Wind Farms – Appendix

Collection 11

Version 1 - BETA

General Coordinator

Washington de Jesus Sant'anna da Franca Rocha (UEFS)

Team

Diego Pereira Costa (GEODATIN/UEFS)

Rodrigo Nogueira de Vasconcelos (GEODATIN/UEFS)

Nerivaldo Afonso Santos (GEODATIN/UEFS)

Rafael Oliveira Franca Rocha (GEODATIN/UEFS)

Soltan Galano Duverger (GEODATIN/UEFS)

Deorgia Tayane Mendes de Souza (UEFS/PPGM)

Jocimara Souza Lobão (UEFS/PPGM)

1 Overview

Wind power converts the kinetic energy of moving air into electricity through wind turbines, whose siting is governed primarily by wind speed, air density, and the persistence of the local wind regime. Turbines are rarely deployed in isolation: they are grouped into wind farms (parques eólicos) that also comprise access roads, mounting platforms and foundations, internal cabling and one or more substations. In the Brazilian context, wind generation has expanded rapidly since the early 2010s, driven by public energy auctions, falling technology costs, and the exceptional wind resource of the semi-arid Northeast, where the trade winds over the Caatinga sustain some of the highest capacity factors in the world (ABEEólica, 2023).

This rapid growth creates a need for accurate, spatially explicit and temporally consistent mapping of wind infrastructure. Beyond the turbines themselves, the associated access roads and earthworks represent a measurable land-use footprint whose expansion should be monitored for land-use planning, environmental trade-off assessment and energy-transition accounting. As climate projections indicate potential reductions in hydroelectric output due to altered precipitation patterns, the strategic coupling of wind with solar generation is becoming increasingly vital for Brazil's long-term energy security and carbon-neutrality goals (Sampaio & González, 2017).

In this sense, wind farms, introduced in MapBiomas Collection 11 as a cross-cutting renewable-energy theme, represent a newly mapped land-use class across Brazil's territory. Because individual turbine platforms are near the limit of what medium-resolution optical sensors can resolve, we adopt a data-fusion approach that combines Sentinel-2 optical imagery with Sentinel-1 synthetic-aperture radar (SAR), together with a deep-learning semantic-segmentation model an Attention U-Net (Ronneberger et al., 2015; Oktay et al., 2018) with a ResNet-34 backbone (He et al., 2016). The model was trained over the all biomes, where most of the country's wind infrastructure is concentrated, and inference was extended to 1,671 wind farms across Brazil over the 2016–2025 series. Below, we describe the technical information and methodological steps adopted for this mapping effort.

2 Methods

2.1 General processing

We developed a deep-learning-based approach to map wind-farm infrastructure across Brazil by integrating Sentinel-2 optical and Sentinel-1 SAR imagery, an Attention U-Net convolutional neural network, and cloud-based processing within Google Earth Engine (GEE) (Gorelick et al., 2017) combined with local GPU processing. For each year, cloud-free optical and radar observations were reduced to an annual median composite and stacked into a nine-channel tensor at 10-metre resolution. Training samples were derived automatically from a national vector cadastre of turbine locations and OpenStreetMap access roads, using buffer-based label generation. An Attention U-Net model was trained on a local workstation with optimized hyperparameters and data-augmentation techniques to improve generalization. The resulting per-farm predictions were post-processed through morphological cleaning, a tower/road/substation separation rule and a temporal-consistency filter, then exported as GEE-compatible assets and resampled to a standardized 30-metre grid. Final maps were validated against known wind-farm records and biome boundaries. All processing steps are illustrated in Figure 1.

