RF Super Resolution: A Deep Learning Approach to Spatial Enhancement for LoRa

Andreas Kuster
College of Computing and Data Science
Nanyang Technological University
Singapore
S220003@e.ntu.edu.sg

Huatao Xu
Computer Science and Engineering
Hong Kong University of Science and Technology
Hong Kong
huatao@cse.ust.hk

Rui Tan
College of Computing and Data Science
Nanyang Technological University
Singapore
tanrui@ntu.edu.sg

Mo Li
Computer Science and Engineering
Hong Kong University of Science and Technology
Hong Kong
lim@cse.ust.hk

Abstract

The analog-to-digital converter (ADC) in a radio frequency (RF) front-end and digital signal processing (DSP) are significant sources of energy consumption, particularly in low-power systems like LoRa. To reliably demodulate weak signals, current systems rely on heavy oversampling - often 8x the signal bandwidth - which imposes a substantial and persistent oversampling tax on the analog front-end and DSP. This paper investigates if this tax can be mitigated by adapting techniques from image and video super resolution. We propose *RF Super Resolution*, a lightweight, real-time neural upscaler for RF signals. Our approach pairs an efficient digital interpolation algorithm with a shallow four-layer CNN. The neural network is trained to learn and correct the residual artifacts introduced by the digital upsampling and noise, effectively mimicking the output of a high-rate analog ADC and denoising filter. We validate our system on a large-scale, over-the-air LoRa study. Our results show that *RF-SR*, given a 2× Nyquist input (250 kHz), restores demodulation performance of a native 8× oversampled (2 MHz) system at half its sampling rate (1 MHz). This effectively removes the analog oversampling requirement, and provides an additional 1.25 dB SNR gain over the oversampled baseline, making it an efficient and effective signal enhancer suitable for gateway integration or post training quantized deployment at the end-node.

CCS Concepts

• **Hardware → Digital signal processing**; **Wireless devices**; • **Computer systems organization → Neural networks**; **Sensor networks**; • **Networks → Wireless access networks**.

Keywords

LoRa, LPWAN, Super Resolution, Machine Learning for Wireless Systems, Signal Denoising, Deep Learning Super Sampling (DLSS)

ACM Reference Format:
Andreas Kuster, Huatao Xu, Rui Tan, and Mo Li. 2026. RF Super Resolution: A Deep Learning Approach to Spatial Enhancement for LoRa. In *The 24th Annual International Conference on Mobile Systems, Applications and Services (MobiSys '26), June 21–25, 2026, Cambridge, United Kingdom.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3745756.3809216

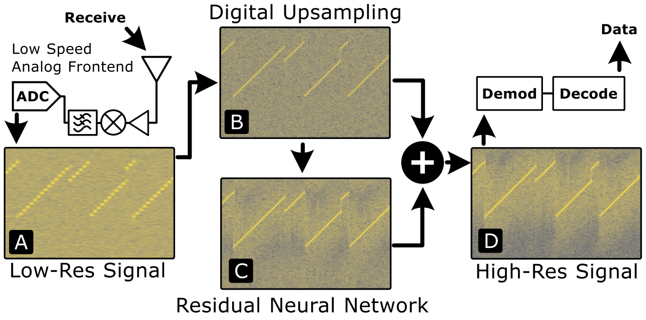

Figure 1: High-level overview of RF-SR. A low-speed, low-cost, low-energy analog frontend (A) is sufficient because the super-resolution upsampler digitally handles both the increase in sampling rate (B) and de-noising (C). The enhanced signal (D) then feeds the standard demodulation and decoding pipeline.

1 Introduction

The motivation for this research originates from visual computing, particularly the domain of image and video super-resolution[14, 27, 32]. Real-time upscaling techniques, such as NVIDIA's Deep Learning Super Sampling (DLSS) and RTX Video Super-Resolution (VSR) [11], have demonstrated two principal benefits for resource-constrained systems: **(1) Quality Enhancement:** A low-resolution, blurry, or artifact-laden image or video stream can be intelligently upscaled to high perceptual quality; and **(2) Compute Reduction:** By rendering demanding tasks (like ray tracing) at low resolution and neurally upscale the output to high resolution, users get a high resolution experience at a fraction of the cost. This *render low, upscale high* paradigm has transformed real-time graphics by

MobiSys '26, Cambridge, United Kingdom

ACM ISBN 979-8-4007-2027-7/26/06
https://doi.org/10.1145/3745756.3809216

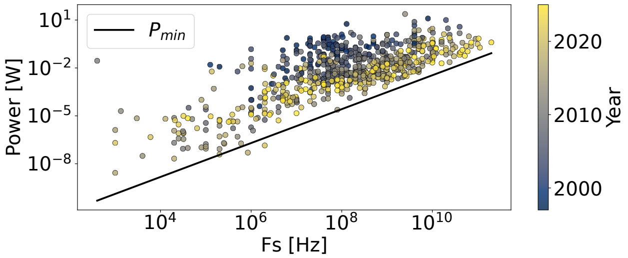

Figure 2: ADC power vs sampling rate scaling.

decoupling output quality from prohibitively expensive computation. Inspired by this success, we ask: **Can a similar principle be applied for radio frequency (RF) signals such as LoRa?** In the RF context, super-resolution can be translated along two axis: **(1) Spatial SR:** Increasing the sampling rate, analogous to increasing the pixel density / resolution; and **(2) Spectral SR:** Increasing the ADC bit resolution, analogous to increasing pixel colour depth.

For our investigation, we target LoRa (Long Range), a low-power wide-area network (LPWAN) technology that leverages Chirp Spread Spectrum (CSS) modulation to achieve robust communication in high-noise, low SNR environments. LoRa's design and operating regime make it an attractive yet challenging platform for neural signal enhancement, due to its noisy and resource-constrained application scenarios.

A key limitation in existing research on neural-enhanced LoRa particularly in the contexts of collision resolution and physical-layer signal improvement is a critical dependence on heavy oversampling. As detailed in Table 1, it is typical to sample signals at 1–2 MHz for channel bandwidths of only 125–250 kHz, resulting in oversampling factors ranging from 4× to 8×. This "oversampling tax" imposes substantial and continuous power burdens throughout the analog frontend, through the ADC to the subsequent data transfer and digital signal processing stages (e.g., FFT has $O(n \cdot log(n)$ and correlation $O(n^2)$ runtime complexity). A large-scale survey of ADCs [35] in Figure 2 shows that power consumption scales superlinearly with sampling rate (F_s), where the minimum power consumption trendline is given by: $P_{\text{min}} = 1.4 \text{ nW} \cdot F_s^{1.06}$ making F_s the primary driver of ADC power.

In addition, prior signal processing-based (e.g., LoRaTrimmer [16]) or deep learning–based approaches to LoRa signal enhancement (e.g., NELoRa [30], GLoRiPHY [41]) are characterized by considerable power and resource requirements. These works demonstrate notable symbol-level SNR improvements (1–3 dB); however, they require computational platforms such as discrete GPUs, with energy and memory consumption orders of magnitude beyond that of commercial LoRa end-nodes (Table 2).

While some works, such as LoRaSeek [36], reduce the memory footprint to accommodate lower-power platforms (e.g., 21.8 MB on a Raspberry Pi), this remains three orders of magnitude higher than our work. Furthermore, LoRaSeek still requires an NVIDIA RTX 3090 for real-time processing; a Raspberry Pi remains 16.5× too slow (12.6× in compressed mode) to achieve real-time performance. Lastly, given the high reliability required by deployment in non-easily serviceable scenarios, the inherent susceptibility of high-performance GPUs to various software and hardware faults (e.g., bit errors, bus instability) [13, 29, 38] compromises their viability for this use case.

Paper	**Fs** [MHz]	**OSF**	**Dataset + SNR**
LoRaPHY [57]	1	8x	OTA[2], no info[3]
SDR-LoRa [7]	-[1]	-	Emulation
LoRa Collision Resolving			
CIC [46]	2	8x	OTA[2], no info[3]
TnB [40]	1	8x	OTA[2], no info[3]
Pyramid [58]	1	8x	OTA[2], w/ AWGN
CoLoRa [50]	-[1]	-	OTA[2], w/ AWGN
Choir [17]	-[1]	-	OTA[2], no info[3]
PCube [54]	-[1]	-	OTA[2], tx pwr + loc
NScale [49]	1	8x	OTA[2], w/ AWGN
CurvingLoRa [31]	1	8x	OTA[2], w/ AWGN
FlipLoRa [56]	1	8x	OTA[2], tx pwr + loc
FTrack [55]	1	4x	OTA[2], tx power
AlignTrack [10]	1	8x	OTA[2], w/ AWGN
OpenLoRa [34]	-[1]	-	OTA[2], tx pwr + loc
mLoRa [52]	2	-	No SNR experiment
LoRa Signal Enhancing			
XCopy [53]	1	4x	OTA[2], location
NELoRa [30]	1	8x	OTA[2], w/ AWGN
LoRaTrimmer [16]	1	8x	OTA[2], w/ AWGN
GLoRiPHY [41]	1	8x	OTA[2], w/ AWGN
DeepSense [9]	1	8,4,2x	OTA[2], w/ AWGN
LoRaSeek [36]	1	8,4,2x	OTA[2], w/ AWGN
LoRa Upsampler			
RF-SR (our work)	0.250	2x	OTA[2], tx pwr + loc

Table 1: Survey of sampling parameters in LoRa research. [1]Not reported. [2]Over-the-air sample collection. [3]SNR variation method not reported

To fill the gap, we propose *RF Super Resolution* (RF-SR), a deep learning–based upsampler and filter for real-time spatial super-resolution RF signal, illustrated in Figure 1. The core idea is to allow the analog front-end and ADC to operate at substantially reduced sampling rates; the resulting low-rate digital signal is then upsampled and processed through a lightweight neural network designed to correct interpolation artifacts and suppress noise. This enables delivery of a high-quality, denoised digital signal to standard demodulation and decoding pipelines without incurring the traditional cost of oversampling.

