This article has been accepted for publication in IEEE Internet of Things Journal. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JIOT.2026.3698589

Learning the Energy-Accuracy Frontier: Data-Driven Optimization in LoRa IoT Networks

Igor Bisio, *Senior Member, IEEE,* Chiara Garibotto, *Senior Member, IEEE,* Fabio Lavagetto, Andrea Sciarrone, *Senior Member, IEEE,* Matteo Zerbino, *Graduate Student Member, IEEE*

***Abstract*—Low-Power Wide-Area Network (LPWAN) technologies such as LoRa are key enablers of large-scale IoT monitoring systems, where long communication range and low energy consumption are essential. These advantages, however, come at the cost of strict throughput limitations, which significantly shape IoT system design. Compressive Sensing (CS) can mitigate this constraint by reducing transmitted data volumes, effectively trading communication load for additional processing at the sensing node and receiver. From an IoT perspective, this shift impacts node lifetime, hardware requirements, and overall network scalability. In this paper, we propose a data-driven optimization framework for LoRa-based IoT sensing systems employing CS. The approach jointly analyzes reconstruction quality and energy consumption through surrogate regression models that capture the interaction between physical-layer parameters and compression levels. This enables efficient multi-objective optimization via Pareto-front analysis and utopia-based selection. Results show that CS does not always dominate the quality-energy trade-off and that unified and stratified surrogate strategies identify closely aligned optimal operating points. Overall, the framework provides a practical and interpretable tool for the design of energy-efficient IoT sensing deployments.**

***Index Terms*—IoT, LoRa, Multiobjective Optimization, Compressive Sensing**

I. Introduction

A. Background and Motivation

When considering IoT deploying scenarios, it is not uncommon to deal with the LoRa communication paradigm. Its low power consumption and the ability to be independent of external infrastructure, such as the core network employed by 4G and 5G technologies, make it a really strong candidate when the amount of data to be transmitted is not prohibitively large. In the context of infrastructural monitoring, LoRa can be exceptionally useful, particularly in places and situations where a broadband connection is not easily achievable (for instance because providing powerline electricity to the structure is not feasible or there is no cellular coverage). Even though infrastructure monitoring usually implies dealing with a continuous data flow, it is possible to use Compressive Sensing (CS) techniques to reduce the amount of data to be sent, lessening the burden on nodes. This work aims at showing the effectiveness of an optimization framework which, starting from estimated quality and energy consumption measures, calculates regression data to enrich the measures themselves and chooses the ideal amount of data to be sent using CS techniques while maximizing the quality of the signal which can be extracted at the receiver side and, at the same time, minimizing the amount of energy expended by nodes.

From an edge-computing perspective, the proposed framework targets the constraints of battery-powered IoT nodes, where both computation and radio transmission must be carefully budgeted. Compressive sensing introduces a lightweight preprocessing step that reduces payload size prior to transmission, slightly increasing computational cost while significantly lowering transmission energy, which remains the dominant factor in low-power communication systems. The surrogate-based approach follows a practical separation between offline training and online inference. Model training and Pareto-front generation are performed offline using pre-collected datasets or periodic measurements, and can be executed on gateways or external computing platforms. At runtime, the edge node operates with minimal complexity: given an estimate of the channel condition, it selects the compression fraction (C_F) and transmission parameters (e.g., SF, C_R) from a small set of precomputed policies or lookup tables. This avoids embedding complex learning models on constrained devices while ensuring fast and predictable decision-making. This structure is also compatible with adaptation to real-world uncertainties, such as hardware variability, channel fluctuations, and interference, through measurement-driven dataset updates or periodic retraining. While these mechanisms were not explicitly validated in the present study, they represent a practical extension that could allow surrogate models to better reflect evolving operating conditions without increasing on-device complexity. Additionally, uncertainty-aware formulations could introduce safety margins to reduce sensitivity to idealized assumptions. Overall, the proposed methodology aligns with typical edge design principles: computationally intensive learning is performed offline, while deployed nodes rely on simple, interpretable, and energy-efficient decision mechanisms.

This work has been partially funded by the European Union - NextGenerationEU, within the framework of the project "RAISE – Robotics and AI for Socioeconomic Empowerment" and has been supported by European Union - NextGenerationEU. This research has been carried out in the framework of the "Mind the Bridge" project, funded by the program POR FESR Liguria 2021-2027. This work has also been supported by the Italian Ministry of Enterprises and Made in Italy (MIMIT) in the framework of the "Development and Cohesion Fund 2014-2020" as per CIPESS Resolution No. 9/2021 - through the House of Emerging Technologies of Genoa (CTE Genova) [CUP B37F23000000008].

The authors are with the DITEN Department, University of Genoa, Italy. C. Garibotto, F. Lavagetto and A. Sciarrone are also with the RAISE Ecosystem, Genoa, Italy. (emails: igor.bisio@unige.it, chiara.garibotto@unige.it, fabio.lavagetto@unige.it, andrea.sciarrone@unige.it, matteo.zerbino@edu.unige.it)

This article has been accepted for publication in IEEE Internet of Things Journal. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JIOT.2026.3698589

B. Contribution

Our contributions are summarized as follows:

- we extend the investigation initiated in our previous works [1], [2] by proposing a multi-objective optimization framework that provides a joint and comprehensive analysis of both compressive sensing and LoRa communications, explicitly addressing the trade-off between signal reconstruction quality and end-device energy consumption.
- we propose a data-driven surrogate modeling approach to approximate quality-energy trade-offs as a function of LoRa physical-layer parameters and compression levels, enabling efficient exploration of the design space without relying solely on analytical or heuristic models;
- we perform a Pareto-based multi-objective analysis with utopia-point selection, providing quantitative insights into when compressive sensing effectively improves energy efficiency and showing that compression is not always the dominant factor shaping optimal operating points;
- we introduce an edge-oriented deployment methodology that separates offline learning from lightweight online decision-making, making the proposed framework suitable for resource-constrained IoT nodes through simple policies or lookup tables rather than on-device learning models.

The remainder of this paper is organized as follows: Section II provides an overview of the related literature. Section III introduces the methodology and workflow followed for this paper. Section IV presents the performance results from the experimental tests. Finally, the last section concludes the paper and suggests directions for future research.

II. Related Works

Various works have investigated energy efficiency and performance optimization in LoRa and LoRaWAN systems. Some of them propose analytical and regression-based models to characterize the impact of transmission parameters, such as spreading factor, payload size, and transmit power, on device energy consumption and lifetime [3]. In parallel, multi-objective optimization frameworks have been introduced to jointly address energy efficiency, throughput, and fairness in LoRa networks, often relying on analytical models or heuristic solvers [4]. Data compression has also been explored in LoRa contexts, primarily through conventional lossless or transform-based techniques, with studies evaluating the resulting trade-offs between payload reduction and energy savings [5]. In addition, standardized mechanisms such as Static Context Header Compression (SCHC) have been proposed to improve protocol efficiency over LPWAN links, further highlighting the relevance of compression in low-power wide-area networks [6].

