WindSense: Low-Cost IoT Sensing System for High-Granularity Wind Monitoring in Electric Distribution Networks

1st Alexandre Antunes Barcelos
Master Degree Program in Automation and Systems, CEFET-MG
Leopoldina, Brazil
alexandrebarcelos@cefetmg.br

2nd Ângelo Oliveira Rocha
Master Degree Program in Automation and Systems, CEFET-MG
Leopoldina, Brazil
angelo@cefetmg.br

3rd Fabiano Pereira Bhering
Master Degree Program in Automation and Systems, CEFET-MG
Leopoldina, Brazil
fabianobhering@cefetmg.br

4th Murillo Ferreira dos Santos
Master Degree Program in Automation and Systems, CEFET-MG
Leopoldina, Brazil
murillo.ferreira@cefetmg.br

5th Paolo Mercorelli
Institute for Production Technology and Systems (IPTS), Leuphana University of Lüneburg, Germany
paolo.mercorelli@leuphana.de

***Abstract*—This paper introduces WindSense, a low-cost Internet of Things (IoT) sensing system designed to mitigate failures in electric power distribution networks resulting from extreme weather events, by addressing the current lack of observational representativeness at low altitudes. The proposed architecture, which employs LoRa communication, MQTT/TLS protocols, and high-frequency sampling, overcomes the constraints of conventional hourly institutional measurements, which typically smooth out transient phenomena and micro-bursts that are critical for the structural integrity of poles and conductors. Validation through a locally deployed prototype, combined with event-centered analyses grounded in official emergency reports, demonstrates that the system not only captures localized dynamic wind behaviors in complex urban environments but also produces risk scores that offer substantial lead times, thereby enabling proactive operational preparedness and more precise post-event forensic assessment.**

***Index Terms*—IoT Sensing, Low-Altitude Wind Monitoring, Electric Distribution Networks.**

I. Introduction

Extreme meteorological phenomena have increasingly compromised the operational resilience of electric power distribution networks, as localized wind gusts and convective storm systems can induce cascading outages via conductor galloping, vegetation encroachment and contact, and structural component failures. Regulatory and operational communications have underscored the growing importance of resilience-enhancing measures and preparedness for periods characterized by intense precipitation and strong winds, particularly under summer convective regimes [1]. In a broader international context, large-scale assessments of weather-induced outages further indicate that a relatively small number of high-impact events can dominate overall risk profiles, thereby motivating the deployment of monitoring and forecasting strategies that enable earlier and more effective operational readiness [2].

A key limitation to operational preparedness is that many decision-making workflows depend on regional numerical weather prediction products whose grid-scale resolution cannot resolve the micro- and local-scale wind fields experienced by assets along streets, distribution feeders, and exposed spans. Even when such products include near-surface wind variables (e.g., 10-m wind), the adequate representativeness of these fields is constrained by the underlying model resolution and physical parameterizations, with deficiencies being especially pronounced in complex topography and densely built environments. Operational systems such as the Global Forecast System (GFS) provide gridded surface and near-surface wind fields as part of their routine outputs [3]. In contrast, state-of-the-art global forecasting systems (e.g., ECMWF IFS) offer high-quality global guidance at kilometer-scale resolution but yield grid-averaged approximations of the near-ground wind environment [4]. For engineering applications at the distribution-network scale, the mismatch between grid-scale winds and the local wind regime near obstacles and infrastructure underscores the need for complementary high-resolution local measurements.

In urban and peri-urban environments, near-surface wind fields are strongly modulated by surface roughness elements and building-induced flow channeling, resulting in pronounced spatial heterogeneity over relatively short horizontal length scales. This behavior is well documented in the literature on pedestrian-level wind and building aerodynamics, where flow distortion, acceleration, and recirculation occur in the vicinity of individual structures and within street canyons [5]. In alignment with international best practices, rigorous observational strategies necessitate careful attention to instrument siting, exposure conditions, and representativeness, thereby emphasizing that local in situ measurements can resolve phenomena that are not reliably captured by gridded datasets or coarse-resolution products [6]. Consequently, enhancing the spatial and temporal resolution of observations at or near the asset scale can substantially improve situational awareness and enable more targeted, operationally relevant preparedness measures.