To achieve this goal, a key difficulty lies in denoising and reconstructing LoRa waveforms from undersampled and noisy signals without introducing artifacts or losing crucial signal features. To overcome this, we design a compact four layer convolutional neural network (CNN) that is both computationally efficient and effective. Specifically, our architecture incorporates complex-valued convolutions to handle the in-phase and quadrature (IQ) components of RF signals to capturing their intrinsic characteristics. Additionally, we introduce a hybrid denoiser loss function, which jointly considers time-domain reconstruction accuracy and frequency-domain denoising objectives. This combined loss ensures robust performance across varying channel and noise conditions, enabling

Device Type	**Power / Memory Consumption**
End-Node	
Semtech SX1276 [45]	Rx: 39.6 mW (12 mA@3.3 V) Tx: 396 mW (120 mA@3.3 V)
Seeed Studio Wio-E5 [43]	Rx: 22.1 mW (6.7 mA@3.3 V) Tx: 366.3 mW (111 mA@3.3 V)
Gateway	
The Things Indoor [48]	≤4.5 W, 160 KB SRAM
Kerlink iZeptoCell [23]	≤3.5 W, 256MB DRAM
SenseCAP M1 [47]	5 W, 1/2/4/8GB DRAM
SenseCAP LoRaWAN [44]	3.6 W, 512MB DRAM
MikroTik wAP LR2 [33]	≤8 W, 64MB DRAM
Signal Enhancer	
NELoRa [30]	CPU + GTX1080Ti (250W TDP) 10684MiB GPU memory
LoRaTrimmer [16]	CPU + RTX2080Ti (250W TDP) 8460MiB GPU memory
GLoRiPHY [41]	CPU + A6000 (300W TDP) 3656MiB GPU memory
DeepSense [9]	CPU + GTX1060 (120W TDP), source not available
LoRaSeek [36]	CPU + RTX3090 (350W TDP, real-time), Raspbbery Pi (5W, sub-real-time), 21.8MiB, source not available
RF-SR (our work)	Single CPU Core (15W TDP)[1] 25.7kB memory + input size
RF-SR INT8 PTQ[4]	NPU (14.2mW)[2], ASIC (<1mW)[3] 6.4kB memory + input size

Table 2: Power consumption of LoRa equipment. [1]Intel Core i7-1355U (TDP of the 10-core package), [2]Alif E3 ARM Cortex-M55/Ethos-U55 HE [1, 2], [3]Compute-In-Memory CNN Accelerator [28], [4]Post Training Quantization [39]

the lightweight network to achieve gains in both denoising and super-resolution. Through this work, we demonstrate that spatial super-resolution for RF signals is both feasible and practical, opening new paths for power-efficient, deployable neural enhancement in next-generation IoT communication systems.

Our key contributions are:

- To the best of our knowledge, we are the first to apply the super-resolution *render-low, upscale high* paradigm from visual computing to RF signals, demonstrating that the analog oversampling tax can be mitigated.
- We demonstrate that with a 4×/8× lower analog ADC sampling rate, our system achieves similar performance to native oversampling on real-world OTA data and a 1.25dB gain on synthetic data.
- We introduce a lightweight, shallow CNN architecture optimized for efficient edge and end-node inference, paired with a novel training tweak that enables robust signal enhancement in complex, low SNR environments.
- We present a comprehensive ablation study of DSP alternatives and quantization (spectral SR), showing that low-precision (4-6 bit) ADCs are sufficient for LoRa, offering a complementary power-saving vector.
- We released the source code for RF-SR to enable reproducibility and future development[1].

The remainder of this paper is structured as follows. Section 2 details the *System Design*. We present our main *Evaluation* on a large-scale over-the-air dataset in Section 3 followed by a comparison with related LoRa signal enhancers. The *Ablation Study* follows in Section 4, analysing DSP alternatives and investigating the effects of quantization. Finally, we cover *Related Work* in Section 5, discuss generalizability in Section 6, and conclude the paper in Section 7.

2 System Design and Implementation

Our system design is inspired by real-time video upscalers like RTX Video Super Resolution [11]. A key insight is that a neural network does not need to learn the difficult task of upsampling from scratch. Instead, we can leverage decades of signal processing research by using an efficient, traditional interpolation algorithm (e.g., polyphase or FFT-based [12]) and then task a small neural network with a much simpler problem: correcting the artifacts induced by interpolation and noise.

2.1 Why Oversampling Works

First, we investigate why oversampling is used by many prior works, and analyse the underlying working principle. Oversampling is effectively a trade-off between bandwidth and resolution, providing a signal-to-noise ratio (SNR) improvement known as *process gain* [24]. To quantify this, the ADC's quantization error $e[n]$ is modelled in this derivation as an uncorrelated, additive white-noise source P_Q [5, 42]. Under this assumption, the total quantization noise power $P_Q = \Delta^2/12$ is distributed uniformly across the full sampling bandwidth ($-F_s/2$ to $+F_s/2$), resulting in a constant Power Spectral Density (PSD) of $\Phi_Q = P_Q/F_s$.

A digital low-pass filter is applied to attenuate noise outside the signal bandwidth BW. The remaining in-band quantization noise power is calculated by integrating the PSD over the signal band:

$$P_{Q,\text{in-band}} = \int_{-BW}^{BW} \Phi_Q \, df = \int_{-BW}^{BW} \frac{P_Q}{F_s} \, df = P_Q \cdot \frac{2 \cdot BW}{F_s} \quad (1)$$

The Oversampling Ratio (OSR) is defined as the ratio of the sampling frequency to the Nyquist rate of the signal, $\text{OSR} = F_s/(2 \cdot BW)$ [5, 42]. Substituting the OSR into Equation 1 yields the fundamental relationship for oversampling converters:

$$P_{Q,\text{in-band}} = \frac{P_Q}{\text{OSR}} \quad (2)$$

This demonstrates that the in-band noise power is inversely proportional to the OSR [42]. Converting to decibels, the process gain is $10\log_{10}(\text{OSR})$ [24]. For a LoRa signal ($BW = 125\,\text{kHz}$) sampled at $F_s = 1\,\text{MHz}$, the OSR is 8, yielding a process gain of $10\log_{10}(8) \approx 9\,\text{dB}$. Since the theoretical dynamic range of a quantizer increases by approximately 6.02 dB per bit [24], this 9 dB improvement corresponds to an increase in resolution of 1.5 bits.

2.2 Network Architecture

The network is designed to be simple and efficient for real-time inference on gateway and end-node hardware. We employ a shallow

---

[1] https://github.com/andreaskuster/RFSuperResolution

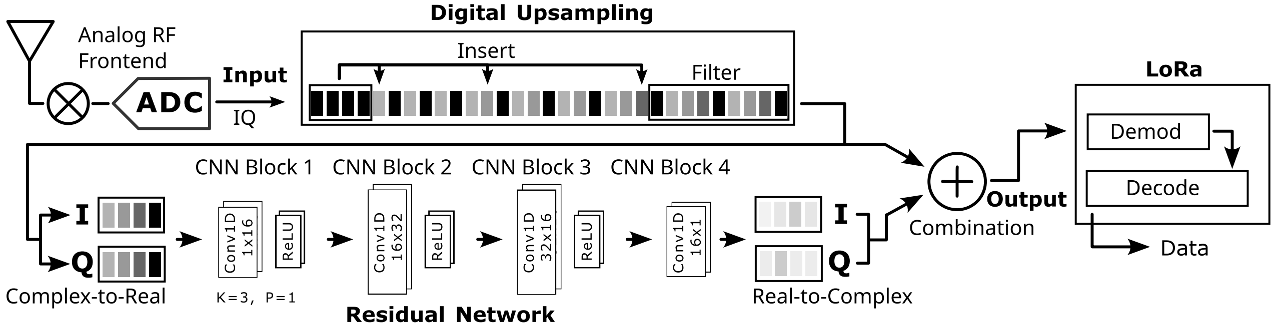

Figure 3: The RF-SR residual-learning architecture. A low-F_s signal is upsampled via standard interpolation. The CNN computes a residual (the artifacts + noise) which is added back to the interpolated signal.

1D Convolutional Neural Network (CNN) that learns an additive residual correction to a traditionally upsampled signal. The model's operation follows a "resample-and-refine" structure. We first upsample the low-rate signal $x[n]$ to $x_{interp}[m]$ using a standard polyphase filter bank [51]. This interpolated signal $x_{interp}[m]$ is spectrally correct but contains interpolation artifacts and noise (analyzed in Section 4). We feed this signal into our lightweight CNN, which learns to compute a residual signal, $\hat{r}[m]$. This residual represents the (complex-valued) error - the artifacts and noise. The final, clean output $\hat{y}[m]$ is the sum of the interpolated signal and the network's predicted residual, as shown in Figure 3.

$$\hat{y}[m] = x_{interp}[m] + \hat{r}[m] = x_{interp}[m] + \text{CNN}(x_{interp}[m]) \quad (3)$$

This structure allows the use of efficient DSP operations to handle the bulk effort of the upsampling, while the lightweight CNN focuses only on learning to correct the errors.

Complex-Valued Convolutions. A key requirement of our CNN is to handle complex-valued input, as RF I/Q data is inherently complex (e.g., $x = I + jQ$). Rather than implementing true complex arithmetic, we use a standard, computationally efficient approach. The real (I) and imaginary (Q) components of a C_{in}-channel complex signal are concatenated to form a $2 \cdot C_{in}$ channel real-valued tensor. A standard 1D convolutional layer then maps this to a $2 \cdot C_{out}$ channel tensor, which is subsequently re-interpreted as a C_{out}-channel complex signal. This allows the model to be efficient and compatible with standard deep learning hardware and libraries.

Residual Network Structure. The residual network, CNN($\cdot$), is a shallow four layer sequential model. It consists of four complex 1D convolutional layers with channel sizes progressing as $1 \rightarrow 16 \rightarrow 32 \rightarrow 16 \rightarrow 1$. All convolutions use a kernel size of 3 with padding of 1 to preserve the signal length at each step. A ReLU activation function is applied after the first three layers.

Model Simplification. The final architecture is the result of a systematic analysis of the design space, where we optimized depth, width, and kernel sizes to align the model's receptive field with the localized nature of LoRa interpolation artifacts. We conducted extensive trials with varied configurations, including gated mechanisms, different activation functions (e.g., ReLU vs. Sigmoid), and dropout regularization to mitigate noise sensitivity. While complex architectures like LSTMs [21] or deep stacks of dilated convolutions (Residual TCNs [3]) can maximize sequential memory, they proved inferior in this high-efficiency/low-SNR context for two reasons. First, increased complexity rendered the models several orders of magnitude more computationally expensive, making them impractical for low-power gateway and end-node deployment. Second, and more critically, excessive capacity led to severe overfitting; instead of correcting deterministic artifacts, these deep networks attempted to *hallucinate* chirps often collapsing demodulation performance. By opting for a shallow CNN with optimized padding and kernel dimensions, we constrained the model to focus its limited parameters on correcting localized discontinuities. This targeted design rationale prevents the destruction of the underlying LoRa signal in high-noise floors, resulting in a more robust and efficient architecture than its higher-parameter counterparts.