The application of compressive sensing (CS) in LPWAN environments has received comparatively limited attention. One notable example is the Nephalai framework [7], which applies CS at the physical layer to reduce data transfer between LPWAN gateways and cloud servers, targeting scalability and fronthaul efficiency rather than end-device energy-quality trade-offs. More recently, CS has been experimentally investigated for LoRa-based image transmission, demonstrating payload reduction capabilities but without a systematic analysis of energy consumption or multi-objective optimization [8]. Comprehensive surveys on LoRaWAN and LPWAN technologies further emphasize the importance of energy efficiency and parameter trade-offs, while broader studies on AI-driven wireless optimization further highlight the growing role of learning-based resource management across future communication systems, including secure and adaptive communication settings [9]–[12]. Despite these contributions, to the best of our knowledge no prior work jointly addresses compressive sensing, LoRa communication, and the multi-objective optimization of reconstruction quality and energy consumption using surrogate regression models and Pareto-based analysis, which constitutes the main focus of this study.

Compressive Sensing (CS) [13], [14] can potentially overcome the limitations of the Nyquist-Shannon sampling theorem by enabling accurate signal reconstruction from a reduced number of measurements. Its theoretical foundations are well established [15] and have been further developed through approaches that combine mathematical theory with engineering practice [16]. Despite the many demonstrated applications, CS still offers significant potential, particularly in structural health monitoring, where source-level data compression is essential for continuous, large-scale infrastructure monitoring. Due to its flexibility, CS has also been extended to the field of deep learning, leading to neural models that replace or complement traditional CS algorithms. In this context, [17] proposes a deep neural network for compressed signal reconstruction, while [18] introduces DCSNet to support Green IoT applications by reducing data transmission. A comparative analysis between classical CS methods and autoencoder-based approaches is presented in [19], highlighting differences in reconstruction accuracy and computational efficiency. The ability of CS to reduce transmitted data naturally aligns with LPWAN technologies such as LoRa, which are designed for energy-efficient, long-range IoT communications. Within this framework, [20] presents an energy-efficient scheduling strategy that integrates CS for data reconstruction.

Other works focus on dynamic transmission parameter optimization: [21] and [22] propose adaptive schemes based on signal characteristics (e.g. SNR, SINR, RSSI) to improve coverage and efficiency, while [23] analyzes energy consumption across different LoRa configurations and proposes parameter selection strategies. In [24], a tree-based LoRa architecture with self-configuring relay nodes is proposed to reduce network-wide energy consumption compared to the conventional star topology. A complementary line of research is presented in [25], which focuses on green LoRa networks powered by hybrid grid-renewable energy sources and formulates resource management and learning-based strategies to minimize gateway energy consumption. Finally, several studies specifically address energy efficiency in LoRa networks. Works such as [26] and [27] formulate nonconvex optimization problems considering user scheduling, spreading factor assignment, and power control, while [28] investigates the impact of frame size on energy consumption.

This article has been accepted for publication in IEEE Internet of Things Journal. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JIOT.2026.3698589

Overall, while prior studies have addressed LoRa parameter optimization, data compression techniques, and, more recently, the use of compressive sensing in LPWAN contexts, these aspects are typically treated in isolation or without a unified optimization perspective. In contrast, this work explicitly integrates compressive sensing and LoRa communication within a data-driven multi-objective framework, enabling a joint analysis of reconstruction quality and energy consumption. Furthermore, the use of surrogate models combined with Pareto-based optimization provides a systematic and computationally efficient way to explore the design space.

III. Methodology

A. Theoretical Foundation

We formulate the design of a LoRa-based compressive sensing system as a multi-objective optimization problem, where the goal is to jointly balance reconstruction quality and energy consumption. Each configuration is defined by a set of physical-layer parameters (*e.g.,* spreading factor and coding rate), channel conditions, and a compression fraction controlling the number of transmitted measurements. Because of this, we consider a decision vector $\overline{n}$, containing compression fraction, LoRa coding rate and LoRa spreading factor, alongside the SNR ξ.

$$\overline{n} = [C_F, C_R, SF] \tag{1}$$

Reconstruction quality is quantified through a scalar metric $Q(\overline{n}, \xi)$, defined as a weighted combination of two complementary components: the normalized root mean squared error (nRMSE) and the cosine similarity between the original and reconstructed signals. This choice allows capturing both amplitude fidelity and spectral alignment. The quality function is defined as

$$\begin{cases} Q(\overline{n}, \xi) = \beta \cdot Q_r(\overline{n}, \xi) + (1-\beta) \cdot Q_c(\overline{n}, \xi) \\ with\ Q_r(\overline{n}, \xi) = 1 - nRMSE(\overline{n}, \xi) \\ and\ Q_c(\overline{n}, \xi) = (cos_{sym}(\overline{n}, \xi) + 1)/2 \end{cases} \tag{2}$$

where $\beta \in [0, 1]$ controls the trade-off between the two terms. The nRMSE is inverted to obtain Q_r, while the cosine similarity is linearly remapped to $[0, 1]$ through Q_c, so that both components are expressed on a common scale and can be consistently combined. Energy consumption is modeled through a proxy $E_{pkt}(\overline{n}, \xi)$ that accounts for both transmission and local processing costs. The energy required for a single transmission attempt is computed as the sum of a transmission term k_{tx}, proportional to the LoRa Time-on-Air (ToA) derived from the standard Semtech formulation [29], and a compression cost k_{comp} associated with the execution of the compressive sensing algorithm:

$$E_{pkt}(\overline{n}, \xi) = k_{tx} \cdot ToA(\overline{n}) + k_{comp} \cdot C_F \tag{3}$$

When confirmed transmissions are considered, an additional reception cost is included to account for acknowledgment windows. To model the effect of channel conditions, the expected number of transmission attempts $E(\overline{n}, \xi)$ is estimated from a success probability function that depends on signal-to-noise ratio, spreading factor, coding rate, payload size p_s, and maximum retransmission attempts N_{max}. The resulting expected energy consumption can be approximated by:

$$E(\overline{n}, \xi) = E_{pkt}(\overline{n}, \xi) \cdot \frac{1-(1-p_s)^{N_{max}}}{p_s} \tag{4}$$

The resulting multi-objective optimization problem aims at minimizing energy consumption $E(\overline{n}, \xi)$ while maximizing reconstruction quality $Q(\overline{n}, \xi)$. The set of Pareto-optimal configurations is analyzed to characterize the trade-offs between compression, communication parameters, and energy efficiency, and representative operating points are selected using a distance-based criterion with respect to the utopia point. The optimization problem can be written as:

$$\begin{cases} \min_{\overline{n}}\ E(\overline{n}, \xi) \\ \max_{\overline{n}}\ Q(\overline{n}, \xi) \\ \text{s.t. } C_F \in [0, 1],\ C_R \in [\frac{4}{5}, \frac{4}{6}, \frac{4}{7}, \frac{4}{8}],\ SF \in [7, ..., 12] \end{cases} \tag{5}$$