In this context, this paper presents WindSense, a low-cost IoT sensing system designed to provide high-granularity wind monitoring close to the level of distribution assets. The system combines near-surface wind sensing and standard meteorological variables with a low-power embedded platform and long-range wireless telemetry. For long-range, low-power communication in unlicensed bands, LoRa-based links are widely recognized for their suitability in IoT monitoring scenarios due to coverage and energy efficiency [7], [8]. The architectural choices follow established principles of distributed sensing systems and wireless sensor networks, including modular sensor nodes, constrained communications, and resilient data logging [9], [10].

The main contribution of WindSense is not the replacement of national forecasting systems or institutional observing networks, but rather the enhancement of operational preparedness through the addition of a scalable layer of localized, high-resolution monitoring. Owing to its low per-unit cost, the system design supports dense sensor deployments, thereby improving spatial coverage and enabling a more refined characterization of near-surface wind dynamics along critical segments of the distribution network. The remainder of this paper describes the system architecture, experimental design, data processing workflow, and event-based analyses that quantify the operational benefits of high-granularity sensing for distribution-network monitoring.

II. System Architecture

A. System Architecture and Data Flow

The WindSense architecture is designed as an integrated pipeline that ensures data integrity from the point of measurement to final visualization. The system follows a layered approach optimized for distributed environmental sensing and interoperability with existing IoT network standards [9], [11]. As illustrated in Figure 1, the framework is divided into three functional tiers:

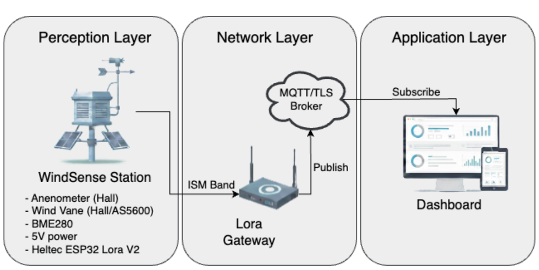

Fig. 1. WindSense layered architecture.

At the perception layer, each edge node performs local acquisition of near-ground atmospheric variables, including wind speed, direction, and ambient thermodynamic data. To ensure data provenance and reproducibility, these nodes generate precise, timestamped records directly at the point of origin.

Following data acquisition, the network layer enables telemetry transfer via LoRa-based sub-GHz communication protocols, selected for their intrinsic robustness to electromagnetic interference and signal attenuation in dense urban environments [7]. The gateway functions as a protocol translation and aggregation node, securely encapsulating and forwarding the acquired field telemetry into the Internet Protocol (IP) domain.

Finally, the application layer manages data distribution through a brokered publish/subscribe pattern using MQTT over TLS, ensuring both secure and lightweight messaging [12]. The resulting management interface provides event-centric visualization and traceable data exports, enabling rigorous auditability of recorded wind phenomena within their specific historical and operational contexts.

B. Data packet format and interoperability

To ensure traceability and interoperability across tools, each WindSense record is serialized as a JSON document [13] and includes an explicit timestamp in an Internet-standard profile for date-time representation (RFC 3339) [14]. This choice enables consistent merging of datasets from heterogeneous sources and reduces ambiguity in time alignment during event-centered analyses.

A typical packet includes the node identifier, timestamp, primary meteorological variables, and operational telemetry (e.g., supply voltage or battery status). When exporting logs for external analysis or documentation, the same records can be rendered into a tabular CSV representation to support broad compatibility with engineering tools and spreadsheets [15].

C. Web dashboard and operational monitoring

The WindSense web dashboard provides an operational view of the system state by subscribing to the MQTT broker and rendering time series, indicators, and derived quantities needed for inspection and reporting. This design decouples visualization from the embedded firmware, allowing the interface to evolve independently as new sensors, nodes, or deployment constraints are introduced.

Figure 2 presents a representative dashboard screen used in this work to demonstrate data visualization and reporting capabilities.

III. Methodology and Experimental Setup

An experimental deployment was carried out with sensors installed at approximately 6 m above ground on a local test pole. While distribution assets can reach up to 12 m depending on local standards and feeder geometry, the experimental installation provides a controlled baseline for verifying end-to-end operation, including sensing and transmission. In addition to the field prototype, the study evaluates event-oriented analyses using institutional meteorological data from Brazil's National Institute of Meteorology (INMET), whose platforms provide access to station data and visualization services [16],

Fig. 2. WindSense web dashboard: example visualization for a node (time series and operational indicators).