2.3 HybridDenoiserLoss Function

Training with a standard Mean Squared Error (MSE, or ℓ_2) loss worked for synthetic data with Additive White Gaussian Noise (AWGN). However, it failed to converge when trained on our large-scale over-the-air dataset. The over-the-air (OTA) data contains complex, non-Gaussian noise profiles from multipath, reflection, and impulsive interference, as well as hardware impairments like clock skew (time-domain shift) and carrier frequency offset (phase/frequency shift).

This failure has two root causes: 1) MSE loss is sensitive to outliers, and a single impulsive noise spike can dominate the gradient. 2) A small, uncorrected time-domain clock skew or phase-domain CFO between the prediction and the target signal results in a massive MSE loss, even if the predicted shape is perfect and could be demodulated by LoRa. This destabilizes training, especially at low SNR. To solve this, we introduce the *HybridDenoiserLoss*, shown in Algorithm 1. It combines two components:

1. **Time-Domain ℓ_1 Loss (MAE):** We replace ℓ_2 with ℓ_1 (Mean Absolute Error) in the time domain. ℓ_1 loss is robust to outliers, mitigating the impulsive noise problem.
2. **Frequency-Domain Magnitude ℓ_1 Loss:** To handle skew and CFO, we compute the ℓ_1 loss on the *magnitude* of the FFT. The FFT magnitude is invariant to time shifts and phase offsets, providing a stable, time-invariant signal representation that *guides* the network to learn the correct spectral shape, even if its time/phase alignment is imperfect.

The total loss is a weighted sum:

$$\mathcal{L}_{\text{Total}} = \lambda_{\text{Time}} \mathcal{L}_{\ell 1\_\text{Time}} + \lambda_{\text{Freq}} \mathcal{L}_{\text{MAE\_Mag}} \tag{4}$$

We found $\lambda_{\text{Time}} = 1.0$ and $\lambda_{\text{Freq}} = 0.5$ to be effective, prioritizing time-domain coherence while using the frequency domain as a robust regularizer.

Algorithm 1 Hybrid Denoiser Loss

```
Input: prediction $Y$, target $T$
Parameters: $\lambda_{\text{Time}}, \lambda_{\text{Freq}}$
Define: $\mathcal{L}_1(A, B) \leftarrow \text{MAE}(A, B)$
Define: $\mathcal{F}\{\cdot\} \leftarrow$ Fast Fourier Transform$(\cdot)$
1. Calculate Time-Domain L1 Loss
$\mathcal{L}_{\text{Time}} \leftarrow \mathcal{L}_1(Y, T)$
2. Calculate Frequency-Domain Magnitude L1 Loss
$Y_{\text{mag}} \leftarrow |\mathcal{F}\{Y\}|$
$T_{\text{mag}} \leftarrow |\mathcal{F}\{T\}|$
$\mathcal{L}_{\text{Freq}} \leftarrow \mathcal{L}_1(Y_{\text{mag}}, T_{\text{mag}})$
3. Combine Losses
$\mathcal{L}_{\text{Total}} \leftarrow (\lambda_{\text{Time}} \cdot \mathcal{L}_{\text{Time}}) + (\lambda_{\text{Freq}} \cdot \mathcal{L}_{\text{Freq}})$
return $\mathcal{L}_{\text{Total}}$
```

2.4 Training and Regularization

The OTA dataset still caused overfitting even with the hybrid loss. To combat this, we added a ℓ_1 regularization (Lasso) penalty to the model's weights. This encourages sparsity, forcing the network to learn only the most critical features of the interpolation artifacts rather than memorizing the high-entropy channel noise. The final loss function is:

$$\mathcal{L}_{\text{Final}} = \mathcal{L}_{\text{Total}} + \lambda_{L1} \sum_{\theta \in \text{CNN}_\theta} |\theta| \tag{5}$$

Our training process is two-stage: first, we train the model on synthetic LoRa signals to learn the fundamental signal structure and interpolation artifacts. Second, we fine-tune the model using the OTA training set to adapt it to real-world channel complexities.

2.5 Computational Feasibility

To quantify the computational feasibility of RF-SR, we explicitly derive the operational load required to sustain a real-time throughput of $R_s = 1\,\text{MSPS}$ (250kSPS 4x upsampling). Our architecture processes the I and Q signal components as independent channels, utilizing two parallel, real-valued convolutional streams. We define the operational load (C_{ops}) in terms of floating-point operations (FLOPs), where each Multiply-Accumulate (MAC) constitutes two operations (one multiplication and one addition).

For the specified four layer network, the total MACs per stream is the sum of products of kernel size (K), input channels (C_{in}), and output channels (C_{out}). The total system load accounts for both I/Q streams and the operations per MAC:

$$C_{\text{ops}} = R_s \times \underbrace{2}_{\text{I/Q streams}} \times \underbrace{2}_{\text{Ops/MAC}} \times \sum_{l=1}^{L} K_l \cdot C_{in,l} \cdot C_{out,l} \tag{6}$$

Substituting the layer dimensions (1⟶16⟶32⟶16⟶1, with $K = 3$), the summation yields 3168 MACs per stream. The resulting load is:

$$C_{\text{ops}} = 1\,\text{MHz} \times 2 \times 2 \times 3168 \approx 12.7\,\text{GOPS} \tag{7}$$

We demonstrate that this performance is attainable across a hierarchy of processing units, starting with a standard commodity laptop equipped with an Intel Core i7-1355U processor (15W total package TDP, ~48GOPS/core). While this general-purpose CPU handles the throughput load, efficient deployment on constrained edge hardware can be achieved through Post Training Quantization (PTQ) [39]. By reducing the model's precision from floating-point to 8-bit integers (INT8), we minimize memory bandwidth and computational requirements, allowing the model to fit within the strict power envelopes of modern high-efficiency edge AI hardware like the *Alif E*3 microcontroller with its ARM Cortex-M55 and Ethos-U55 High-Efficency (HE) Neural Processing Unit (NPU) [2]. In this configuration, the INT8 optimization allows the NPU to sustain the required throughput with approximately 12% utilization, consuming at most 14.2 mW [1]. Furthermore, for dedicated ASIC implementations, we project the power consumption (P_{dynamic}) using the reported 40 TOPS/W [28], illustrating the potential for sub-milliwatt operation:

$$P_{\text{dynamic}} \approx \frac{12.7\,\text{GOPS}}{40\,\text{TOPS/W}} = 0.3175\,\text{mW} \tag{8}$$

Our performance evaluation of the PTQ model in Figure 18 shows no degradation of up to -20dB and minimal degradation thereafter, making this a suitable choice.

3 Evaluation

3.1 Experimental Setup

We conducted a large-scale data collection of 10,000 LoRa packet samples on our university campus. We placed a transmitter[2] in a fixed indoor lab location. We then moved receivers[3] to various indoor and outdoor locations to capture diverse channel conditions, including indoor Line-of-Sight (LoS), short/mid-range Non-Line-of-Sight (NLoS), and 800m *over-the-hill* scenarios, as shown in Fig. 5(a). These locations specifically encompass both freestanding outdoor environments and dense urban building blocks to ensure robustness against complex multipath and shadowing. To cover a large dynamic range, we varied the transmit power and used RF attenuators on the receiver. All packets used LoRa parameters: SF = 12, BW = 125 kHz, 16-byte random payload, 8 preamble symbols, and CR = 4/5; the 125 kHz bandwidth was used as it is the de-facto standard (Table 1), and SF12 was chosen to demonstrate/deal with the worst-case SNR required for long-range communication. We employ the open-source *SDR-LoRa* platform [8], which achieves commercial-grade SNR performance for packet demodulation and decoding, allowing us to report internal demodulation metrics and implement modifications like windowing, zero-padding, and testing symbol-only enhancement techniques such as *LoRaTrimmer* [16], NeLoRa [30] and GLoRiPHY [41] in an end-to-end, real-world fashion. Furthermore, we use 250kSPS sampling rate as baseline despite the 125kHz bandwidth, in order to avoid aliasing for non-perfecly centered signal (CFO). To account for hardware impairments across

---

[2] Ettus Research B205mini [18]

[3] Ettus Research N210 [19]

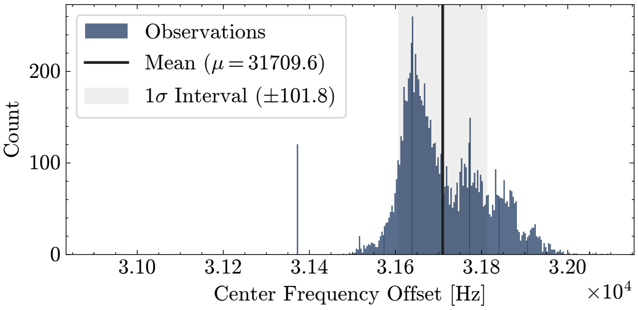

Figure 4: Over-the-air CFO distribution. Shaded area shows the $\pm 1\sigma$ interval ($\sigma \approx 101.8$ Hz), capturing hardware variance and environmental instability.

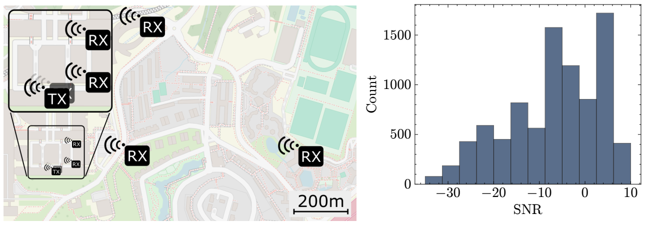

(a) Data collection node placement. **(b) Histogram of packet SNRs.**

Figure 5: Overview of the over-the-air dataset.

diverse environments, our dataset captures significant Carrier Frequency Offset (CFO) instability, with a measured mean of 31.7 kHz and a standard deviation of 101.8 Hz (see Figure 4). These variations, ranging from 30.9 kHz to 32.1 kHz, reflect the impact of heterogeneous placement and fluctuating environmental conditions on oscillator stability. Our observations align with real-world COTS hardware behaviour [59], which demonstrates that SX1276 modules experience significant temperature-dependent frequency drift and can vary by hundreds of Hertz across different days. The Signal-to-Noise Ratio (SNR) is calculated using the measured power during packet reception ($P_{\text{signal+noise}}$) and the noise power (P_{noise}) captured when no signal is transmitted:

$$\text{SNR}_{\text{dB}} = 10 \log_{10}\left(\frac{P_{\text{signal}}}{P_{\text{noise}}}\right) \approx 10 \log_{10}\left(\frac{P_{\text{signal+noise}} - P_{\text{noise}}}{P_{\text{noise}}}\right) \quad (9)$$

The resulting SNR distribution of the collected packets is shown in Figure 5(b). The full dataset was randomly split into an 80% training set and a 20% test set.