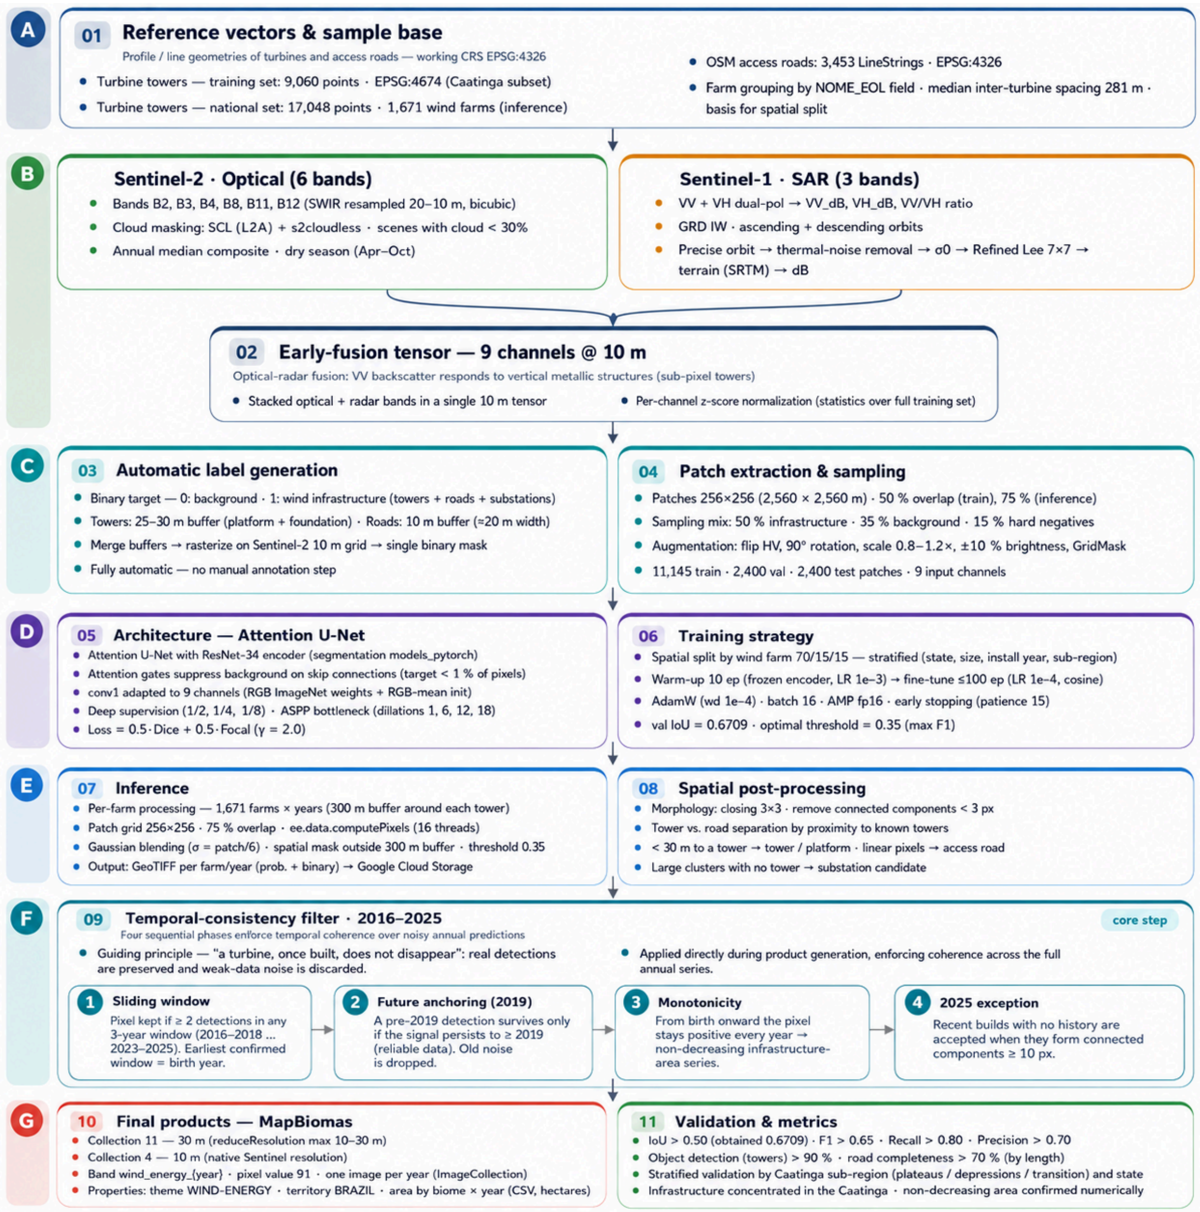

Figure 1. *Processing workflow. All processing steps are represented using distinct colors to illustrate each stage of the methodology, from the input vectors to the final MapBiomas products.*

The project adopts `EPSG:4326` as its working coordinate reference system; buffering and rasterization operations that must be performed in meters use a temporary UTM projection and are then returned to `EPSG:4326`.