[17]. This enables consistent comparisons between the high-frequency WindSense measurements and hourly reference observations from surrounding stations, supporting engineering-oriented conclusions about the value of increased granularity near the ground.

The validation of the proposed methodology is organized into two primary investigative tracks: (i) an event-centered workflow and (ii) a prototype deployment and regional comparison. These tracks facilitate the evaluation of system performance and scalability prior to large-scale deployment. In both tracks, a quantitative risk score is derived from the calculated wind statistics and assessed against explicit, valid alert criteria, primarily focusing on signal persistence and minimum lead-time thresholds.

A. Track 1: Event-centered workflow (ISE + institutional meteorological series)

The first track establishes a reproducible experiment workflow anchored to official emergency reports and institutional meteorological series. This process involves comprehensive temporal harmonization, quality checks, and the definition of analysis windows aligned with specific weather events.

To establish a reproducible workflow, we analyze critical events using (i) official ISE documentation as the formal record of the occurrence and its time boundaries, and (ii) institutional meteorological series provided by INMET as external reference observations for the region [16], [17]. The workflow starts by standardizing timestamps to an interoperable representation (ISO 8601 / Internet timestamps) to ensure consistent time alignment across heterogeneous sources [14], [18]. Next, basic quality checks are applied to flag gaps, duplicates, and invalid values, preserving traceability of all transformations performed over the raw data [6].

To ensure interpretational operationality, each event is evaluated using fixed temporal windows anchored to the reference timestamp T (event onset). This allows quantifying the system's predictive efficacy by addressing a critical diagnostic metric: whether, for a predefined threshold, the system generates an alert with a statistically significant lead time prior to T.

The analysis focuses on wind-related descriptors commonly used in environmental monitoring (e.g., moving maxima, moving means, persistence above thresholds), following measurement and reporting practices consolidated in instrument and observation guides [6].

B. Track 2: Prototype deployment and regional comparison

The second track focuses on the system's empirical performance through a local prototype deployment in the municipality of Leopoldina. This deployment allows for a direct comparison between WindSense measurements and data from nearby institutional stations. The methodology addresses explicitly the discrepancies in sampling resolution, demonstrating how local, high-frequency sensing can capture transient behaviors and micro-bursts that are typically smoothed out or lost in coarser, low-resolution institutional datasets.

We evaluate WindSense's empirical behavior using a local prototype that generates high-frequency sensor data (5-minute sampling), enabling direct comparisons with nearby institutional stations used as regional references.

A key methodological step in this track is harmonizing sampling resolutions: INMET series are provided at 1-hour intervals, while WindSense samples every 5 minutes. Therefore, comparisons are performed by aligning each INMET timestamp to the nearest WindSense observation (or, when required, by aggregating WindSense within a small tolerance window around the institutional timestamp), preserving the physical meaning of each matched point and preventing bias from time offsets [14], [18]. This alignment strategy supports the central claim of the prototype track: local, high-frequency sensing can capture short-lived wind variability close to the ground that may be attenuated in coarser institutional datasets, especially when assessing event-adjacent dynamics.

C. Risk score construction and valid-alert criterion

In both tracks, the wind signal is summarized into a risk score $s(t)$ derived from short-horizon wind statistics computed on sliding windows (e.g., intensity, persistence, and rate-of-change proxies). The score is evaluated against explicit operational criteria so that an "alert" is not triggered by isolated fluctuations. In particular, an alert is considered *valid* only when it satisfies: (a) persistence (the score remains above a threshold L for a minimum number of consecutive samples) and (b) minimum lead time (the first valid alert occurs at least δ minutes before T). This definition makes the alert logic auditable and consistent with WindSense's 5-minute sampling and the reporting needs of engineering operations.

Whenever the analysis is treated as a classification problem (alert vs. no-alert within a defined positive time window), performance reporting follows standardized definitions of precision, recall, and F_1 to avoid ambiguity in interpretation and comparison [19]. When continuous-value agreement is assessed between series (e.g., WindSense vs. institutional stations), correlation and error metrics are reported with clear definitions, including Pearson correlation [20] and mean absolute error, which is recommended as a robust measure for average model performance in several practical contexts [21].