3.2 Oversampling Performance Comparison

We first evaluate our system on the captured OTA test set. Figure 6 plots the Packet Error Rate (PER) vs. SNR. We compare five key configurations: Std LoRa (250 kHz), which is the 2× Nyquist signal; the oversampled signal, including 4× (500 kHz), 8× (1 MHz), and the "gold standard" 16× (2 MHz) signal; and finally, RF-SR, which uses the 250 kHz signal as input and digitally upscales it to 1 MHz. The results show that higher sampling rates improve performance. The 2× (250 kHz) signal performs poorly, 4× (500 kHz) is better, and 8×/16× plateaus. This confirms why the papers in Table 1 default to 8× oversampling. The key result is that our RF-SR system, using only the low-rate 250 kHz input, achieves a PER curve that is comparable to the native 1 MHz/2 MHz baseline. This demonstrates that we can mitigate the expensive analog oversampling with a computationally efficient digital upscaler without loss in demodulation performance.

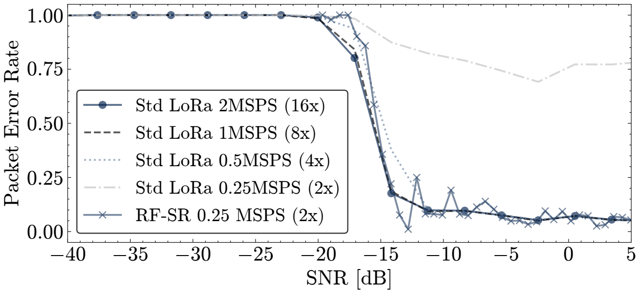

Figure 6: End-to-end demodulation/decoding PER vs. SNR performance on the OTA test set.

3.3 LoRa Enhancer Comparison

To provide a comprehensive comparison to LoRa enhancer works, we evaluated both neural network based [30, 41] and DSP enhancers [16] using AWGN, on which they where trained and reported (Table 1) in Figure 8. We co-profiled execution time and power consumption[4] using a unified framework based on NVML (GPU) [37] and Intel RAPL (CPU) [22]. Certain related works, such as DeepSense [9] and LoRaSeek [36], could not be included in this benchmark as their source code is not publicly available. The SNR performance is reported at the point were the packet error rate exceeds 20%. The INT8 energy is based on *YOLO-Fastest* object detection benchmark on the Alif E3 MCU [1] for the model, and the SX1276 end-node [45] for demodulation.

In this scenario with the simpler, more predictable noise model, RF-SR achieves a gain of 1.25dB compared to the 16x oversampled signal, showcasing that it is not only able to correct interpolation artifacts but also partially denoise the signal. The memory consumption for RF-SR is dominated by feeding the 1MSPS (upsampled) signal into the standard LoRa software-based packet decoder [8] (full packet). In comparison to standard LoRa, RF-SR lies right between 1 and 2MSPS Std LoRa for both energy and memory, adding minimal memory, and a slight amount of energy overhead, while achieving a 1.25dB, 1.3dB, 2.9dB SNR gain relative to 2/1/0.5MSPS. While the post training INT8 quantized model maintains similar performance to the full-fledged RF-SR model, it uses substantially less energy and memory, increasing RX power consumption of a standard LoRa SX1276 module by 30% running on the NPU or less than 2% running an a dedicated ASIC. LoRaTrimmer showed the strongest performance with 7.9dB gain over the gold standard, however with 386Ws, it has the highest energy consumption. Both NeLoRa and GLoRiPHY show a marginal gain of 0.3dB above the 16x oversampled signal, which is 0.95dB lower compared to RF-SR. This comes at a slightly lower energy budget than LoRaTrimmer for

---

[4] Intel Core i9-10980XE, Nvidia GeForce RTX 3090

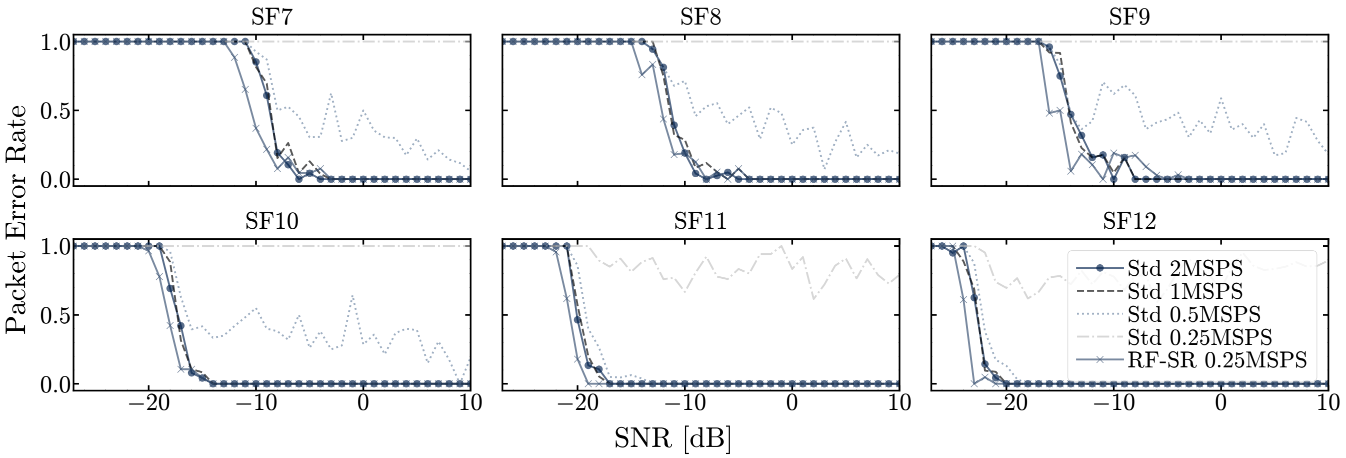

Figure 7: End-to-end Packet Error Rate (PER) vs. SNR across LoRa spreading factors *SF*7–*SF*12. RF-SR (0.25 MSPS input) consistently maintains a lower error rate and achieves superior SNR margins compared to the native 2 MSPS oversampled baseline across all spreading factors.

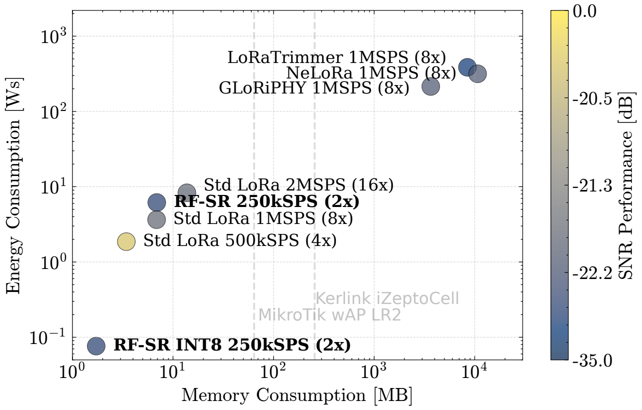

Figure 8: LoRa enhancer energy, memory and performance comparison.

both NeLoRa and GLoRiPHY, but a higher memory consumption for NeLoRa. The memory and energy consumption for all three works is significantly above the available memory on COTS LoRa gatways [23, 33] for real-time processing, and orders of magnitude higher than RF-SR.

3.4 Impact of Spreading Factor

To evaluate the robustness of RF-SR across diverse LoRa operational modes, we expand our evaluation to include the full range of spreading factors (*SF*7-*SF*12). Different SFs vary significantly in symbol duration and noise sensitivity; lower SFs result in shorter airtime and fewer samples per symbol for the neural network to process. Despite being optimized for the most signal-rich *SF*12 cases, our results in Fig. 7 demonstrate that the *render-low, upscale-high* paradigm generalizes effectively across all configurations. RF-SR consistently outperforms the native 2 MSPS oversampled baseline while operating on a significantly reduced 0.25 MSPS input. Specifically, we observe gains of 0.8 dB for *SF*7 and *SF*11, 1.0 dB for *SF*8 and *SF*10, and 1.25 dB for *SF*12, with a peak improvement of 2.2 dB for *SF*9. This indicates that the lightweight CNN successfully corrects interpolation artifacts regardless of the specific chirp slope or time-domain duration.

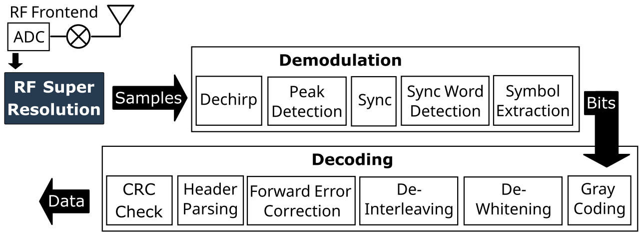

Figure 9: The LoRa receive pipeline, showing our RF-SR module inserting a digital super-resolution step after the ADC.

3.5 LoRa-Internal Performance Analysis

While the PER results in Section 3.2/3.3 validate our system at the packet level, a deeper analysis using the Symbol SNR metric is required to confirm that our digital upscaler truly replicates the internal confidence of a native high-sampling-rate system.

The LoRa receive pipeline, based on open-source implementations [7, 57] and with RF-SR integration, is shown in Figure 9. After the analog front-end and digitization, the signal undergoes dechirping, window alignment, and FFT-based symbol extraction. The output of the FFT is fed to a peak-finding algorithm to determine the symbol value. This is the core demodulation step. This pipeline allows us to see not just if a packet fails, but how close it was to succeeding or failure. After de-chirping and the FFT, the 2^{SF} (e.g., 4096) output bins correspond to the possible symbol values. We introduce the *Symbol SNR* (SymSNR) metric. As shown in Figure 10,

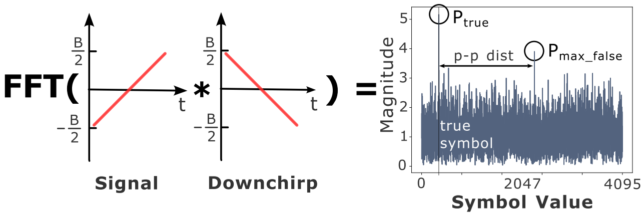

Figure 10: Definition of the Symbol SNR metric: the ratio of the energy in the true FFT peak to the energy in the highest false peak.