2.2 Image acquisition

Two Sentinel missions, Sentinel-2 (Drusch et al., 2012) and Sentinel-1 (Torres et al., 2012), accessed through the GEE catalog, provide the input signal. They are fused early at the level of the input tensor into nine channels at 10 m and composited annually by the median over the dry-season window (April–October in the Caatinga), which maximizes the number of clean scenes.

Sentinel-2 — optical (6 bands). Bands B2, B3, B4, B8, B11 and B12 are used, with the 20 m SWIR bands resampled to 10 m by bicubic interpolation. Clouds are masked using the Scene Classification Layer (SCL) of the L2A product together with the `s2cloudless` probability mask, and an annual median is computed from scenes with less than 30% cloud cover.

Sentinel-1 — SAR (3 bands). The dual-polarization GRD IW product (VV and VH, ascending and descending orbits) is processed through the standard chain — application of the precise orbit file, thermal-noise removal, radiometric calibration to σ^0, Refined Lee (Lee, 1981) speckle filtering (7×7), range-Doppler terrain correction using SRTM, and conversion to decibels — yielding three derived bands: VV_dB, VH_dB and the VV/VH ratio.

The VV polarization is particularly sensitive to vertical metallic structures such as turbine towers, complementing the optical signal where the target is close to sub-pixel. All nine channels are normalized by a per-channel z-score, with the statistics computed over the entire training dataset.

2.3 Training samples and label generation

The mapping is posed as a binary semantic-segmentation problem — background versus wind infrastructure. Turbines, access roads and substations are not separated during segmentation because at 10 m their spectral–textural signatures are too similar; the distinction is made afterwards, in post-processing, by spatial proximity.

Class	Code	Definition
Background	`0`	Everything that is not wind infrastructure
Wind infrastructure	`1`	Turbines (platform / foundation) + access roads + substations

The starting point is a set of point locations of wind turbines and line vectors of the access roads. Labels are generated automatically, without a manual digitizing step:

- **Turbines:** a 25–30 m buffer is applied to each tower point, approximating the platform and rectangular foundation (~60–80 × 30–40 m).
- **Roads:** a 10 m buffer (20 m total width ≈ 2 pixels) is applied to the OpenStreetMap line features.
- **Merge and rasterization:** the turbine and road buffers are merged and rasterized onto the Sentinel-2 10 m grid, producing a single binary mask.

Training samples were taken over the Caatinga biome. From the composites and masks, 256×256-pixel patches (2,560 × 2,560 m) were extracted, with 50% overlap for training and 75% for inference. Patch sampling was balanced as 50% infrastructure, 35% background and 15% hard negatives, and augmented with 90° rotations, horizontal/vertical flips, 0.8–1.2× scaling, ±10% brightness/contrast, Gaussian noise and GridMask; arbitrary-angle rotation was avoided because it introduces artefacts in the SAR channels. The resulting dataset comprised 11,145 training, 2,400 validation and 2,400 test patches of nine channels each.

2.4 CNN model

An Attention U-Net convolutional neural network with a ResNet-34 encoder was trained to perform pixel-wise classification of wind infrastructure (Figure 2). The architecture follows the encoder–decoder design with skip connections to preserve spatial detail, with two adaptations for a rare, near-sub-pixel target:

- **Attention gates** on the skip connections suppress background activations, since the positive class accounts for under 1% of the pixels.
- **Deep supervision** adds auxiliary outputs at ½, ¼ and ⅛ scale, and an ASPP module at the bottleneck (dilations 1, 6, 12, 18) widens the receptive field.