IV. Results and Discussion

This section consolidates the findings from the two investigative tracks into an operational interpretation: given a critical-event reference time T, we assess (i) whether a wind-derived risk score $s(t)$ would trigger a *valid* alert (persistence k and minimum lead time δ), and (ii) what changes when measurements move from regional/hourly context to local/5-minute sensing at the infrastructure-relevant height (pole level, up to 12 m).

A. Experiment 1: Event-centered risk-score behavior aligned to T (ISE workflow)

Figure 3 summarizes the event-centered analysis for Experiment 1, utilizing fixed temporal windows aligned to the event onset T. It standardizes the operational assessment across heterogeneous events by evaluating three performance dimensions: (i) *detectability*, defined as the exceedance of signal $s(t)$ over threshold L within the pre-T horizon; (ii) *lead time*, representing the interval between the initial valid threshold crossing and T; and (iii) *stability*, characterizing the distinction between transient peaks and sustained elevations above L:

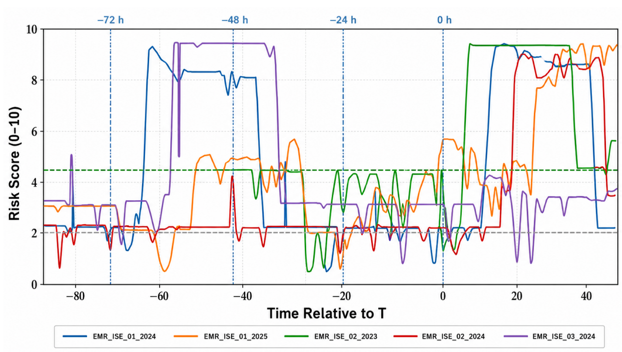

Fig. 3. ISE event-centered wind-risk score using fixed windows aligned to T (event start). Vertical markers indicate operational horizons (e.g., $T-72$, $T-48$, $T-24$, and T).

By anchoring the temporal reference to T across standardized horizons (e.g., $T-72$, $T-48$, $T-24$, and T), the visualization eliminates variance introduced by inconsistent scaling. This consistency allows for the direct interpretation of the signal's capacity to support preparedness protocols, such as crew staging and feeder prioritization, within explicit operational time budgets. The methodology adheres to interoperable time-indexing conventions and auditable criteria detailed in Section III-C [6], [14], [18].

A key quantitative output of this experiment is the *lead time* of the *first valid alert* as a function of threshold L. Table I reports the lead time (in hours) under the adopted valid-alert rule (persistence k consecutive confirmations and minimum lead time δ). Two operational patterns emerge. First, the scoring logic demonstrates significant predictive capacity for some events (tens to more than 100 hours), indicating that when pre-event wind statistics rise in a sustained manner, the workflow provides actionable preparedness margins. Second, the absence of lead time for several event–threshold combinations shows that a single global threshold is not uniformly anticipatory across all cases; in practice, this is expected when events differ in meteorological regime and local exposure, and it motivates cross-event calibration and validation (e.g., leave-one-event-out) to estimate generalization and avoid overfitting thresholds to a specific event. When framed as a binary alerting task within a predefined positive time window around T, reporting should follow standardized definitions of precision, recall, and F_1, ensuring consistent interpretation across events and thresholds [19].

TABLE I
Lead time (hours) of the first valid alert per event and threshold.

Event	L=4 (h)	L=5 (h)	L=6 (h)
`EMR_ISE_02_2023`	–	–	–
`EMR_ISE_01_2024`	83.0	–	–
`EMR_ISE_02_2024`	56.0	–	–
`EMR_ISE_03_2024`	145.0	145.0	145.0
`EMR_ISE_01_2025`	43.2	–	–

Table I indicates that when the institutional wind series exhibits sustained pre-event strengthening, the score can offer meaningful anticipation. However, non-detections ("–") in the pre-T horizon show that institutional series may not consistently expose micro- and local-scale wind behavior relevant at the distribution-asset height. These results suggest that scoring performance depends on the spatial representativeness of the input meteorological data, but that the observational representativeness of the input series can constrain what is detectable in advance. This interpretation aligns with guidance emphasizing exposure, siting, and representativeness in wind observations [6].

