	21.09	0.0483	0.0198
1	Overhead-to-Underground Transition	396	465	23.98	0.0193	0.0107
2	General Protection Box	662	974	50.23	0.0238	0.0114
3	Transformer Substation	80	91	4.69	0.2926	0.2353

Bounding-box scales differ substantially across classes. Transformer Substation has the largest spatial extent, with a mean normalized box area of 0.2926 and a median of 0.2353, while Overhead-to-Underground Transition and General Protection Box are typically much smaller objects. This size heterogeneity is visible in Figure 3 and should be considered when defining input resolution, anchor strategies, or small-object evaluation protocols.



**Figure 1.** Annotation characteristics in the official benchmark. Left: histogram of the number of objects per image. Right: histogram of normalized bounding-box areas.

### 3. Methods

The methodological details below combine what can be verified directly from the supplied archive with author-provided information about image acquisition and annotation. Fields that remain pending for final submission are explicitly marked where necessary.

#### 3.1. Image Sources and Acquisition Context

The images were acquired during 2019 in different localities across the geography of Spain. Most samples were captured directly in situ with a mobile phone under naturally varying illumination and viewpoint conditions, while a smaller number were obtained through Google Maps as complementary visual material. This acquisition strategy produced realistic heterogeneity in framing, scale, and background clutter, which is valuable for robust detection benchmarking

The archive therefore contains heterogeneous street-level and facade-level RGB scenes, including close-up views of cabinets and transition elements as well as wider contextual images containing sidewalks, facades, poles, and surrounding urban furniture. Illumination, viewing angle, object-to-camera distance, and image resolution vary markedly across samples. This diversity is beneficial for real-world detector robustness, particularly in municipal and utility inspection settings where image capture is rarely standardized.

Before public release, the authors should still document inclusion and exclusion criteria, whether any geospatial metadata are distributed, and the legal basis for redistributing any third-party imagery. This is especially important for samples originating from Google Maps or other mapping services, since repository publication may require attribution and/or exclusion of images whose redistribution terms are not compatible with an open downloadable dataset.

#### 3.2. Annotation Format and Class Ontology

Each label file follows the standard YOLO line format: `<class_id> <x_center> <y_center> <width> <height>`, where `x_center` and `y_center` denote the box center and `width` and `height` denote box size, all normalized to the image dimensions. All images were manually annotated in LabelImg using axis-aligned bounding boxes and then exported in YOLO-compatible text format. In the cleaned release, the class mapping is documented as 0: Inspection Chamber; 1: Overhead-to-Underground Transition; 2: General Protection Box; and 3: Transformer Substation.

Based on visual inspection of the archive, Inspection Chamber corresponds to a street-access chamber or utility pit, Overhead-to-Underground Transition to an overhead-to-underground service transition element, General Protection Box to a wall-mounted utility protection cabinet, and Transformer Substation to a transformer-center or substation-type enclosure.

#### 3.3. Curation, Standardization, and Split Generation

The curation report included in the archive indicates that raw file names were not fully consistent and that a standardized naming convention was applied to the curated release. The benchmark samples distributed in the public repository were renamed sequentially (for example, `img_000001.jpg` and `img_000001.txt`) to ensure deterministic image-label pairing and easier automation. The final GitHub release separates the primary assets into `images/`, `labels/`, and `metadata/`, which simplifies download, version control, and reproducible reuse.

After manual annotation in LabelImg, quality control verified one-to-one pairing between images and labels, class-ID range consistency, YOLO line format, and normalized coordinate geometry. Only samples that passed these checks were retained in the published benchmark. Annotation defects were removed from the public release, while their exclusion is documented in the curation report for transparency.

The official benchmark uses 698 images for training, 150 for validation, and 149 for testing. In the public GitHub release, these partitions are recorded in `metadata/` rather than encoded as separate image folders. The `metadata/` directory also includes `classes.txt`, `data.yaml`, annotation schema, and

diagnostic figures. This structure keeps the benchmark compact while preserving all information needed to reproduce the recommended protocol.

### 3.4. Technical Validation

Technical validation focuses on annotation integrity and release transparency. Following LabelImg-based annotation, the exported YOLO files were programmatically checked for formatting and coordinate validity. During curation of the supplied archive, four image-label pairs were identified as invalid because at least one bounding box extended beyond the normalized image domain after de-normalization to image coordinates. These pairs were removed from the cleaned release. The released benchmark subsets themselves contain no empty labels and only valid image-label pairs.

Internal quality-control checks did not detect missing image-label pairs, ambiguous matches, or duplicate groups in the average-hash duplicate screen. Although this does not guarantee the absence of all near-duplicates, it provides a useful first-pass control against accidental redundancy. Figure 3 further shows that most bounding boxes occupy a small fraction of the image area, with the Transformer Substation class representing a notable large object exception.