Multi-objective optimization in this work arises naturally from the fundamental tension between communication/processing efficiency and reconstruction fidelity in resource constrained IoT nodes. In particular, increasing the compression strength and selecting low-overhead LoRa physical layer settings (*e.g.,* lower spreading factor and lower coding rate) reduces transmission time on air and therefore energy consumption, but may degrade the quality of the recovered signal at the receiver, especially under harsh noise conditions, while at the same time making retransmissions necessary, thus expending more energy. Conversely, conservative transmission settings and milder compression improve robustness and reconstruction accuracy while increasing energy cost. We then adopt a Pareto-based viewpoint in which each feasible configuration maps to a point in the objective space defined by energy and reconstruction quality. The Pareto front is then the set of non-dominated trade-offs: configurations for which no other configuration achieves both higher (or equal) quality and lower (or equal) energy simultaneously. This representation is particularly valuable in practical deployments because it quantifies how much quality must be sacrificed to gain a certain energy reduction (and vice versa), and it allows system designers to apply application-dependent preferences. In our pipeline, the "ground-truth" Pareto front is computed by exhaustive enumeration over the discrete dataset for each SNR level, providing a robust reference trade-off curve. Surrogate modeling extends this approach by learning continuous approximations of the objective surfaces, enabling dense sampling over compression ratios and joint parameter exploration beyond the original measurement grid. The resulting framework supports fair comparison across different surrogate regressors (structured convex surrogates, tree-based models, and probabilistic uncertainty-aware models) by referring to the same datasets, objective definitions, and Pareto extraction procedure, thereby isolating the impact of the regression family on the attainable trade-off sets.

This article has been accepted for publication in IEEE Internet of Things Journal. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JIOT.2026.3698589

B. Regression Models

In this paper, we evaluate three regression models under the same conditions. First, a convex Max-Affine (MA) surrogate, implemented as a log-sum-exp soft maximum of affine functions, is used to regress Q and E. This model enforces a structured functional form that promotes stability and mitigates spurious oscillations, although at the cost of limited expressiveness in highly non-convex settings. It can be interpreted as a simplified form of Input-Convex Neural Network (ICNN) [30]. Second, Random Forest (RF) regressors are considered as a strong nonparametric baseline. They effectively capture nonlinear interactions in small tabular datasets, but their piecewise-constant nature may lead to less smooth objective surfaces and small discontinuities in the resulting trade-off curves. Third, Gaussian Process Regression (GPR) is adopted as a probabilistic surrogate well suited to limited data regimes. It provides smooth predictions together with uncertainty estimates, enabling both a standard mean-based surrogate and an uncertainty-aware variant (GPR-UA), which penalizes configurations in regions of low confidence. In this work, the MA surrogate is primarily motivated by its convexity and simplicity. Convexity prevents spurious local extrema in the predicted objectives, leading to well-behaved Pareto fronts, while the shallow structure ensures robust training and low inference cost. These properties are particularly suitable for edge-oriented IoT scenarios, where the goal is to derive stable and lightweight parameter-selection policies rather than highly expressive models. More complex ICNN architectures were not considered, as they would introduce additional tuning complexity without clear benefits for the considered problem, where the input space is low-dimensional and the underlying relationships are largely smooth. Across all methods, we adopt the same evaluation protocol: for each SNR, a dense candidate set is generated, objective pairs are computed, Pareto fronts are extracted, and surrogate fronts are compared against the ground truth.

C. Convexity and Optima Search

In multi-objective optimization, convexity plays an important role when Pareto-optimal solutions are obtained through continuous optimization, as it guarantees well-behaved trade-offs and enables the use of scalarization techniques with strong optimality guarantees. However, in this work Pareto fronts are not computed via continuous optimization, but through exhaustive enumeration of a finite design space. For each SNR, all admissible parameter combinations are evaluated (or predicted on a dense grid), and Pareto-optimal points are identified via non-dominance filtering. In this setting, convexity is not required for correctness, as the Pareto front is obtained directly by comparison. Moreover, the design space is discrete and low-dimensional, making exhaustive search computationally feasible and eliminating the need for structural assumptions on the objective functions. The resulting Pareto front is therefore exact with respect to the evaluated configurations. Nonetheless, convex surrogates such as the max-affine model are included for their regularization properties. By enforcing a structured approximation, they produce smooth and well-behaved objective surfaces, typically leading to compact and stable Pareto fronts, although potentially at the cost of under-representing fine-grained trade-offs. In contrast, non-convex models such as Random Forest and Gaussian Process regression offer greater flexibility in capturing complex patterns, but may introduce irregularities in the predicted objective surfaces. This can result in fragmented or less stable Pareto fronts, which require careful interpretation and, in some cases, additional validation (e.g., comparison with the ground-truth front or uncertainty-aware selection). Overall, the combination of convex and non-convex surrogates allows balancing structural regularity and expressive power, providing a more comprehensive view of the energy-quality trade-off and of the reliability of the resulting Pareto-optimal configurations.

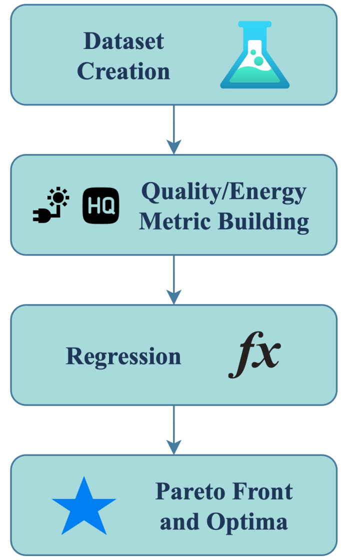

Fig. 1: A diagram showing the followed workflow.

D. Workflow

*1) **Dataset & Design Quantities**:* The objective of the work is to study how to jointly choose the compressive sensing rate and LoRa physical layer parameters while taking the channel quality into account so that an IoT node can minimize its energy consumption while maintaining as high a reconstruction quality as possible at the receiver. The design variables are the Compression Fraction (C_F), that is the fraction of original signal samples kept by compressive sensing, the LoRa Spreading Factor (SF), the LoRa Coding Rate (C_R), and the Signal-to-Noise Ratio (SNR) experienced on the channel. For each combination of these parameters a simulated dataset was built, starting from real world signals, with the help of the LoRaPHY tool [31], that contains, after the full transmit-receive-reconstruct chain, normalized RMSE and cosine similarity measures between the original and reconstructed signals. On top of that, two scalar indices, a quality index Q and an energy proxy E representing the approximated energy usage of the node, are built. The multi-objective goal is to maximize Q and minimize E, firstly done by an exhaustive grid search over the discrete dataset, and then by training surrogate models that approximate the relationship between the design variables and (Q, E). The surrogates allow further exploration of the parameter space, showing how far it is possible to push compression while keeping energy and quality under control.

The dataset is built as follows. There are several discrete SNR levels, from $-18dB$ to $-10dB$. The LoRa spreading factor SF ranges from 7 to 12, and the coding rate is represented as an integer C_R between 1 and 4, which corresponds to code rates $4/5$, $4/6$, $4/7$, $4/8$ in LoRa notation. The compression fraction C_F covers from 1% to 100% of the original samples, but in the final processed dataset we keep only multiples of 5% for simplicity. For each quadruple of SNR, SF, C_R and C_F, the simulation loops over multiple runs for statistical reliability: compressive sensing is applied to a reference signal at the given compression fraction, the compressed representation is transmitted over LoRa with the chosen SF and C_R, channel noise (AWGN) corresponding to the selected SNR is added, the receiver reconstructs the signal, and the normalized RMSE and cosine similarity with respect to the original signal are computed.