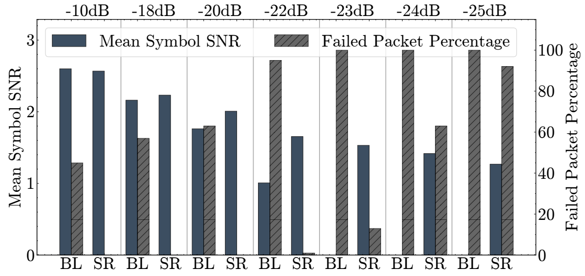

Figure 11: Symbol SNR and packet failure comparison between the baseline (BL) and RF-SR (SR).

we define SymSNR as the ratio of the energy in the *true* symbol's bin (P_{true}) to the energy in the *highest false* bin ($P_{\text{max\_false}}$).

$$\text{Symbol SNR} = \frac{P(\text{bin}_{\text{true}})}{P(\text{bin}_{\text{max\_false}})} \tag{10}$$

A high SymSNR indicates high demodulation confidence. A SymSNR < 1.0 results in bit errors, as the wrong peak/symbol is chosen. Comparing the baseline to the RF-SR upsampled performance in Figure 11, we can observe that even at -10dB, while the SymSNRs are similar, a substantial amount of packet decoding failed for the baseline, gradually reaching 100% at -23dB. On the other hand, RF-SR manages to distribute the noise more evenly, visible at the much more steady reduction of SymSNR from -10dB to -25dB compared to the baseline, and by managing to stay packet error free until -20dB, with a faster jump of packet loss reaching 90% at -25dB.

4 Ablation Study

We now analyse *why* our system works by ablating its core DSP and machine learning components, comparing them against other common DSP techniques. Following this, we investigate quantization to explore the "spectral" dimension of super-resolution and determine if high bit resolution is necessary.

4.1 AI Tax

While RF-SR offers performance gains, it is essential to quantify the specific value added by the neural network versus classical digital signal processing. Since our model first utilizes a standard polyphase interpolator, we evaluate whether the computational *tax* of the subsequent four-layer CNN is justified by the resulting improvement in demodulation. Figure 12 compares the Packet Error Rate (PER) of native sampling rates against both a purely interpolated signal (Interp 4x) and the full RF-SR pipeline. At the 20% PER threshold, upsampling a 0.25 MSPS signal to 1 MSPS via interpolation alone achieves -21.5 dB, which performs better than native 0.5 MSPS (-21.1 dB) but worse than native 1 MSPS (-22.2 dB). In contrast, the full RF-SR model reaches -23.3 dB, providing a 1.8 dB gain over interpolation alone and a 1.1 dB improvement over the native 1 MSPS oversampled baseline. The justification for this AI tax depends on the specific use case: high-performance gateways benefit from this critical link margin, while extremely resource-constrained end-nodes may prefer pure interpolation to minimize memory and energy overhead.

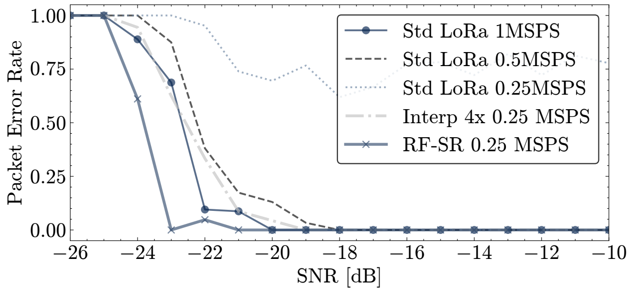

Figure 12: PER comparison between native sampling, digital interpolation, and RF-SR at SF12.

4.2 How Interpolation Can Help

The 2× Nyquist Problem. Figure 13(a) shows a spectrogram of a LoRa signal sampled at F_s=250 kHz. The signal is severely distorted. The "vertical lines" are spectral leakage artifacts that occur when the chirp "wraps" from +BW/2 to −BW/2. This instantaneous frequency/phase discontinuity in the time domain scatters energy across all frequency bins. This is not a signal, but a FFT artifact. Furthermore, the low F_s provides poor time resolution, "blurring" the chirp over time.

Short FFT Window. This artifact can be mitigated by using a shorter FFT window (Figure 13(f)), which improves time resolution. However, this creates an even worse problem: the frequency resolution is heavily degraded, making the FFT bins too coarse to distinguish symbols. Thus, we either have to trade time or frequency resolution, which are both essential, especially at low SNR.

Interpolation. The fundamental problem is a lack of time-domain samples in general, but particularly at the sharp frequency jump. Digital interpolation (Figure 13(k)) mitigates this dilemma. By mathematically creating intermediate time-domain samples, it allows us to simultaneously achieve the high time resolution (to prevent artifacts) and long FFT windows (high frequency resolution) to find peaks that only a higher sampling rate normally provides.

This digital upsampling is a foundational technique in multi rate signal processing [12]. It is a well-established method used in hardware-level applications, such as relaxing analog filter specifications in high-speed Digital-to-Analog Converters (DACs) [25] and as a core component of ΣΔ (Sigma-Delta) ADCs [5, 42]. The goal is to mathematically create intermediate samples, often to avoid issues like aliasing or interpolation artifacts [20]. However,

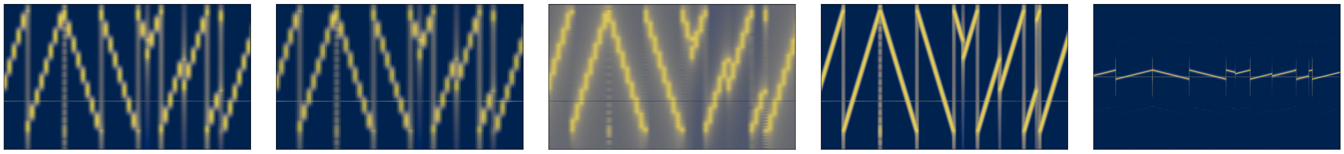

(a) FFT1024^1 **(b) Zero-Padding 8x^1** **(c) No Window^1** **(d) Overlapping FFT^1** **(e) Polyphase Interp.^{3}**

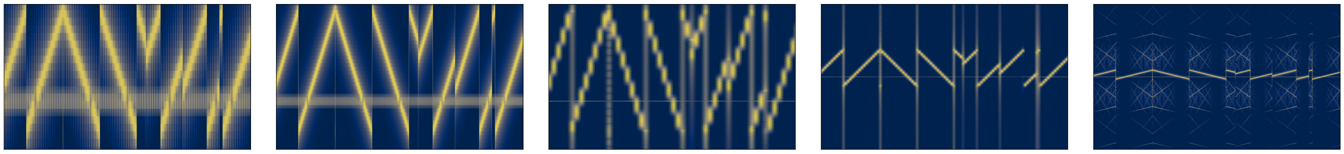

(f) FFT32^1 **(g) FFT64^1** **(h) Hann Window^1** **(i) No Filter^2** **(j) PChip Interp.^{3}**

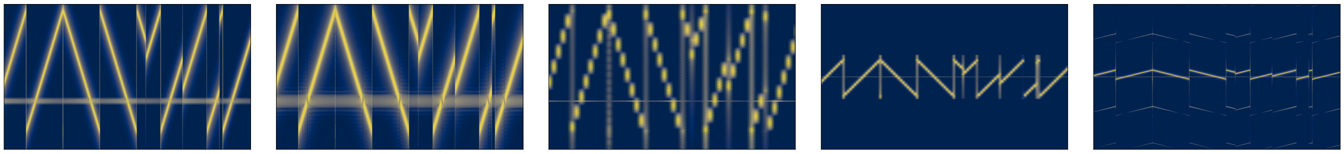

(k) Interpolation 8x^1 **(l) Zero-Pad 8x + FFT64^1** **(m) Blackman-Harris^1** **(n) FIR Filter^2** **(o) Linear Interp.^{3}**

Figure 13: Ablation Study. Column 1 (a, f, k): Demonstrates the challenge of low sampling frequency and the mitigation using digital interpolation. Column 2 (b, g, l): Illustrates the effect of zero-padding on signal representation in both time and frequency domains. Column 3 (c, h, m): Compares the effects of different FFT windowing technique on reducing spectral artifacts. Column 4 (d, i, n): Shows the benefit of overlapping FFT windows and the limited effectiveness of Low-Pass Filters (LPF) in removing in-band artifacts. Column 5 (e, j, o): Compares the performance of different digital interpolation algorithms. X-Axis: 0.22s, Y-axis:^{1}125kHz, ^{2}500kHz, ^{3}2MHz

the choice of algorithm is critical; Figure 13(e) shows that while polyphase based resampling performs well, simpler methods like Linear (Figure 13(o)) or non-RF optimized methods like PCHIP (Figure 13(j)) introduce artifacts.

4.3 Alternative DSP Techniques

We evaluated other common DSP techniques usind with LoRa to mitigate artifacts from low-F_s sampling, and how they relate to RF-SR.

Zero-Padding. Zero-padding in the time domain increases frequency resolution in the FFT, as demonstrated in Figure 14, and has been used in LoRa demodulation [57] to improve packet reception. However, it does *not* create new time-domain samples, so it does not increase the time resolution. As seen in Figure 13(b), it fails to remove the chirp-wrap artifacts. Pairing zero-padding with a shorter FFT window eliminates the warp and improves spectral resolution (Figure 13(l), but not to the same level as we achieve with interpolation (Figure 13(k))

FFT Windowing. The given artifacts are a form of spectral leakage (Figure 13(c), no windowing). Applying a windowing function (e.g., Hann in Figure 13(h)) tapers the signal at both boundaries, which successfully reduces these artifacts. However, windowing achieves this by discarding signal energy towards both start and end of the signal to make it cyclic, which lowers the final FFT peak energy. This effect is even stronger in Blackman-Harris (Figure 13(m)). This signal energy loss is a problem in a low-SNR system and must be carefully traded off, as we will discuss in the symbol SNR study in Figure 18.

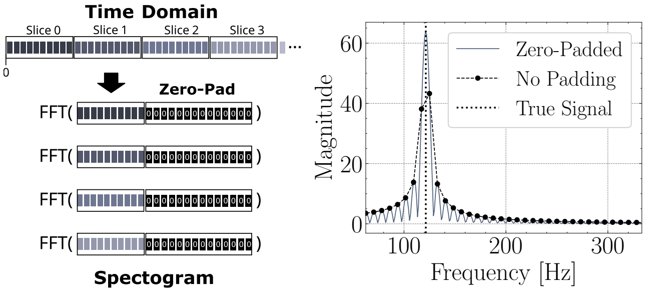

Figure 14: Zero-padding principle and effect on the FFT.