The ResNet-34 backbone is instantiated through `segmentation_models_pytorch`. Its first convolution is adapted from three to nine input channels: the B2/B3/B4 channels inherit the ImageNet RGB weights, and the remaining six channels are initialized with the mean of those RGB weights.

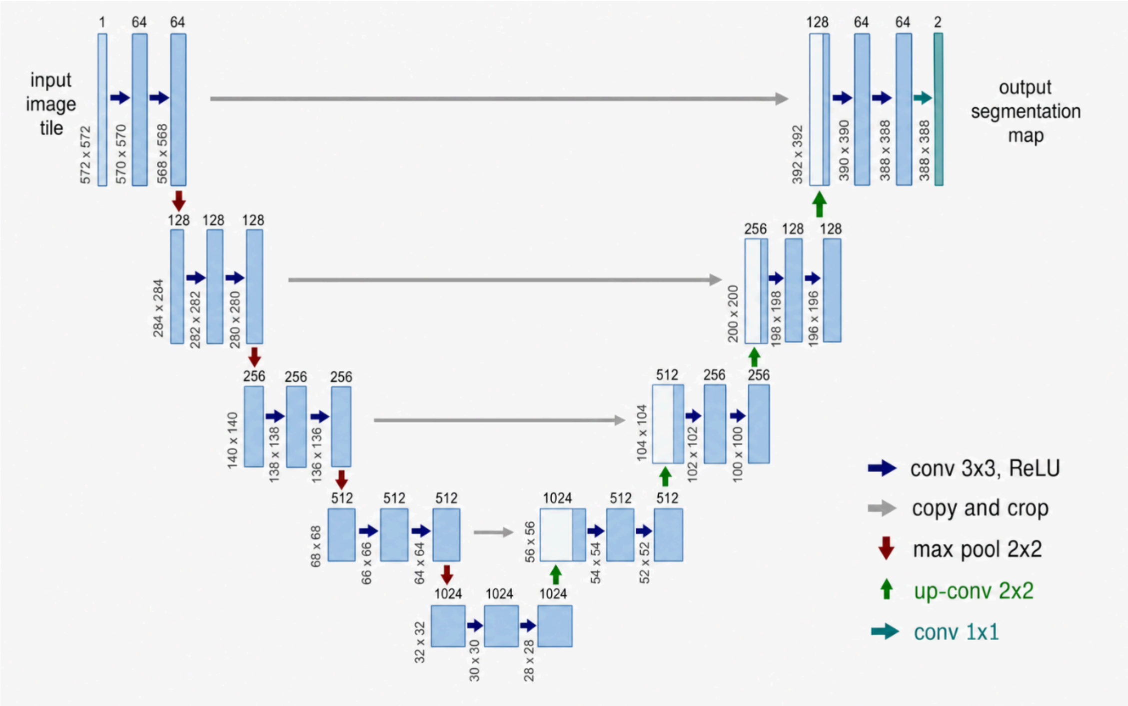

Figure 2. *Attention U-Net with a ResNet-34 encoder. The model processes 256 × 256 × 9 input tiles and outputs a binary mask distinguishing wind infrastructure from the background. Attention gates on each skip connection and an ASPP bottleneck adapt the network to a sparse, near-sub-pixel target. Fonte: RONNEBERGER et al. 2015*

Loss function. The network is optimized with a compound loss that balances region overlap and hard-pixel focus:

$$L = 0.5 \cdot \text{DiceLoss} + 0.5 \cdot \text{FocalLoss}\ (\gamma = 2.0)$$

The Dice term (Milletari et al., 2016) optimizes IoU directly for the rare class, while the Focal term (Lin et al., 2017) down-weights obvious background and concentrates learning on edges and difficult pixels.

2.5 Training strategy

Data were split spatially by wind farm never by patch into 70% training, 15% validation and 15% test, stratified by state, farm size, installation year and Caatinga sub-region, with roughly five hold-out farms never seen during training. Training proceeded in two transfer-learning phases: a 10-epoch warm-up with the encoder frozen (training the decoder and attention gates, LR = 1e-3), followed by fine-tuning of the whole model for up to 100 epochs (LR = 1e-4, cosine annealing).