2) ***Quality & Energy Indices:*** From these raw metrics we derive a scalar quality index Q, as shown in Eq. 2. We first renormalize the RMSE across the whole dataset, even if it is already normalized at the signal level, so that the renormalized RMSE lies between zero and one over the dataset. We then map the cosine similarity, originally between minus one and plus one, into the interval from zero to one by a simple affine transformation, as seen in Eq. 2. Using these two ingredients we define two partial quality scores: one that rewards low RMSE, and one that is simply the mapped cosine similarity. The final quality Q is a combination of these two partial scores, controlled by a parameter β between zero and one. When $\beta = 0.5$, which is what we used for all shown experiments, the amplitude fidelity promoted by RMSE and the shape similarity encouraged by cosine similarity have equal weight. The outcome is a scalar Q between zero and one where higher is better and that blends information from both RMSE and cosine similarity. Considering both metrics at the same time allows for an improved evaluation process for the reconstructed signals, ensuring that both amplitude and shape information are correctly recovered.

While quality was quantified via simulation with the LoRaPHY tool, energy consumption was estimated thanks to a custom-made model, summarized in Eq. 3 and 4. This proxy is aimed at capturing, at least qualitatively, the energy consumption of a real LoRa node. The energy required for each transmission attempt is modeled as the sum of a transmission term, proportional to the LoRa time-on-air computed using the standard Semtech formulation, and a local processing term accounting for the computational cost of compressive sensing. In addition, when confirmed transmissions are considered, a fixed reception cost is included to model the energy spent during acknowledgment windows. To account for channel impairments, the expected number of transmission attempts is estimated from a link success probability that depends on signal-to-noise ratio, spreading factor, coding rate, and payload size. The total energy proxy is then obtained as the expected energy over all transmission attempts, and finally normalized across the dataset to yield values of E in the [0,1] range. This normalization is important because it puts Q and E on comparable scales for the multi-objective analysis.

3) ***Optimum Computation, Unified & Stratified Approaches:*** With Q and E defined, we first compute an optimum that does not involve any machine learning model, but is instead derived exhaustively from the dataset itself. First we find all non-dominated points when considering Q and E, then, for each quadruplet of SNR, SF, C_R, and C_F, we compute the Euclidean distance in the quality-energy plane to an utopia point defined as $Q = 1$ and $E = 0$. In other words, we measure how far each configuration is from the ideal situation of perfect quality and zero energy consumption. This table is our ground truth, based only on the discrete grid of compression fractions that we simulated.

To generalize beyond the discrete grid and to explore the solution space more densely, we introduce surrogate models based on Max-Affine Convex Neural Networks (MA), Random Forest regression (RF), and Gaussian Process Regression (GPR). Both a unified and a stratified approach were used for every regression method. In the unified setting, a single surrogate model is trained for each SNR level, jointly mapping the design parameters (spreading factor, coding rate, and compression fraction) to the quality and energy metrics. This approach captures global trends across the entire configuration space, at the cost of increased model complexity. In the stratified variant, instead, independent surrogate models are trained for each (SF,C_R) pair, with the compression fraction as the only continuous input variable. For each SNR, this results in multiple low-dimensional regressors, each modeling the relationship between compression and the corresponding quality and energy metrics for a fixed transmission configuration. The predicted points from all models are then merged into a single design cloud, from which the Pareto front and optimal operating points are extracted in the same manner as for the unified approach.

In order to improve the physical consistency and practical usefulness of the surrogate-based optimization, a refinement was introduced to the max-affine regression framework. Since the combined quality metric Q is explicitly defined in the range [0,1] by construction, unconstrained regression models may occasionally extrapolate beyond this interval, especially near the boundaries of the training data or when using smooth approximations of piecewise-linear functions. To mitigate this issue without overly restricting the expressiveness of the model, a soft penalty term was added to the training loss, discouraging predictions larger than one or smaller than zero while still allowing the optimizer to fit the data accurately.

Second, the resolution of the compression parameter during the optimization phase was decoupled from the resolution of the original dataset. While the dataset is discrete, due to the cost of generating and storing simulations at many compression ratios, the surrogate model enables continuous interpolation over the compression dimension. By evaluating the trained model on a much finer grid of compression values, it becomes possible to generate denser and smoother Pareto fronts and to identify trade-off points that lie between the discrete samples of the dataset. This refinement is particularly important for accurately locating knee points and utopia-distance optima, which may not coincide with the original sampling grid. This does not introduce artificial information, rather, it leverages the regression model to interpolate between known configurations under the assumption of smoothness, which is explicitly encoded in the surrogate architecture.

*4) **Observations**:* More broadly, these modifications highlight the fundamental motivation for using regression-based surrogate models instead of relying exclusively on the raw dataset. Direct optimization on the dataset restricts the analysis to a finite and coarse set of operating points, potentially missing better trade-offs that exist between sampled configurations. In contrast, surrogate models transform the problem into a continuous design space, enabling dense exploration, sensitivity analysis, and the incorporation of additional preferences such as compression bias or uncertainty awareness. This is especially relevant in system-level design problems, where the goal is not merely to select the best among pre-enumerated options, but to understand the structure of the trade-off surface and to derive robust parameter-selection policies that generalize beyond the exact conditions under which the dataset was generated. In this sense, regression is not used to replace the dataset, but to extract and exploit its underlying structure in a way that supports principled multi-objective optimization and edge-oriented deployment.

The feasibility and implications of both the unified and stratified modeling strategies depend strongly on the surrogate family, especially in the more complex unified case. Random forest regressors readily accommodate heterogeneous input features and nonlinear interactions, making them well suited for unified regression with minimal architectural changes. Gaussian process regression can also be extended to the unified setting and offers the additional benefit of uncertainty quantification, but its computational cost scales poorly with the number of training samples, which may limit its practicality for large datasets. In contrast, max-affine models introduce a more subtle trade-off: while they guarantee convexity and stability, enforcing convexity with respect to all input dimensions may be overly restrictive. If the surrogate is required to be convex simultaneously in compression, spreading factor, coding rate, and SNR, the model may oversmooth variations that are genuinely non-convex along some of these dimensions, effectively averaging distinct operating regimes and potentially degrading local accuracy. From a network-level perspective, the current formulation focuses on single-node optimization, where the trade-off between reconstruction quality and energy consumption is analyzed independently for each device. While this is appropriate for understanding fundamental system behavior, it does not explicitly capture multi-node interactions such as channel contention, interference, or coordinated resource allocation. Extending the proposed framework to multi-node scenarios would require incorporating shared communication constraints and possibly coupling the optimization across devices, which may impact both scalability and model complexity. Nevertheless, the surrogate-based approach provides a suitable foundation for such extensions, as it can be integrated with higher-level network optimization or scheduling mechanisms.