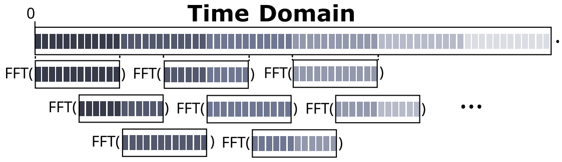

Figure 15: FFT window overlapping principle.

Overlapping FFT. The conventional approach to spectral analysis involves slicing the continuous sample stream into sequential, non-overlapping window blocks. While it is simple and low compute, this method inherently ties the time resolution to the length of the window, meaning better frequency resolution comes at the

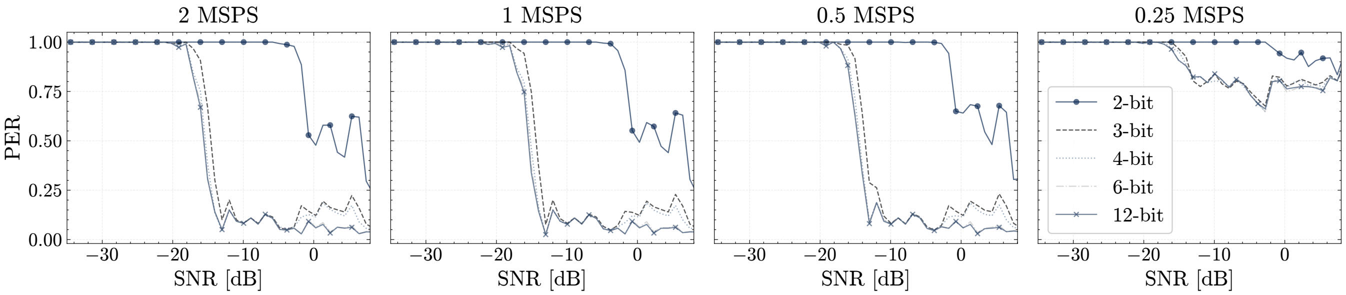

Figure 16: PER vs. SNR for varying ADC bit depths at different sampling rates (F_s). Performance is robust, showing that high-precision ADCs are not necessary for effective LoRa demodulation.

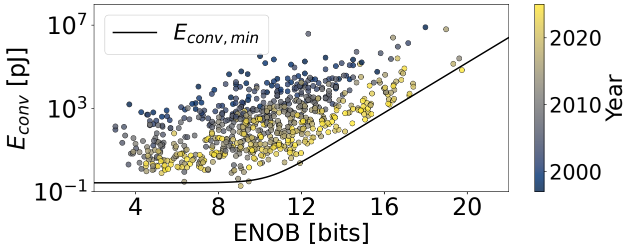

Figure 17: ADC conversion energy survey.

expense of poorer time resolution. This can be mitigated by using overlapping FFT windows (Figure 15). Instead of processing sequential slices, the FFT computation is performed over windows that overlap. This technique significantly increases the temporal density of the spectral estimates, effectively improving the time resolution without sacrificing the spectral (frequency) resolution. While this technique provides a higher refresh rate of the spectral data, it does so at the cost of increased computational load proportional. Furthermore, the fundamental problem, of having to few samples at the fast wrap time persists in (Figure 13(d)).

Filtering. A sharp digital low-pass filter can remove out-of-band noise (Figure 13(i)/13(n)). However, the chirp-wrap artifact is *in-band* spectral leakage. Filtering is ineffective against it.

4.4 Quantization

Finally, we explore the *spectral* dimension of super-resolution: bit precision. There is a physical relation between sampling rate and bit resolution, most notably in ΣΔ (Sigma-Delta) ADCs [6]. These ADCs use extreme oversampling (high OSR) to shape the quantization noise, allowing a simple 1-bit quantizer to achieve a very high Effective Number of Bits (ENOB) after digital filtering. For an L^{th}-order ΣΔ modulator, the in-band quantization noise power P_q scales with OSR [42]:

$$P_q \propto \left(\frac{1}{\text{OSR}}\right)^{2L+1} \tag{11}$$

Looking at it from an energy per conversion angle, the survey in Figure 17 shows that conversion energy stays relatively flat for ADCs with an Effective Number of Bits (ENOB) ≤ 10, but grows exponentially with ENOB as illustrated by the trendline:

$$E_{\text{conv,min}} = 0.27\text{ pJ} + 0.145\text{ aJ} \cdot 4^{\text{ENOB}} \tag{12}$$

The formula for ENOB, where SNDR (Signal-to-Noise and Distortion Ratio) is the metric for total dynamic performance [24, 26], is

$$\text{ENOB} = \frac{\text{SNDR} - 1.76}{6.02} \tag{13}$$

Equation 11, 12 and 13 imply a fundamental link between sampling rate and bit resolution.

We tested this by taking our high-rate 2 MHz signals and quantizing them from 12-bits down to 2-bits. The results, shown in Figure 16 were unexpected. Performance is stable down to 6 bits, and 4 bits shows minor degradation. Performance only collapses at 2 bits. This is a key finding, showing that high precision is not required. LoRa's CSS modulation and dechirp/FFT pipeline is a massive correlation process that provides significant processing gain. This gain is designed to pull the signal peak out from under the *thermal noise floor* (P_N). Our results show that as long as the *quantization noise floor* (P_Q) is below the thermal noise floor, it has no impact on demodulation. For LoRa, the noise floor of a 4-bit or 6-bit ADC is already good enough. This finding, combined with our primary result, suggests that the optimal front-end for LoRa is not a high-speed, high-precision ADC, but rather a low-speed (we can digitally upsample with RF-SR) and low-precision ADC, compounding the power savings.

4.5 SymSNR Performance Comparison

In Figure 18, we compare different signal enhancement scenarios from the ablation study; *W* (Windowing), *Z* (Zero-Padding), *F* (Filtering), and *Q* (INT8 Quantization), and their performance at increasing SNR. Beside Symbol SNR, we also showcase the count of *failed packet decoding*. This is crucial, as SymSNR cannot be determined for failed packets due to the missing time and frequency synchronization, obscuring the true time decoding capability.

Baseline and Filtering Performance. The *Baseline* at 250 kHz sampling rate exhibits poor performance, with significant packet decoding failures starting at SNR = −10 dB and a complete breakdown from −22 dB onwards. *Strong Signal FIR Filtering (F)* shows similar performance to the baseline around −10 dB. However, its

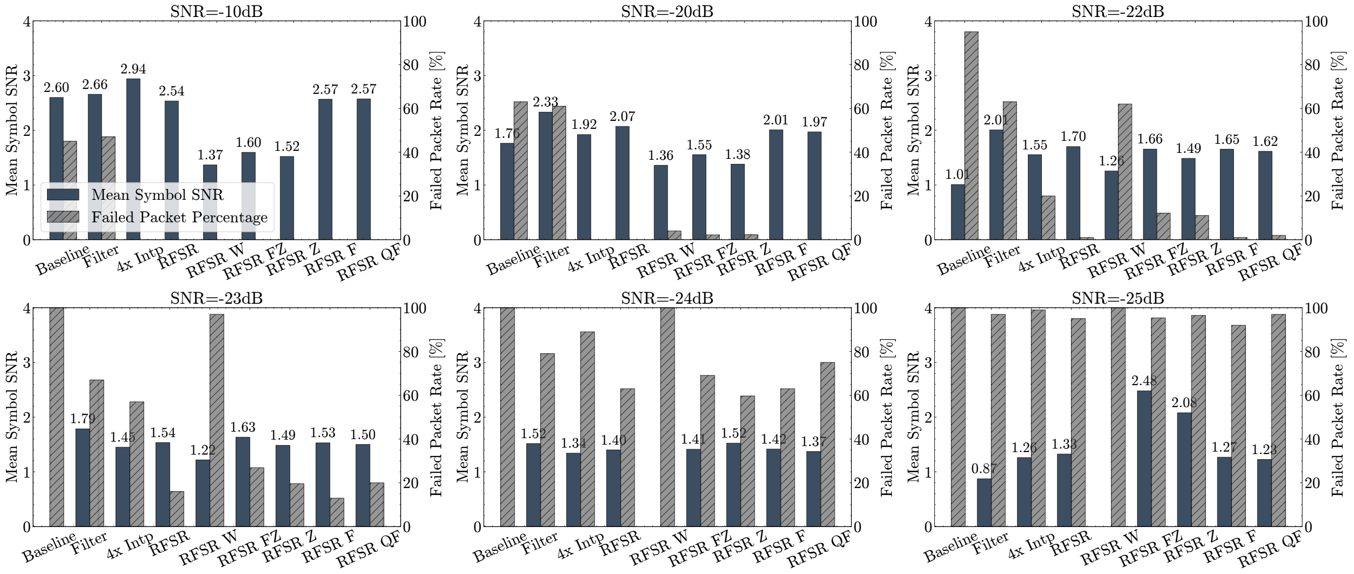

Figure 18: A LoRa-internal symbol SNR and packet error comparison for baseline (250kSPS), filter (FIR), 4xInterpolation (Resample Poly), RF-SR (W=windowing, F=filter, Z=zero padding, Q=8bit integer quantization).

relative performance improves as SNR increases, proving to be useful by strongly rejecting *out-of-band noise* but struggling with in-band noise. When applied as pre-filtering on top of RFSR (*RFSR F*), it demonstrably improves performance, particularly in low SNR cases.

Interpolation vs RFSR. *Polyphase Interpolation* alone offers a significant performance improvement over the baseline, even exceeding the RFSR performance in high SNR cases. However, the *RFSR* model, which builds upon interpolation but actively *denoises and removes artifacts*, shows superior performance when considering low SNR cases (for which it was optimized), successfully maintaining decoding capability down to −23 dB.

Windowing. *Windowing (W)* proves to be detrimental. By removing valuable signal information near the boundaries, it fails in cases where other techniques still succeed (e.g., −22 dB and −23 dB).

Zero-Padding. *Zero-Padding (Z)* shows similar, but slightly worse performance than filtering. Interestingly, the combination of both filtering and zero-padding (*RFSR FZ*) shows consistently worse performance than each method separately.

Best Performing Model and Quantization. The most consistently best performing configuration is RFSR together with FIR pre-filtering *(RFSR F)*. To meet the realistic milliwatt power targets for end-node deployment (as introduced in Table 2), we evaluated the model performance when quantized down to INT8 (*RFSR QF*). Up until −22 dB, the quantized model performs almost equivalent to its floating-point counterpart, with only slight degradation down to −25 dB. This demonstrates the high usefulness and efficiency of the model in this resource-constrained configuration.