B. Experiment 2: Regional station-based score around the critical event

The same event-centered logic to institutional stations in the Leopoldina region was applied for Experiment 2, focusing on the critical event with reference time T= 04 Nov 2025, 15:00. The objective is twofold: (i) to characterize how a regional observational network expresses pre-event risk signals using the same score definition, and (ii) to provide an interpretable baseline for comparing with local pole-level sensing. The analysis preserves time-harmonization practices (ISO 8601/RFC 3339) and uses consistent horizons to avoid misinterpretation due to time offsets or inconsistent plotting scales [14], [18].

Figure 4 shows the aligned station-based risk score and the corresponding wind behavior in an extended window ($T-98$ to T+12 hours). This representation serves to evaluate whether the regional signal intensified with sufficient lead time and exhibited temporal persistence. In practice, station-based scores tend to reflect broader-scale forcing and exposure conditions at the station sites; therefore, pre-T elevations can represent regional storm approach and synoptic/mesoscale strengthening, but they can under-represent local accelerations and short-lived microbursts that occur near obstacles and within urban roughness fields. This is consistent with established evidence that near-ground winds in built environments are strongly shaped by roughness elements and flow channeling, producing significant variability over short distances [5], [6].

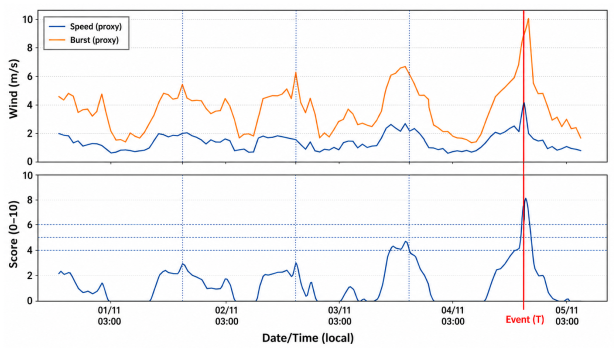

Fig. 4. Regional station-based risk score and wind behavior around T using an extended pre-event horizon ($T-98$) and post-event window ($T+12$). The vertical marker indicates T (04 Nov 2025, 15:00).

Operationally, Experiment 2 clarifies the role of institutional stations: they provide indispensable regional situational awareness and context for storm evolution, but they are not designed to guarantee representativeness at the distribution-asset level in terms of height and exposure. As a result, even when the regional signal is correctly elevated, the magnitude, timing, and intermittency of pole-level gust stress can remain partially unobserved. This observation motivates the prototype track: adding a dense, low-cost perception layer at pole level to directly measure the wind regime experienced by the infrastructure, complementing institutional networks rather than replacing them [6], [17].

C. Prototype Analysis: Local Sensing

The prototype track evaluates WindSense's empirical behavior through a local deployment, generating observations at a 5-minute sampling interval at approximately 6 m above ground level. A primary methodological challenge involves the resolution mismatch between the 5-minute WindSense data and the typically hourly institutional series. Consequently, comparative analyses necessitate rigorous timestamp harmonization (utilizing nearest-neighbor matching or constrained tolerance-window aggregation) to maintain the physical integrity of matched data points and mitigate artifacts derived from temporal misalignment [14], [18].

Figure 5 illustrates the information-resolution gap inherent in these disparate sampling frequencies. The engineering implications are significant: transient gust structures, rapid ramps, and localized peaks (critical factors for conductor oscillation, vegetation interference, and mechanical stress) are effectively captured and audited only through high-density sampling at the asset height. Increased temporal granularity thus transcends mere data detailing; it fundamentally shifts the observability of transient phenomena in the event-adjacent horizon, mainly where near-ground wind fields are modulated by structural obstacles and surface roughness [5], [6].

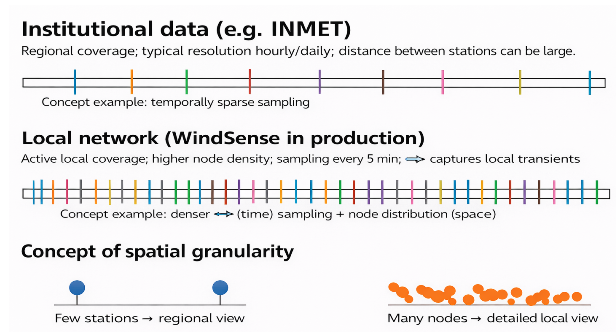

Fig. 5. Information granularity comparison between institutional hourly observations and WindSense 5-minute sampling at pole level (infrastructure-relevant height).