Altogether, we have three sets of optima for each combination of SNR, SF and C_R: the “real” optimum directly from the dataset grid, the optimum predicted by the stratified models, and the optimum predicted by the unified models, as well as results for single-objective optimization. For each of these we know the chosen compression fraction, quality, energy, and the distance to the utopia point. This allows us to quantify how close the surrogates are to the “true” Pareto-optimal configurations in terms of both parameters and objective values.

IV. RESULTS

The effectiveness of compressive sensing strongly depends on the underlying sparsity or compressibility of the signal, and this directly shapes the energy-quality trade-offs observed in the Pareto fronts. In this work we employed three main types of signals, a really sparse signal coming from accelerometric readings of an experimental structure built to showcase the effects of structural vibrations; a medium-sparsity signal from a publicly available dataset related to a railway bridge [32], [33], and a low-sparsity signal coming from a road bridge near Genoa, Italy. For highly sparse signals (*i.e.,* signals with a small number of significant coefficients in the chosen basis), reconstruction quality typically remains high over a wide range of aggressive compression ratios, meaning that Q degrades slowly as the compression fraction decreases. In this regime, the Pareto front tends to shift toward low energy solutions, because strong compression produces significant time on air savings while preserving acceptable fidelity. Because of this, surrogate-based optimization is expected to select stronger compression more often, particularly when combined with robust LoRa settings (high SF/C_R) under low SNR conditions. Such results are related to the experimental structure (shaker), as seen in Figure 2a, 2b and in the first part of Table I. For moderately or weakly sparse signals, such as those shown in Figures 2c to 2f as well as the latter part of Table I, the relationship becomes more sensitive: the region of compression ratios where reconstruction remains stable is narrower and Q begins to fall more noticeably once compression becomes too aggressive. In such a context, the Pareto front is biased toward higher energy solutions, because maintaining acceptable quality often requires transmitting more samples and/or choosing more robust LoRa configurations, especially at low SNR. Overall, across the three sparsity regimes one should expect not only different Pareto front locations but also different “stability” of the optimal policy: highly sparse signals admit wider safe operating regions (hence more consistent parameter recommendations), whereas weakly sparse signals lead to more brittle optima that may depend strongly on SNR.

Regarding the unified and stratified regression strategies, despite the conceptual differences between the two of them, the resulting Pareto fronts and selected optimal configurations are remarkably similar across the examined SNR levels and surrogate models. Both unified and stratified approaches consistently identify comparable trade-offs between energy consumption and reconstruction quality, and lead to closely

This article has been accepted for publication in IEEE Internet of Things Journal. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JIOT.2026.3698589

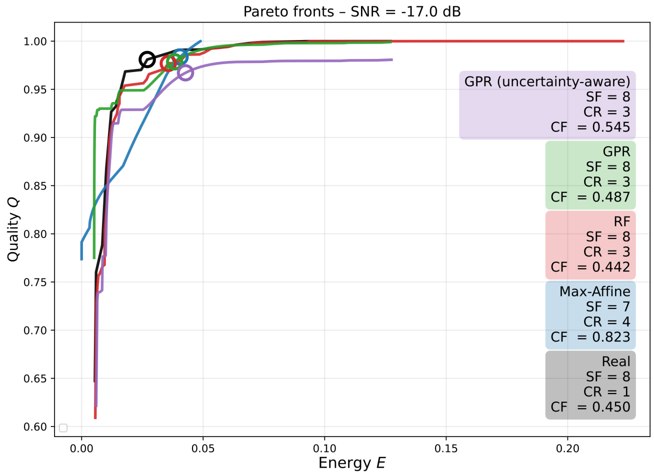

(a) Experimental Structure, -17dB.

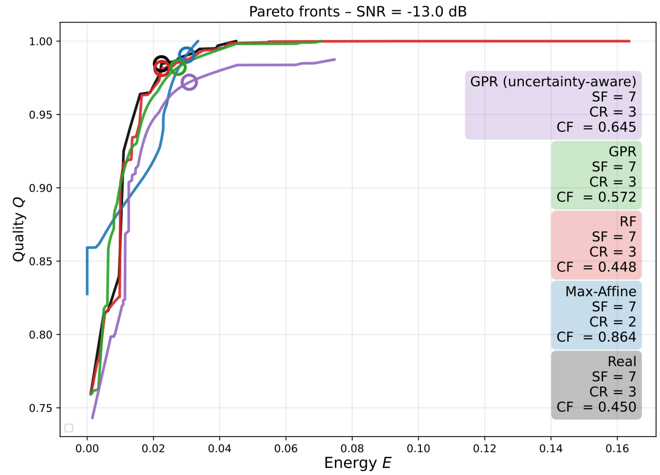

(b) Experimental Structure, -13dB.

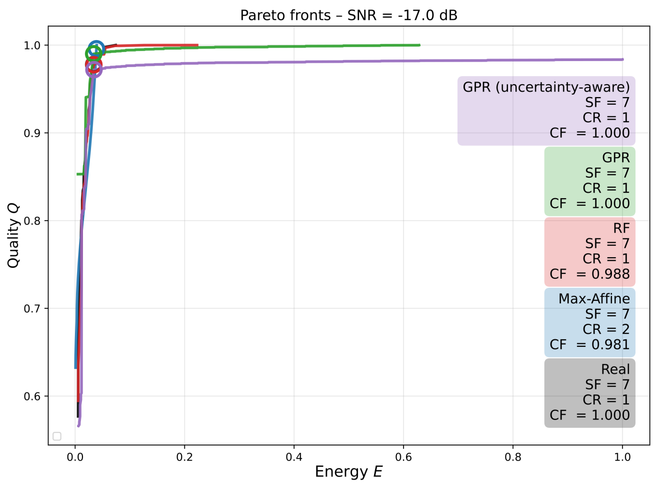

(c) Road Bridge, -17dB.

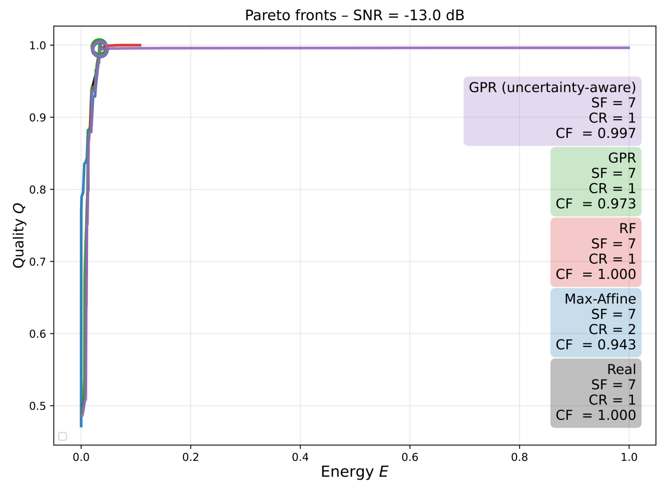

(d) Road Bridge, -13dB.