Peak-to-Peak Distance Distribution. To better quantify the benefit of zero-padding, particularly its contribution to improved spectral resolution, we investigated its effect on symbol errors. The conjecture is that higher spectral resolution should lead to fewer symbol errors where the erroneous peak is near the true peak due to the higher frequency resolution. To evaluate this, we introduce the *peak-to-peak distance*: the distance between the ground truth peak and the erroneous symbol peak (Figure 10), and plot its distribution in Figure 19. For the SNR = −20 dB case (left panel), the RFSR model exhibits a lower frequency of bit errors with a very small peak-to-peak distance. This suggests that the zero-padding component is operating as expected by making the symbol peaks more distinguishable, thus preventing small errors. However, this benefit diminishes at even lower SNR cases, specifically SNR = −24 dB

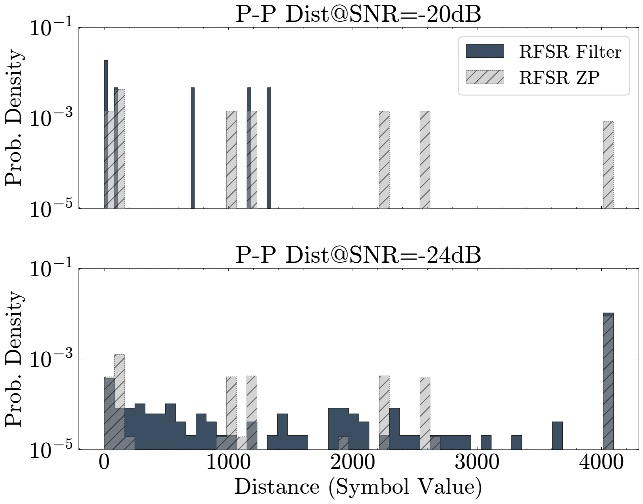

Figure 19: Symbol value distance distribution from true to false peak for wrongly demodulated symbols (bit errors).

(right panel). Here, both the zero-padded and non-zero-padded variants show a similar distribution of erroneous peaks, particularly a high number of peaks with a distance very close to 1 or 4095 (which is ±1 apart in the cyclical LoRa spectrum). This suggests that at extremely low SNR, the noise dominates, overwhelming the resolution gains provided by zero-padding.

5 Related Work

Our work is inspired by and contributes to several distinct domains.

Visual Super-Resolution. Our core methodology comes from visual computing. Foundational works like SRCNN [14], VDSR [27], and EDSR [32] proved that CNNs could learn complex upscaling mappings. More relevant to our system-level goals are real-time upscalers like NVIDIA's DLSS and VSR [11], which introduced the efficient paradigm of using classical interpolation followed by a neural *residual-correction* network [11]. To the best of our knowledge, we are the first to adapt and apply this super-resolution and residual-learning paradigm to the RF domain for sampling rate enhancement.

Neural-Enhanced LoRa/LPWANs. A growing body of work applies heavy signal proessing or deep learning to LoRa, primarily falling into two categories: *heavyweight denoisers* (e.g., NELoRa [30], GLoRiPHY [41], LoRaTrimmer [16]) and *diversity-based systems* (e.g., Charm [15], OPR [4], Choir [17]). Our approach is distinct in four main aspects: **(1) Problem Solved:** While existing denoisers typically maintain the standard 8× oversampling, we investigate the *super-resolution* paradigm to reduce sampling rates, aiming to lower ADC, DSP and memory requirements; **(2) Practicality:** Previous neural enhancers generally utilize high-performance GPUs. However, given the reliability required in non-easily serviceable scenarios, the susceptibility of such hardware to faults (e.g., bit errors, bus instability) [13, 29, 38] presents a challenge. Our work instead targets the constrained power and memory budgets of edge gateways and end-nodes by employing a shallow, quantized CNN; **(3) Overhead:** Distinct from diversity schemes that rely on multiple synchronized nodes or gateways, our system focuses on enhancing single packets from a single receiver; and **(4) Robustness:** We extend evaluation beyond synthetic AWGN noise models by validating on large-scale over-the-air datasets to address complex real-world channel conditions.

LoRa Collision Resolution. A major focus of the LoRa research community has been multi-packet reception (e.g., [10, 40, 46, 49, 50, 55, 56, 58]). These works have dramatically improved LoRa's network capacity. However, as shown in Table 1, their implementations universally rely on the 8× "oversampling tax." Our work is complementary; RF-SR could be used as a preprocessing step in these systems to reduce their frontend power consumption.

Multirate DSP and ADC Design. Our work is grounded in the foundational theory of multirate signal processing [12], which provides the classical interpolation methods we use. It directly engages with the ADC design trade-offs established [6, 42]. Our findings in Section 4.4 suggest that the conventional wisdom of trading high F_s for high ENOB (the $\Sigma\Delta$ approach) may be an inefficient design for spread-spectrum systems, which does not require high precision.

6 Discussion

Generalizability to Other RF Protocols While this work is validated on LoRa, the "render-low, upscale-high" paradigm is fundamentally protocol-agnostic and suggests a promising venue for any RF application currently burdened by a high oversampling tax. This lightweight neural correction approach could potentially be adapted to other spread-spectrum, high-bandwidth, or resource-constrained systems where high-rate ADCs are traditionally used to gain SNR or process gain.

7 Conclusion

The energy cost of high-rate ADCs, driven by a *de facto* 8× oversampling standard in the LoRa research community, represents a critical and unaddressed bottleneck. We challenged this paradigm by proposing RF-SR, a system that adapts the *render low, upscale high* principle from visual computing to the RF domain. Our solution combines efficient digital interpolation with a shallow, four layer complex-valued CNN that learns to correct only the residual artifacts and noise. To enable training on real-world signals, we introduced the *HybridDenoiserLoss*, which balances time-domain and frequency-magnitude errors to provide a stable gradient in the presence of clock skew and non-Gaussian over-the-air noise. Based on a large-scale, 10000-sample OTA dataset, we demonstrated that RF-SR processing a 2×-Nyquist (250 kHz) signal achieves similar demodulation performance as a native 8× oversampled (1 MHz) system. Furthermore, our model provides an additional 1.25dB SNR improvement. Our ablation study revealed that high-precision ADCs are not essential, offering an orthogonal power-saving vector. In contrast to prior signal enhancers that rely on heavy computational resources, RF-SR is memory and compute efficient to be deployed on commodity gateways and, when quantized to INT8, on resource-constrained end-nodes. This lightweight and practical design frees IoT systems from the "oversampling tax," enabling a new class of low-power frontends built on low-speed, low-precision ADCs without sacrificing communication range or reliability.

Acknowledgments

We would like to thank our shepherd and the anonymous reviewers for their insightful feedback and constructive comments, which greatly improved the quality of this paper. This research is supported in part by the National Research Foundation, Singapore, under its AI Singapore Programme (AISG Award No: AISG4-GC-2023-006-1B), in part by Singapore Ministry of Education under its AcRF Tier-1 grant RT14/22, and by Hong Kong GRF under Grant 16204224 and 16210425.

A Artifact Appendix

The research artifacts accompanying this paper are available via DOI: https://doi.org/10.5281/zenodo.19451778.

References

[1] Alif Semiconductor. 2023. *How Hardware Accelerated Machine Learning Transforms the Landscape For Embedded IoT Devices.* Technical Report. Alif Semiconductor. https://alifsemi.com/download/AWPR0007 Version 1.0, AWPR0007 November 2023.

[2] Alif Semiconductor. 2025. *Ensemble E3 Series Datasheet.* https://alifsemi.com/download/ADTS0007 ADTS0007 v2.12—May 2025.

[3] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. 2018. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv:1803.01271* (2018).

[4] Artur Balanuta, Nuno Pereira, Swarun Kumar, and Anthony Rowe. 2020. A cloud-optimized link layer for low-power wide-area networks. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services* (Toronto, Ontario, Canada) *(MobiSys '20)*. Association for Computing Machinery, New York, NY, USA, 247–259. doi:10.1145/3386901.3388915

[5] B.E. Boser and B.A. Wooley. 1988. The design of sigma-delta modulation analog-to-digital converters. *IEEE Journal of Solid-State Circuits* 23, 6 (Dec 1988), 1298–1308. doi:10.1109/4.90025

[6] B.E. Boser and B.A. Wooley. 1988. The design of sigma-delta modulation analog-to-digital converters. *IEEE Journal of Solid-State Circuits* 23, 6 (Dec 1988), 1298–1308. doi:10.1109/4.90025

[7] Fabio Busacca, Stefano Mangione, Sergio Palazzo, Francesco Restuccia, and Ilenia Tinnirello. 2024. SDR-LoRa, an open-source, full-fledged implementation of LoRa on Software-Defined-Radios: Design and potential exploitation. *Computer Networks* 241 (2024), 110194. doi:10.1016/j.comnet.2024.110194

[8] Fabio Busacca, Stefano Mangione, Ilenia Tinnirello, Sergio Palazzo, and Francesco Restuccia. 2022. SDR-LoRa: dissecting and implementing LoRa on software-defined radios to advance experimental IoT research. In *Proceedings of the 16th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization* (Sydney, NSW, Australia) *(WiNTECH '22)*. Association for Computing Machinery, New York, NY, USA, 24–31. doi:10.1145/3556564.3558239

[9] Justin Chan, Anran Wang, Arvind Krishnamurthy, and Shyamnath Gollakota. 2019. DeepSense: Enabling Carrier Sense in Low-Power Wide Area Networks Using Deep Learning. arXiv:1904.10607 [cs.NI] https://arxiv.org/abs/1904.10607

[10] Qian Chen and Jiliang Wang. 2023. AlignTrack: Push the SNR Limit of LoRa Collision Decoding. *IEEE/ACM Transactions on Networking* 31, 5 (Oct 2023), 2070–2085. doi:10.1109/TNET.2023.3235041

[11] Brian Choi. 2023. RTX Video Super Resolution. https://blogs.nvidia.com/blog/rtx-video-super-resolution/ NVIDIA Blog. Accessed: 2025-11-15.

[12] R.E. Crochiere and L.R. Rabiner. 1983. *Multirate Digital Signal Processing*. Prentice-Hall.

[13] Shengkun Cui, Archit Patke, Hung Nguyen, Aditya Ranjan, Ziheng Chen, Phuong Cao, Brett Bode, Gregory Bauer, Catello Di Martino, Saurabh Jha, Chandra Narayanaswami, Daby Sow, Zbigniew T. Kalbarczyk, and Ravishankar K. Iyer. 2025. Characterizing GPU Resilience and Impact on AI/HPC Systems. arXiv:2503.11901 [cs.DC] https://arxiv.org/abs/2503.11901

[14] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. 2014. Learning a Deep Convolutional Network for Image Super-Resolution. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer International Publishing, Cham, 184–199.