Hyperparameter	Value
Batch size	16

Hyperparameter	Value
Optimizer	AdamW (weight decay 1e-4)
Scheduler	CosineAnnealingWarmRestarts
Precision	AMP (fp16-mixed)
Early stopping	patience 15 (mean IoU)
Hardware	NVIDIA RTX 4070 Super, 12 GB

The best model reached a validation IoU of 0.6709 for the infrastructure class. The operating threshold was set to 0.35, optimized on the validation set to maximize the F1 score (0.5331 at that threshold).

2.6 Inference

Inference is performed per wind farm rather than per MGRS tile: 1,671 farms are processed across the years of the series. A 300 m buffer is drawn around each individual turbine; because the median inter-turbine distance is 281 m, these buffers overlap and thereby also cover the connecting access roads. Within each farm:

- a grid of 256×256 patches is built with 75% overlap (stride 64 px);
- imagery is downloaded through `ee.data.computePixels` using 16 parallel threads;
- patch predictions are recombined by Gaussian blending (σ = patch/6), so that patch centres weigh more than the edges;
- pixels lying outside the 300 m buffer of any turbine are set to zero;
- the 0.35 threshold is applied to obtain the binary mask.

The inference output is one GeoTIFF per farm and year (band 1 = probability, band 2 = binary), uploaded through Google Cloud Storage to a prediction `ee.ImageCollection` with `year`, `park_name`, `state` and `threshold` properties.

3 Complementary database

The methodology relies on a complementary vector database of wind infrastructure, used both to generate labels and to guide per-farm inference. It draws on the wind-generation facilities registered in ANEEL’s Generation Information System, SIGA (ANEEL, 2024), and comprises turbine point locations and access-road lines:

Data	Records	Format / CRS	Use
Turbines (training)	9,060	Point · EPSG:4674	Caatinga subset — training
Turbines (complete)	17,048	Point · EPSG:4674	1,671 farms — inference
Access roads (OSM)	3,453	LineString · EPSG:4326	infrastructure labels

Turbines are grouped into wind farms through the `NOME_EOL` attribute, which also defines the spatial split. Their national distribution is strongly concentrated in the Northeast: Bahia (618 farms), Rio Grande do Norte (376), Piauí (169) and Ceará (168), followed by Rio Grande do Sul (139), Paraíba (114), Pernambuco (61), Maranhão (18), Santa Catarina (12), Minas Gerais (9) and other states (7). The geometry of the farms is regular, with a median inter-turbine distance of 281 m (mean 311 m, P90 474 m) a property exploited by the 300 m inference buffer.

4 Post-processing

Post-processing converts the raw annual probability masks into clean, semantically labelled and temporally consistent infrastructure maps, in three stages.

Morphological cleaning. A 3×3 morphological closing (`cv2.MORPH_CLOSE`) fills small gaps, connected components smaller than 3 px are removed, and reconnection of linear segments shorter than 100 m is planned for a future version.

Tower / road / substation separation. Although the model output is a single class, the elements are separated by spatial rules: infrastructure within 30 m of a known turbine is labelled tower/platform; the remaining linear pixels are labelled access road; and large clusters with no associated turbine are flagged as candidate substations.

Temporal-consistency filter. Raw annual predictions are noisy and, in the older part of the series (before 2019, where Sentinel-2 L2A coverage is sparse), unreliable. The filter enforces temporal coherence over the full 2016–2025 series under a simple physical principle once built, a wind turbine does not disappear through four operations:

1. **Sliding window.** a pixel is a candidate if it is detected in at least 2 of any 3 consecutive years (2016–2018 … 2023–2025); its birth year is the earliest confirmed window.
2. **Future anchoring (2019 anchor).** a detection born before 2019 survives only if the signal persists to 2019 or later (the reliable-data period); older noise that vanishes is discarded, while a real turbine is kept and correctly dated to its early birth year.

3. **Monotonicity.** from its birth year onward, a pixel remains positive in every subsequent year, which prevents real detections from disappearing and makes the mapped area non-decreasing.
4. **2025 exception.** recent constructions with no prior history are accepted when they form connected components of at least 10 px.