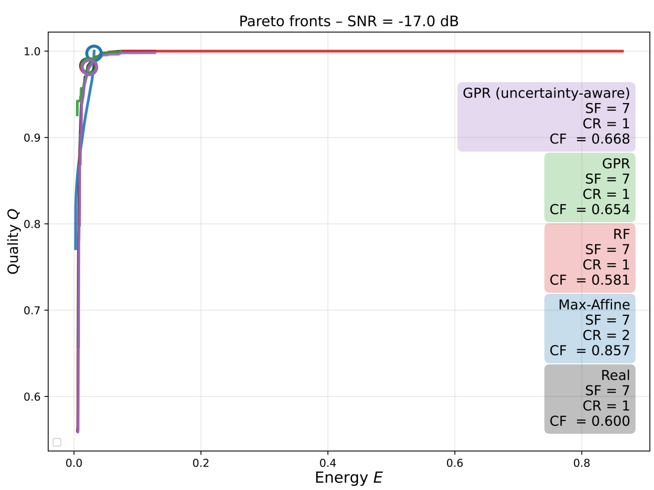

(e) Railway Bridge, -17dB.

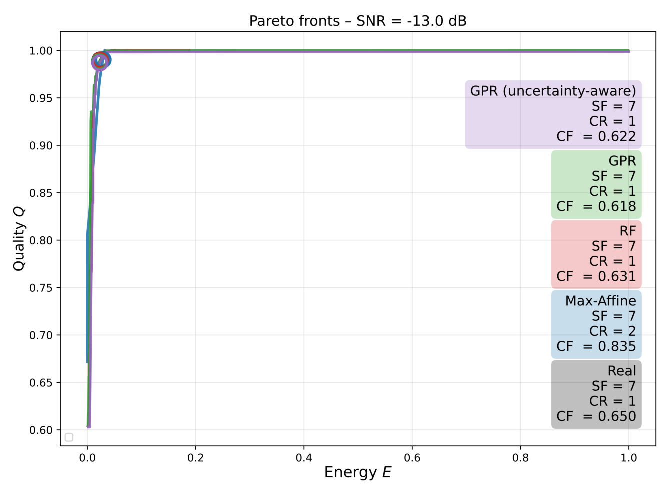

(f) Railway Bridge, -13dB.

Fig. 2: Experimental results in the unified setting.

aligned compression ratios and transmission parameter selections. Because of this, for the sake of conciseness, both Tables I and II as well as Figures 2 and 3 are related to results coming only from the unified approach. This behavior can be explained by the structured nature of the underlying dataset. For a fixed SNR, the dependence of both quality and energy on the compression fraction is relatively smooth and monotonic for each (SF,C_R) combination. As a result,

This article has been accepted for publication in IEEE Internet of Things Journal. This is the author's version which has not been fully edited and
content may change prior to final publication. Citation information: DOI 10.1109/JIOT.2026.3698589

TABLE I: Methods comparison and attained results for selected configurations in the unified experimental setting.

Experimental Structure, -17dB						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	0.9809	0.0270	0.450	1	8	0.0331
MA	0.9836	0.0405	0.823	4	7	0.0437
RF	0.9769	0.0357	0.442	3	8	0.0425
GPR	0.9786	0.0383	0.487	3	8	0.0439
GPR_UC	0.9671	0.0427	0.545	3	8	0.0539
Experimental Structure, -13dB						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	0.9845	0.0224	0.450	3	7	0.0272
MA	0.9906	0.0300	0.864	2	7	0.0315
RF	0.9814	0.0226	0.448	3	7	0.0292
GPR	0.9822	0.0274	0.572	3	7	0.0327
GPR_UC	0.9720	0.0308	0.645	3	7	0.0416
Road Bridge, -17dB						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	0.9779	0.0339	1.000	1	7	0.0405
MA	0.9958	0.0389	0.981	2	7	0.0391
RF	0.9777	0.0341	0.988	1	7	0.0408
GPR	0.9901	0.0339	1.000	1	7	0.0353
GPR_UC	0.9720	0.0344	1.000	1	7	0.0443
Road Bridge, -13dB						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	0.9943	0.0339	1.000	1	7	0.0344
MA	0.9938	0.0347	0.943	2	7	0.0353
RF	0.9952	0.0342	1.000	1	7	0.0345
GPR	0.9975	0.0331	0.973	1	7	0.0332
GPR_UC	0.9946	0.0343	0.997	1	7	0.0347
Railway Bridge, -17dB						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	0.9829	0.0220	0.600	1	7	0.0279
MA	0.9974	0.0315	0.857	2	7	0.0316
RF	0.9822	0.0225	0.581	1	7	0.0287
GPR	0.9830	0.0236	0.654	1	7	0.0291
GPR_UC	0.9813	0.0244	0.668	1	7	0.0308
Railway Bridge, -13dB						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	0.9898	0.0234	0.650	1	7	0.0256
MA	0.9902	0.0287	0.835	2	7	0.0304
RF	0.9896	0.0239	0.631	1	7	0.0260
GPR	0.9881	0.0225	0.618	1	7	0.0254
GPR_UC	0.9870	0.0230	0.622	1	7	0.0264

TABLE II: Comparison of attained results when performing multi-objective optimization versus optimizing only for Energy or Quality, unified setting.

Multiobjective Optimization						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	0.9809	0.0270	0.450	1	8	0.0331
MA	0.9836	0.0405	0.823	4	7	0.0437
RF	0.9769	0.0357	0.442	3	8	0.0425
GPR	0.9786	0.0383	0.487	3	8	0.0439
GPR_UC	0.9671	0.0427	0.545	3	8	0.0539
Q Max						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	1.0000	0.1271	1.000	3	9	0.1271
MA	1.0000	0.0489	0.982	4	7	0.0489
RF	1.0000	0.2228	0.996	3	10	0.2228
GPR	0.9991	0.1271	1.000	3	9	0.1271
GPR_UC	0.9807	0.1275	1.000	3	9	0.1289
E Min						
Method	**Q**	**E**	C_F	C_R	S_F	d_{utopia}
Dataset	0.6473	0.0055	0.050	1	7	0.3528
MA	0.7738	0.0000	0.084	3	9	0.2262
RF	0.6090	0.0056	0.050	1	7	0.3910
GPR	0.7751	0.0052	0.055	1	7	0.2249
GPR_UC	0.6216	0.0060	0.050	1	7	0.3785

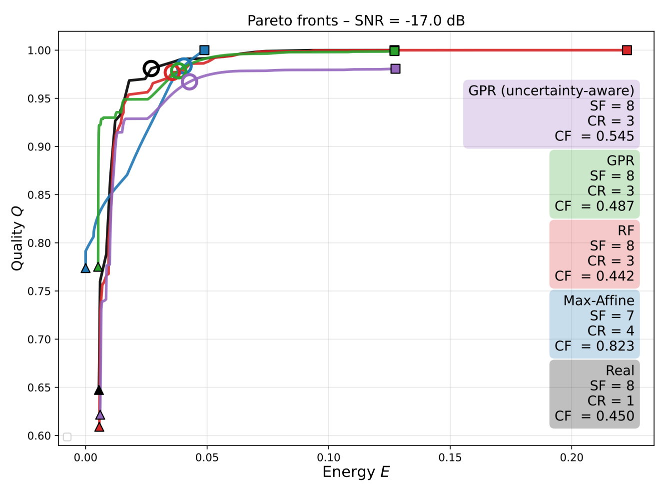

Fig. 3: Multi-objective (circles), Energy-only (triangles), and Quality-only (squares) optima, unified setting, experimental structure, -17dB.

the unified surrogate is able to effectively capture the same relationships learned by the ensemble of stratified models, while implicitly leveraging shared structure across different configurations. From a practical standpoint, these results suggest that the additional complexity introduced by stratified surrogate modeling does not provide a significant advantage in the considered scenario. Given the higher implementation overhead and reduced scalability of the stratified approach, the unified surrogate formulation represents a more convenient and equally effective choice for the optimization pipeline presented in this work.