The release should therefore be interpreted as a realistic but imperfect benchmark. Researchers are encouraged to report per-class precision-recall metrics, inspect false positives for visually similar urban structures, and account explicitly for class imbalance when comparing detection systems.

## 4. User Notes

The dataset can be used directly in YOLO-style training pipelines by pointing the training framework to the provided metadata/data.yaml file and the corresponding images/ and labels/ directories. Users seeking unbiased benchmark comparisons should preserve the official train/val/test partition defined in metadata/ and report results on the designated test set. Beyond benchmarking, the dataset can serve as seed data for semi-automated mapping of visible electrical distribution assets at municipal, national, and potentially international scale, thereby supporting digital grid inventory, EV charging rollout studies, and smart-city digital twins.

Since the current imagery comes from Spanish localities only, cross-country transfer should be validated explicitly before using the dataset for worldwide asset-mapping workflows.

**Author Contributions:** IGA: Formal analysis, Visualization, Writing – original draft, Conceptualization, Investigation, Validation, Writing –review & editing, Methodology, Software. HCS: Validation, Formal analysis, Writing – review & editing, Investigation, Software, Methodology, Visualization. FVF: Writing – review & editing, Writing – original draft, Formal analysis, Methodology, Conceptualization, Visualization, Validation, Investigation.

**Funding:** No external funding was reported.

**Data Availability Statement:** The dataset described in this draft should be deposited in Urban-Electrical-Distribution-Dataset at <https://huggingface.co/datasets/Garcia-Atutxa/Urban-Electrical-Distribution-Dataset/tree/main>. The recommended release contains cleaned public dataset folders (images/train, images/val, images/test; labels/train, labels/val, labels/test).

**Acknowledgments:** GPT-5.4, developed by OpenAI, was used to assist with the grammatical correction and refinement of the manuscript's writing.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## References

1. Silva NS e, Castro R, Ferrão P. Smart Grids in the Context of Smart Cities: A Literature Review and Gap Analysis. *Energies*. 27 de febrero de 2025;18(5). doi:10.3390/en18051186

2. Lin TY, Maire M, Belongie S, Bourdev L, Girshick R, Hays J, et al. Microsoft COCO: Common Objects in Context [Internet]. arXiv; 2015 [citado 17 de marzo de 2026]. Disponible en: <http://arxiv.org/abs/1405.0312> doi:10.48550/arXiv.1405.0312
3. Santos A, Junior JM, Silva J de A, Pereira R, Matos D, Menezes G, et al. Storm-Drain and Manhole Detection Using the RetinaNet Method. *Sensors*. 9 de agosto de 2020;20(16). doi:10.3390/s20164450
4. Nguyen VN, Jenssen R, Roverso D. Automatic autonomous vision-based power line inspection: A review of current status and the potential role of deep learning. *Int J Electr Power Energy Syst*. julio de 2018;99:107-20. doi:10.1016/j.ijepes.2017.12.016
5. Sharma P, Saurav S, Singh S. Object detection in power line infrastructure: A review of the challenges and solutions. *Eng Appl Artif Intell*. 1 de abril de 2024;130(C). doi:10.1016/j.engappai.2023.107781
6. Garcia-Atutxa I, Martínez-Más J, Bueno-Crespo A, Villanueva-Flores F. Early-fusion hybrid CNN-transformer models for multiclass ovarian tumor ultrasound classification. *Front Artif Intell*. 15 de octubre de 2025;8. doi:10.3389/frai.2025.1679310
7. Garcia-Atutxa I, Villanueva-Flores F, Barrio ED, Sanchez-Villamil JL, Martínez-Más J, Bueno-Crespo A. Artificial intelligence for ovarian cancer diagnosis via ultrasound: a systematic review and quantitative assessment of model performance. *Front Artif Intell*. 5 de noviembre de 2025;8. doi:10.3389/frai.2025.1649746
8. Zaman M, Puryear N, Abdelwahed S, Zohrabi N. A Review of IoT-Based Smart City Development and Management. *Smart Cities*. 19 de junio de 2024;7(3):1462-501. doi:10.3390/smartcities7030061
9. Wolniak R, Stecula K. Artificial Intelligence in Smart Cities—Applications, Barriers, and Future Directions: A Review. *Smart Cities*. 9 de junio de 2024;7(3):1346-89. doi:10.3390/smartcities7030057
10. İnci M, Çelik Ö, Lashab A, Bayındır KÇ, Vasquez JC, Guerrero JM. Power System Integration of Electric Vehicles: A Review on Impacts and Contributions to the Smart Grid. *Appl Sci*. 6 de marzo de 2024;14(6). doi:10.3390/app14062246

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.