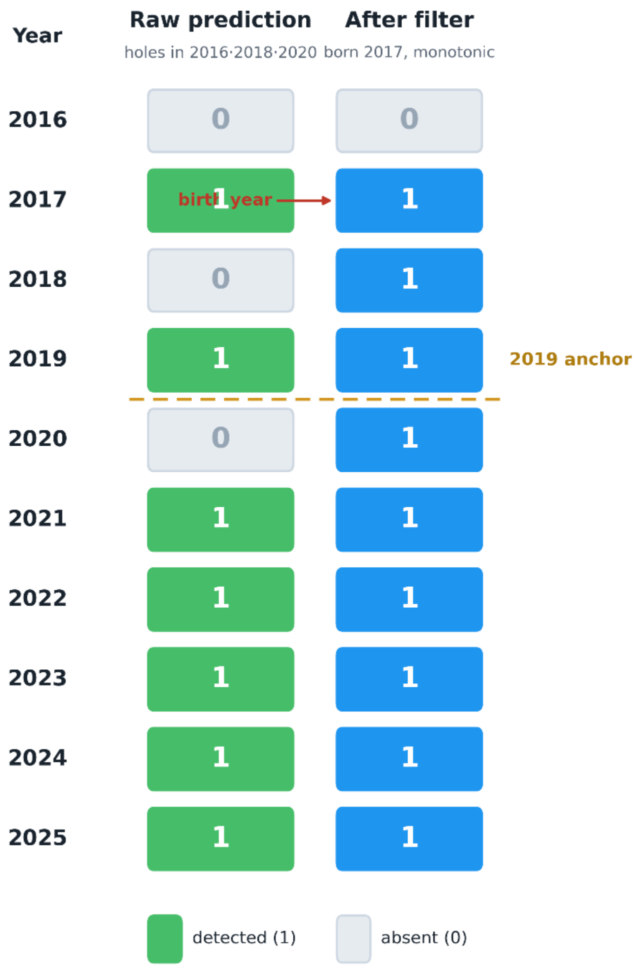

Figure 3. Temporal-consistency filter applied to a single pixel. The raw prediction (top) has spurious holes in 2016, 2018 and 2020; after filtering (bottom) the pixel is born in 2017 and remains positive thereafter. The dashed line marks the 2019 anchor.

A deliberate design consequence is that the mapped area is non-decreasing: this is physically correct because turbines are not removed, but it means the product does not represent decommissioning an acceptable simplification over the 2016–2025 horizon, where decommissioning is extremely rare.

5 Integration with biomes and cross-cutting themes

The temporal filter is applied directly during product generation, without an intermediate stage. Each year becomes one image in an `ee.ImageCollection`, with band `wind_energy_{year}` and pixel value 91 (wind infrastructure). Two products are generated:

Product	Resolution	Asset (ImageCollection)	Metadata
Collection 11	30 m	`.../COLLECTION-11/RENEWABLE-ENERGY/wind-energy`	`collection: 11`
Collection 4	10 m (Sentinel)	`.../LAND-COVER-10M/COLLECTION-4/RENEWABLE-ENERGY/wind-energy`	`collection: 4`

Each image carries the common properties `theme: WIND-ENERGY`, `source: geodatin`, `version: '2'`, `territory: BRAZIL` and `year`. The 30 m product aggregates the native 10 m mapping through `reduceResolution(max)`, while the 10 m product is native. Infrastructure area per biome and year is exported in hectares (CSV), which allows the monotonic behaviour of the series to be confirmed numerically; the overwhelming majority of the mapped infrastructure is concentrated in the Caatinga.

Validation and metrics. The mapping was evaluated against the target thresholds below, with validation stratified by Caatinga sub-region (chapadas, depressions and the Cerrado transition) and by state:

Metric	Target	Observation
IoU (infrastructure)	> 0.50	0.6709 obtained on validation
F1 (infrastructure)	> 0.65	recall-oriented
Recall (infrastructure)	> 0.80	avoid missing infrastructure
Precision (infrastructure)	> 0.70	—
Per-object detection (turbines)	> 90%	after separation
Road completeness	> 70%	by length

Composites, labels, training and inference over the test farms are complete; temporal validation against known construction dates and a full quantitative comparison remain in progress.