The obtained results also show that, in the considered scenarios, the role of compression is not as dominant as might be expected. In particular, increasing the compression level does not necessarily lead to a substantial improvement in the trade-off between reconstruction accuracy and energy consumption. In some cases, stronger compression can be applied without incurring a significant performance loss, while in others the associated energy savings remain marginal. This behavior indicates that compression is not always a critical factor in determining the optimal operating point. More generally, the findings suggest that compression's effectiveness depends on the balance between communication costs and local processing overhead, and must therefore be evaluated jointly with the other system parameters rather than optimized in isolation. From a broader perspective, different conclusions may be drawn if the optimization target is shifted from purely system-level metrics to application-oriented quality-of-experience (QoE) requirements. For instance, when specific signal characteristics or frequency components are deemed more relevant, alternative operating points along the Pareto front may become preferable. In such cases, the notion of optimality can be redefined, either by adopting different selection criteria (*e.g.,* weighted distances or non-uniform preference structures instead of the minimum-distance-to-utopia rule) or by directly modifying the underlying performance metrics, such as reconstruction quality and energy consumption. This flexibility enables the system to prioritize perceptual or task-specific objectives over global efficiency, leading to different compression and transmission strategies better aligned with the desired QoE.

When analyzing the single-objective optimization results, two consistent and expected behaviors emerge. When the objective is to maximize reconstruction quality, the optimal solution almost invariably corresponds to the absence of compression. Conversely, when the objective is to minimize energy consumption, the optimal solution consistently favors the maximum admissible compression. These outcomes are in full agreement with the initial assumptions of this paper: compression inevitably degrades reconstruction fidelity, while reducing the payload size directly lowers the transmission energy. The fact that the proposed framework reproduces these intuitive extremes further validates the strategies applied for both the dataset and the optimization pipeline.

A more subtle behavior is observed for the physical-layer parameters, namely SF and C_R. While it is true that the general behaviour is to raise LoRa parameters for quality maximization and to lower them for energy minimization, maximizing quality does not systematically lead to the selection of the most robust physical-layer configuration (*i.e.,* highest SF and C_R), nor does minimizing energy always result in the lowest SF and C_R. Instead, intermediate or non-extreme values of these parameters can be selected, depending on the signal-to-noise ratio (SNR) and the optimization criterion. The apparent discrepancy arises from the specific definition of the quality metric adopted in this work. Reconstruction quality is evaluated through signal-level metrics, such as normalized RMSE and cosine similarity, which quantify the fidelity of the recovered signal conditional on successful reception. These metrics do not explicitly account for packet loss or decoding failures. In other words, increasing SF or C_R does not improve the quality of the reconstructed signal itself, but rather increases the probability that the packet is successfully and correctly delivered. If the transmission is already reliable at a given SNR, additional robustness at the physical layer does not translate into higher quality under the adopted definition.

Consequently, in the quality maximization problem, the optimizer is effectively free to select any combination of SF and C_R that ensures successful delivery, as long as the compression level is minimal. This explains why configurations with non-maximum SF and C_R may be selected even when the objective is purely to maximize quality. From the optimizer's perspective, increasing robustness beyond what is strictly necessary does not provide any benefit in terms of the objective function. A similar reasoning applies, in a more nuanced form, to the energy minimization problem. While lower SF and C_R generally reduce the time-on-air of a single transmission, they also reduce transmission robustness. When confirmed uplink transmissions and acknowledgments are considered, this loss of robustness may increase the expected number of retransmissions. As a result, the expected energy consumption increases. A configuration with slightly higher SF or C_R may incur a higher cost per transmission attempt, yet result in a lower expected energy consumption due to a reduced probability of retransmission. This trade-off naturally leads to optimal solutions characterized by intermediate values of SF and C_R, particularly at low or moderate SNRs. Such configurations strike a balance between transmission duration and reliability, minimizing the expected energy expenditure rather than the energy of a single transmission attempt.

The absence of extreme SF and C_R values in the single-objective optima is not an artifact of the surrogate models or the optimization procedure. Instead, it is a direct consequence of the adopted objective definitions and the explicit treatment of SNR and retransmissions. Had the quality metric incorporated packet loss (*e.g.,* by assigning zero quality to failed transmissions or by weighting reconstruction quality by the probability of successful delivery) the optimizer would naturally favor more robust physical-layer configurations when maximizing quality. Similarly, if energy were defined solely in terms of time-on-air without accounting for retransmissions, the minimum-energy solution would trivially correspond to the lowest SF and C_R.

V. Conclusion and Future Developments

This work presented a data-driven optimization framework for jointly analyzing reconstruction quality and energy consumption in LoRa-based sensing systems employing compressive sensing. Surrogate regression models were used to approximate the underlying trade-offs across physical-layer parameters and compression levels, enabling an efficient multi-objective analysis via Pareto fronts and utopia-based selection. The results show that, within the considered scenarios, compression does not always play a dominant role in shaping the quality-energy trade-off, and that unified and stratified surrogate modeling strategies lead to closely aligned optimal operating points. Overall, the proposed framework provides a

This article has been accepted for publication in IEEE Internet of Things Journal. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JIOT.2026.3698589

flexible and interpretable tool for exploring design trade-offs in low-power wireless sensing systems.

Future work may extend the proposed framework by exploring alternative compressive sensing design choices, including different sensing matrices and reconstruction strategies, which could improve reconstruction quality at higher compression levels and potentially shift the Pareto-optimal operating points. In addition, the definition of reconstruction quality could be refined by adopting alternative or perception-oriented metrics, or by adjusting the relative weighting of existing measures to better reflect application-specific priorities, such as accurate frequency localization, while also taking into account successful delivery of information. Such a formulation would bridge the gap between mathematical Pareto optimality and experience-oriented system design, and may lead to more robust operating points in challenging channel conditions. Another important direction concerns experimental validation through real-world deployments. Implementing the proposed framework on a physical sensing node would enable a more comprehensive assessment under realistic operating conditions, including hardware constraints and channel impairments, and would provide empirical validation of the energy consumption models adopted in this study.

References

[1] I. Bisio, C. Garibotto, F. Lavagetto, A. Sciarrone, and M. Zerbino, "Distributed multiobjective optimization for edge computing in resource-constrained social iot networks," *IEEE Internet of Things Journal*, vol. 13, no. 8, pp. 15 549–15 558, 2026.