From a quantitative reporting perspective, two complementary evaluation modes are recommended and consistent with the methodology defined earlier. First, when comparing continuous-valued agreement between WindSense and institutional series (after harmonization), correlation and error metrics provide interpretable summaries of alignment; Pearson correlation and mean absolute error (MAE) are widely used and should be reported with explicit definitions to avoid ambiguity [20], [21]. Second, when evaluating *alerting* behavior, performance should be framed in terms of valid-alert criteria (persistence and lead time) and, when treated as classification, should report precision, recall, and F_1 using standardized definitions [19]. Together, these reporting practices ensure that the prototype results remain auditable: the system is evaluated not only by how closely it tracks regional references at matched timestamps, but also by whether it reveals transient dynamics that are operationally meaningful at the pole level and whether those dynamics can support earlier preparedness signals.

D. Discussion

The results across the three experiments yield a consistent operational interpretation. Experiment 1 demonstrates that an event-centered score, governed by auditable valid-alert rules, can achieve substantial lead times; however, its anticipatory performance is non-uniform when constrained by coarse or non-local observations. Experiment 2 indicates that while regional stations provide effective contextual awareness of storm evolution near T, they frequently under-represent the intermittency and amplitude of near-ground gusts (the primary mechanical stressors for distribution assets). The prototype track elucidates this discrepancy: at the infrastructure height (pole level), high-frequency sampling significantly increases the probability of capturing transient peaks and rapid ramps, generating traceable, representative measurements that complement institutional datasets.

Consequently, the primary operational takeaway is that the performance bottleneck is often not the scoring logic or forecasting capability, but rather the *observational representativeness* at the asset level. WindSense addresses this architectural gap via a low-cost, auditable sensing and delivery pipeline (node $\rightarrow$ gateway $\rightarrow$ MQTT/TLS $\rightarrow$ dashboard), enabling high-density deployments where wind stress is physically experienced. This infrastructure complements existing institutional networks, supporting both proactive preparedness (through pre-T readiness actions) and post-event forensic analysis based on time-stamped, local evidence [6], [12], [17].

V. Conclusion

This work demonstrates that the primary bottleneck in operational preparedness for wind-induced failures in electric distribution networks is not forecasting methodology *per se*, but rather the lack of *observational representativeness* at the infrastructure height. While regional meteorological products provide essential synoptic and mesoscale context, their spatial density and sampling resolution are often insufficient to capture the transient, obstacle-modulated wind behavior experienced by distribution assets in complex urban and peri-urban environments [5], [6].

Through an event-centered methodology anchored to the reference onset T, this study established that a wind-derived risk score governed by auditable, valid-alert rules can yield significant lead times, provided the input series captures sustained pre-event strengthening. However, the non-uniformity of this anticipatory capability across events—when relying on coarse or non-local observations—highlights a critical gap. The WindSense prototype deployment demonstrates that increasing temporal granularity and locating sensors at the pole level go beyond mere resolution enhancement; they fundamentally shift the observability of the event-adjacent horizon.

By sampling at 5-minute intervals at infrastructure-relevant heights, WindSense captures transient peaks, rapid ramps, and micro-scale gust structures frequently attenuated in hourly regional datasets. This provides utilities with traceable, time-stamped evidence of the localized wind regimes that directly induce mechanical stress on assets, thereby supporting both proactive preparedness and post-event forensic analysis.

Furthermore, this work demonstrates that these observability gains can be achieved through a low-cost, scalable IoT architecture. The integration of LoRa telemetry, MQTT/TLS data pipelines, and modular sensing nodes enables dense deployments along critical feeders at an affordable cost. Unlike conventional automatic weather stations (which are optimized for climatological representativeness and involve high capital expenditure), WindSense nodes are specifically engineered for pole-level installation, autonomous operation, and high-density deployment.

In summary, WindSense does not seek to replace institutional networks, but to fill a critical observational blind spot. Adding a scalable layer of high-granularity local sensing enables utilities to transition from reliance on regional proxies to incorporating infrastructure-representative measurements into their decision-making processes.

Acknowledgments

The authors thank CEFET-MG and Leuphana Universität Lüneburg for their collaboration and support.