6 Limitations

Limitation	Mitigation
Sub-pixel turbines at 10 m	Label buffering + SAR (strong backscatter from metal)
Narrow, fragmented roads	Morphology + linear reconnection
Confusion with antennas / rural roads	Hard-negative mining + cluster rule
Exposed soil in the Caatinga resembling roads	SWIR bands + temporal median composite
No decommissioning in the product	Design decision (monotonicity); revisit for long series

The mapping is deployed in phases: Phase 1 (complete) training over the Caatinga and inference across Brazil; Phase 2 extension to the Cerrado and Pampa by transfer learning; Phase 3 national refinement. The processing stack comprises PyTorch, segmentation-models-pytorch and Lightning for modelling; rasterio/geopandas for geoprocessing; and Google Earth Engine with Google Cloud Storage for compositing, asset management and distribution.

All scripts used in this study are publicly available in the GitHub repository: https://github.com/mapbiomas/brazil-wind-farm.

7 References

ABEEólica – Associação Brasileira de Energia Eólica (2023). Boletim Anual de Geração Eólica. São Paulo, Brasil.

ANEEL – Agência Nacional de Energia Elétrica (2024). Sistema de Informações de Geração da ANEEL (SIGA). Available at: https://www.gov.br/aneel

Drusch, M.; Del Bello, U.; Carlier, S.; Colin, O.; Fernandez, V.; Gascon, F.; Hoersch, B.; Isola, C.; Laberinti, P.; Martimort, P.; et al. (2012). Sentinel-2: ESA's Optical High-Resolution Mission for GMES Operational Services. Remote Sensing of Environment, 120, 25–36. https://doi.org/10.1016/j.rse.2011.11.026

Gorelick, N.; Hancher, M.; Dixon, M.; Ilyushchenko, S.; Thau, D.; Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. Remote Sensing of Environment, 202, 18–27. https://doi.org/10.1016/j.rse.2017.06.031

He, K.; Zhang, X.; Ren, S.; Sun, J. (2016). Deep Residual Learning for Image Recognition. IEEE CVPR, 770–778. https://doi.org/10.1109/CVPR.2016.90

Lee, J.-S. (1981). Refined filtering of image noise using local statistics. Computer Graphics and Image Processing, 15(4), 380–389. https://doi.org/10.1016/S0146-664X(81)80018-4

Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; Dollár, P. (2017). Focal Loss for Dense Object Detection. IEEE ICCV, 2980–2988. https://doi.org/10.1109/ICCV.2017.324

Milletari, F.; Navab, N.; Ahmadi, S.-A. (2016). V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. Fourth International Conference on 3D Vision (3DV), 565–571. https://doi.org/10.1109/3DV.2016.79

Oktay, O.; Schlemper, J.; Le Folgoc, L.; Lee, M.; Heinrich, M.; Misawa, K.; Mori, K.; McDonagh, S.; Hammerla, N.Y.; Kainz, B.; et al. (2018). Attention U-Net: Learning Where to Look for the Pancreas. arXiv:1804.03999. https://arxiv.org/abs/1804.03999

Ronneberger, O.; Fischer, P.; Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. arXiv:1505.04597. https://arxiv.org/abs/1505.04597

Sampaio, P.G.V.; González, M.O.A. (2017). Photovoltaic solar energy: Conceptual framework. Renewable and Sustainable Energy Reviews, 74, 590–601. https://doi.org/10.1016/j.rser.2017.02.081

Torres, R.; Snoeij, P.; Geudtner, D.; Bibby, D.; Davidson, M.; Attema, E.; Potin, P.; Rommen, B.; Floury, N.; Brown, M.; et al. (2012). GMES Sentinel-1 mission. Remote Sensing of Environment, 120, 9–24. https://doi.org/10.1016/j.rse.2011.05.028