[2] ——, "Optimized resource orchestration for cs-enabled shm in lora networks," *IEEE Networking Letters*, pp. 1–1, 2026.

[3] J. Lorincz, M. Kusačić, E. Čusto, and Z. Blažević, "A comprehensive multiple linear regression modeling and analysis of lora user device energy consumption," *Journal of Sensor and Actuator Networks*, vol. 15, no. 1, 2026. [Online]. Available: https://www.mdpi.com/2224-2708/15/1/5

[4] H. Zhang, Y. Song, M. Yang, and Q. Jia, "Modeling and optimization of lora networks under multiple constraints," *Sensors*, vol. 23, no. 18, 2023. [Online]. Available: https://www.mdpi.com/1424-8220/23/18/7783

[5] "Data compression in lora networks: A compromise between performance and energy consumption," vol. 14. [Online]. Available: https://journals-sol.sbc.org.br/index.php/jisa/article/view/3000

[6] A. Minaburo, L. Toutain, and R. Andreasen, "Static Context Header Compression (SCHC) for the Constrained Application Protocol (CoAP)," RFC 8824, Jun. 2021. [Online]. Available: https://www.rfc-editor.org/info/rfc8824

[7] J. Liu, W. Xu, S. Jha, and W. Hu, "Nephalai: towards lpwan c-ran with physical layer compression," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '20. New York, NY, USA: Association for Computing Machinery, 2020. [Online]. Available: https://doi.org/10.1145/3372224.3419193

[8] F. Chaparro B., M. Pérez, and D. Mendez, "A communication framework for image transmission through lpwan technology," *Electronics*, vol. 11, no. 11, 2022. [Online]. Available: https://www.mdpi.com/2079-9292/11/11/1764

[9] J. Haxhibeqiri, E. De Poorter, I. Moerman, and J. Hoebeke, "A survey of lorawan for iot: From technology to application," *Sensors*, vol. 18, no. 11, 2018. [Online]. Available: https://www.mdpi.com/1424-8220/18/11/3995

[10] M. Haldi Widianto, A. Ramadhan, A. Trisetyarso, and E. Abdurachman, "Energy saving on iot using lora: a systematic literature review," *International Journal of Reconfigurable and Embedded Systems (IJRES)*, vol. 11, pp. 25–33, 03 2022.

[11] B. Mao, F. Tang, Y. Kawamoto, and N. Kato, "Ai models for green communications towards 6g," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 1, pp. 210–247, 2022.

[12] B. Mao, Y. Wu, J. Liu, H. Guo, J. Wang, and N. Kato, "Optimizing secrecy rate for federated learning model aggregation with intelligent reflecting surface toward 6g ubiquitous intelligence," *IEEE Transactions on Cognitive Communications and Networking*, vol. 11, no. 2, pp. 1258–1267, 2025.

[13] E. J. Candes and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.

[14] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[15] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Birkhäuser New York, NY, 2013.

[16] S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019.

[17] Y. An, Z. Xue, and J. Ou, "Deep learning-based sparsity-free compressive sensing method for high accuracy structural vibration response reconstruction," *Mechanical Systems and Signal Processing*, vol. 211, p. 111168, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0888327024000669

[18] M. Zhang, H. Zhang, D. Yuan, and M. Zhang, "Compressive sensing and autoencoder based compressed data aggregation for green iot networks," in *2019 IEEE Global Communications Conference (GLOBECOM)*, 2019, pp. 1–6.

[19] J.-Y. Jeong, M. Ozger, and W.-H. Lee, "Compressed sensing vs. auto-encoder: On the perspective of signal compression and restoration," *IEEE Access*, vol. 12, pp. 41 967–41 979, 2024.

[20] Y. Wu, Y. He, and L. Shi, "Energy-saving measurement in lorawan-based wireless sensor networks by using compressed sensing," *IEEE Access*, vol. 8, pp. 49 477–49 486, 2020.

[21] H. Haif, A. Arous, and H. Arslan, "A decentralized dynamic relaying-based framework for enhancing lora networks performance," *IEEE Internet of Things Journal*, vol. 11, no. 12, pp. 22 183–22 196, 2024.

[22] I. P. Manciu, s. V. Nadoleanu, D. I. Săcăleanu, L. A. Perisoară, and R. C. Constantinescu, "Dynamic adaptation algorithm of lora communication parameters in wireless sensor networks," in *2024 IEEE 30th International Symposium for Design and Technology in Electronic Packaging (SIITME)*, 2024, pp. 286–290.

[23] R. S. Benatti, C. P. de Souza, and O. Baiocchi, "An optimization method based on lora parameters for energy consumption reduction," in *2021 5th International Symposium on Instrumentation Systems, Circuits and Transducers (INSCIT)*, 2021, pp. 1–5.

[24] Y. H. Tehrani, A. Amini, and S. M. Atarodi, "A tree-structured lora network for energy efficiency," *IEEE Internet of Things Journal*, vol. 8, no. 7, pp. 6002–6011, 2021.

[25] R. Hamdi, E. Baccour, A. Erbad, M. Qaraqe, and M. Hamdi, "Lora-rl: Deep reinforcement learning for resource management in hybrid energy lora wireless networks," *IEEE Internet of Things Journal*, vol. 9, no. 9, pp. 6458–6476, 2022.

[26] B. Su, Z. Qin, and Q. Ni, "Energy efficient resource allocation for uplink lora networks," in *2018 IEEE Global Communications Conference (GLOBECOM)*, 2018, pp. 1–7.

[27] ——, "Energy efficient uplink transmissions in lora networks," *IEEE Transactions on Communications*, vol. 68, no. 8, pp. 4960–4972, 2020.

[28] Adnan, M. Rizal, and A. A. Ilham, "Performance of lora gateway based energy consumption and different frame sizes," in *2018 2nd East Indonesia Conference on Computer and Information Technology (EIConCIT)*, 2018, pp. 159–162.

[29] *High-speed CAN transceiver*, SX1276/77/78/79 - 137 MHz to 1020 MHz Low Power Long Range Transceiver, 5 2020, rev. 7. [Online]. Available: https://semtech.my.salesforce.com/sfc/p/#E0000000JelG/a/2R0000001Rbr/6EfVZUorrpoKFfvaF_Fkpgp5kzjiNyiAbqcpqh9qSjE

[30] B. Amos, L. Xu, and J. Z. Kolter, "Input convex neural networks," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 2017, pp. 146–155.

[31] "LoRaPHY," https://github.com/jkadbear/LoRaPHY, accessed: 2025-02-22.

[32] K. Maes and G. Lombaert, "Monitoring railway bridge kw51 before, during, and after retrofitting," *Journal of Bridge Engineering*, vol. 26, no. 3, p. 04721001, 2021. [Online]. Available: https://ascelibrary.org/doi/abs/10.1061/%28ASCE%29BE.1943-5592.0001668

[33] ——, "Monitoring data for railway bridge kw51 in leuven, belgium, before, during, and after retrofitting," Apr. 2020. [Online]. Available: https://doi.org/10.5281/zenodo.3745914