References

[1] Agência Nacional de Energia Elétrica (ANEEL), "Aneel reforça medidas de resiliência e comunicação para o setor elétrico no verão 2025/2026," https://www.gov.br/aneel/pt-br/assuntos/noticias/2025/aneel-reforca-medidas-de-resiliencia-e-comunicacao-para-o-setor-eletrico-no-verao-2025-2026, 2025, acesso em: 2026-01-27.

[2] R. J. Campbell, "Weather-related power outages and electric system resiliency," Congressional Research Service, Tech. Rep. R42696, 2012, acesso em: 2026-01-27. [Online]. Available: https://digital.library.unt.edu/ark:/67531/metadc227867/

[3] NOAA NCEP, "Global forecast system (gfs) user guide / documentation," https://www.nco.ncep.noaa.gov/pmb/products/gfs/, 2026, acesso em: 2026-01-27.

[4] European Centre for Medium-Range Weather Forecasts (ECMWF), "Forecast user guide: The ecmwf forecasting system," https://www.ecmwf.int/en/forecasts/documentation-and-support/forecast-user-guide/forecasting-system, 2026, acesso em: 2026-01-27.

[5] B. Blocken, T. Stathopoulos, and J. P. A. J. van Beeck, "Pedestrian-level wind conditions around buildings: Review of wind-tunnel and cfd techniques and their accuracy for wind comfort assessment," *Building and Environment*, vol. 100, pp. 50–81, 2016.

[6] World Meteorological Organization (WMO), "Guide to instruments and methods of observation (wmo-no. 8)," https://library.wmo.int/, 2026, acesso em: 2026-01-27.

[7] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, "Long-range communications in unlicensed bands: The rising stars in the iot and smart city scenarios," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 4, pp. 3017–3041, 2016.

[8] F. Adelantado, X. Vilajosana, P. Tuset-Peiró, B. Martínez, J. Melià-Seguí, and T. Watteyne, "Understanding the limits of lorawan," *IEEE Communications Magazine*, vol. 55, no. 9, pp. 34–40, 2017.

[9] L. Atzori, A. Iera, and G. Morabito, "The internet of things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[10] J. Yick, B. Mukherjee, and D. Ghosal, "Wireless sensor network survey," *Computer Networks*, vol. 52, no. 12, pp. 2292–2330, 2008.

[11] ——, "Wireless sensor network survey," *Computer Networks*, vol. 52, no. 12, pp. 2292–2330, 2008.

[12] OASIS, "Mqtt version 5.0," OASIS Standard, 2019. [Online]. Available: https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.html

[13] T. Bray, "RFC 8259: The javascript object notation (JSON) data interchange format," RFC Editor, 2017. [Online]. Available: https://www.rfc-editor.org/rfc/rfc8259.html

[14] G. Klyne and C. Newman, "RFC 3339: Date and time on the internet: Timestamps," RFC Editor, 2002. [Online]. Available: https://www.rfc-editor.org/rfc/rfc3339.html

[15] Y. Shafranovich, "RFC 4180: Common format and MIME type for comma-separated values (CSV) files," RFC Editor, 2005. [Online]. Available: https://www.rfc-editor.org/rfc/rfc4180.html

[16] I. N. de Meteorologia (INMET), "Plataforma de mapas do inmet," https://mapas.inmet.gov.br/, Tech. Rep., 2026, acesso em: 2026-01-27.

[17] ——, "Bdmep – banco de dados meteorológicos para ensino e pesquisa," https://bdmep.inmet.gov.br/, Tech. Rep., 2026, acesso em: 2026-01-27.

[18] International Organization for Standardization, "Iso 8601-1:2019 – date and time – representations for information interchange," ISO, 2019, norma técnica (página oficial). [Online]. Available: https://www.iso.org/standard/70907.html

[19] M. Sokolova and G. Lapalme, "A systematic analysis of performance measures for classification tasks," *Information Processing & Management*, vol. 45, no. 4, pp. 427–437, 2009.

[20] K. Pearson, "Notes on regression and inheritance in the case of two parents," *Proceedings of the Royal Society of London*, vol. 58, pp. 240–242, 1895.

[21] C. J. Willmott and K. Matsuura, "Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance," *Climate Research*, vol. 30, no. 1, pp. 79–82, 2005.