[15] Adwait Dongare, Revathy Narayanan, Akshay Gadre, Anh Luong, Artur Balanuta, Swarun Kumar, Bob Iannucci, and Anthony Rowe. 2018. Charm: Exploiting Geographical Diversity through Coherent Combining in Low-Power Wide-Area Networks. In *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. 60–71. doi:10.1109/IPSN.2018.00013

[16] Jialuo Du, Yunhao Liu, Yidong Ren, Li Liu, and Zhichao Cao. 2024. LoRaTrimmer: Optimal Energy Condensation with Chirp Trimming for LoRa Weak Signal Decoding. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking* (Washington D.C., DC, USA) *(ACM MobiCom '24)*. Association for Computing Machinery, New York, NY, USA, 1104–1118. doi:10.1145/3636534.3690681

[17] Rashad Eletreby, Diana Zhang, Swarun Kumar, and Osman Yağan. 2017. Empowering Low-Power Wide Area Networks in Urban Settings. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication* (Los Angeles, CA, USA) *(SIGCOMM '17)*. Association for Computing Machinery, New York, NY, USA, 309–321. doi:10.1145/3098822.3098845

[18] Ettus Research. [n. d.]. USRP B205mini-i (1x1, 70 MHz - 6 GHz). Web Page. https://www.ettus.com/all-products/USRP-B205mini-i/ Accessed: 2025-11-16.

[19] Ettus Research. 2023. N200/N210. Ettus Knowledge Base. https://kb.ettus.com/N200/N210 Accessed: 2025-11-16.

[20] William Harlan. 1998. Avoiding interpolation artifacts in Stolt migration. *Stanford Exploration Project* (1998). Stanford Exploration Project Report 98 (1998), 241-247.

[21] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (11 1997), 1735–1780. arXiv:https://direct.mit.edu/neco/article-pdf/9/8/1735/813796/neco.1997.9.8.1735.pdf doi:10.1162/neco.1997.9.8.1735

[22] Intel Corporation. 2026. *Intel 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2*.

[23] Kerlink. [n. d.]. Wirnet iZeptoCell Serie - Indoor LoRaWAN Gateway. Web Page. https://www.kerlink.com/product/wirnet-izeptocell-serie/ Accessed: 2025-11-16.

[24] Walt Kester. 2008. *Taking the Mystery out of the Infamous Formula, "SNR = 6.02N + 1.76dB," and Why You Should Care*. Tutorial MT-001. Analog Devices. https://www.analog.com/media/en/training-seminars/tutorials/MT-001.pdf Rev. A.

[25] Walt Kester. 2009. *Oversampling Interpolating DACs*. Tutorial MT-017. Analog Devices. https://www.analog.com/media/en/training-seminars/tutorials/mt-017.pdf Rev. A (October 2008).

[26] Walt Kester. 2009. *Understand SINAD, ENOB, SNR, THD, THD + N, and SFDR so You Don't Get Lost in the Noise Floor*. Tutorial MT-003. Analog Devices. https://www.analog.com/media/en/training-seminars/tutorials/mt-003.pdf Rev. A (October 2008).

[27] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. 2016. Accurate Image Super-Resolution Using Very Deep Convolutional Networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1646–1654. doi:10.1109/CVPR.2016.182

[28] Adrian Kneip, Martin Lefebvre, Pol Maistriaux, and David Bol. 2025. IMAGINE: An 8-to-1b 22nm FD-SOI Compute-In-Memory CNN Accelerator With an End-to-End Analog Charge-Based 0.15-8POPS/W Macro Featuring Distribution-Aware Data Reshaping. *IEEE Transactions on Circuits and Systems for Artificial Intelligence* 2, 3 (Sep. 2025), 222–235. doi:10.1109/TCASAI.2025.3576323

[29] Apostolos Kokolis, Michael Kuchnik, John Hoffman, Adithya Kumar, Parth Malani, Faye Ma, Zachary DeVito, Shubho Sengupta, Kalyan Saladi, and Carole-Jean Wu. 2025. Revisiting Reliability in Large-Scale Machine Learning Research Clusters. arXiv:2410.21680 [cs.DC] https://arxiv.org/abs/2410.21680

[30] Chenning Li, Hanqing Guo, Shuai Tong, Xiao Zeng, Zhichao Cao, Mi Zhang, Qiben Yan, Li Xiao, Jiliang Wang, and Yunhao Liu. 2021. NELoRa: Towards Ultra-low SNR LoRa Communication with Neural-enhanced Demodulation. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems* (Coimbra, Portugal) *(SenSys '21)*. Association for Computing Machinery, New York, NY, USA, 56–68. doi:10.1145/3485730.3485928

[31] Chenning Li, Xiuzhen Guo, Longfei Shangguan, Zhichao Cao, and Kyle Jamieson. 2022. CurvingLoRa to Boost LoRa Network Throughput via Concurrent Transmission. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. USENIX Association, Renton, WA, 879–895. https://www.usenix.org/conference/nsdi22/presentation/li-chenning

[32] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. 2017. Enhanced Deep Residual Networks for Single Image Super-Resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.

[33] MikroTik. 2022. wAP LR2 kit. Web Page. https://mikrotik.com/product/wap_lr2 Accessed: 2025-11-16.

[34] Manan Mishra, Daniel Koch, Muhammad Osama Shahid, Bhuvana Krishnaswamy, Krishna Chintalapudi, and Suman Banerjee. 2023. OpenLoRa: Validating LoRa Implementations through an Extensible and Open-sourced Framework. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 1165–1183. https://www.usenix.org/conference/nsdi23/presentation/mishra

[35] Boris Murmann. [n. d.]. ADC Performance Survey 1997-2025. [Online]. Available: https://github.com/bmurmann/ADC-survey.

[36] Khang Nguyen, Yidong Ren, Jialuo Du, Jingkai Lin, Maolin Gan, Shigang Chen, Mi Zhang, Chunyi Peng, and Zhichao Cao. 2025. *LoRaSeek: Boosting Denoising Ability in Neural-enhanced LoRa Decoder via Hierarchical Feature Extraction*. Association for Computing Machinery, New York, NY, USA, 712–726. https://doi-org.remotexs.ntu.edu.sg/10.1145/3680207.3765241

[37] NVIDIA Corporation. 2026. *NVIDIA Management Library (NVML) API Reference Guide*. https://docs.nvidia.com/deploy/nvml-api/index.html Accessed: 2026-03-23.

[38] George Ostrouchov, Don Maxwell, Rizwan A. Ashraf, Christian Engelmann, Mallikarjun Shankar, and James H. Rogers. 2020. GPU lifetimes on titan supercomputer: survival analysis and reliability. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Atlanta, Georgia) *(SC '20)*. IEEE Press, Article 41, 14 pages.

[39] PyTorch Contributors. 2025. Post Training Quantization (PTQ). https://docs.pytorch.org/TensorRT/tutorials/ptq.html. v1.4.0+7d1d80773, Accessed: Nov 2025.

[40] Raghav Rathi and Zhenghao Zhang. 2022. TnB: resolving collisions in LoRa based on the peak matching cost and block error correction. In *Proceedings of the 18th International Conference on Emerging Networking EXperiments and Technologies* (Roma, Italy) *(CoNEXT '22)*. Association for Computing Machinery, New York, NY, USA, 401–416. doi:10.1145/3555050.3569132

[41] Kanav Sabharwal, Soundarya Ramesh, Jingxian Wang, Dinil Mon Divakaran, and Mun Choon Chan. 2024. Enhancing LoRa Reception with Generative Models: Channel-Aware Denoising of LoRaPHY Signals. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems* (Hangzhou, China) *(SenSys '24)*. Association for Computing Machinery, New York, NY, USA, 507–520. doi:10.1145/3666025.3699354

[42] Richard Schreier and Gabor C. Temes. 2017. *Understanding Delta-Sigma Data Converters* (2nd ed.). IEEE Press. doi:10.1002/9781119258308

[43] Seeed Studio. [n. d.]. *Wio-E5 LoRa Wireless Module - Powered by STM32WLE5 Datasheet*. Datasheet. Seeed Studio. https://files.seeedstudio.com/products/317990687/res/LoRa-E5%20module%20datasheet_V1.1.pdf Version 1.1. Accessed: 2025-11-16.

[44] Ltd. Seeed Technology Co. [n. d.]. LoRaWAN Gateway and Wireless Sensor Catalog, Version V1.6. Web Page. https://sensecap-docs.seeed.cc/pdf/LoRaWAN%

20Gateway%20and%20Wireless%20Sensor%20Catalog-V1.6.pdf Accessed: 2025-11-16.

[45] Semtech Corporation. 2020. *SX1276/77/78/79 - 137 MHz to 1020 MHz Low Power Long Range Transceiver*. Datasheet. Semtech Corporation. https://www.semtech.com/products/wireless-rf/lora-connect/sx1276 Accessed: 2025-11-16.

[46] Muhammad Osama Shahid, Millan Philipose, Krishna Chintalapudi, Suman Banerjee, and Bhuvana Krishnaswamy. 2021. Concurrent interference cancellation: decoding multi-packet collisions in LoRa. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference* (Virtual Event, USA) *(SIGCOMM '21)*. Association for Computing Machinery, New York, NY, USA, 503–515. doi:10.1145/3452296.3472931

[47] Seeed Studio. 2021. SenseCAP M1 LoRaWAN Indoor Gateway. Web Page. https://www.seeedstudio.com/SenseCAP-M1-LoRaWAN-Indoor-Gateway-US915-p-5023.html Accessed: 2025-11-16.

[48] The Things Industries. 2019. *The Things Indoor Gateway Datasheet*. Datasheet. The Things Industries. https://www.thethingsnetwork.org/docs/gateways/thethingsindoor/TTIG_datasheet.pdf Accessed: 2025-11-16.

[49] Shuai Tong, Jiliang Wang, and Yunhao Liu. 2022. Combating Packet Collisions Using Non-Stationary Signal Scaling in LPWANs. *IEEE/ACM Transactions on Networking* 30, 3 (June 2022), 1104–1117. doi:10.1109/TNET.2021.3131704

[50] Shuai Tong, Zhenqiang Xu, and Jiliang Wang. 2023. CoLoRa: Enabling Multi-Packet Reception in LoRa Networks. *IEEE Transactions on Mobile Computing* 22, 6 (June 2023), 3224–3240. doi:10.1109/TMC.2021.3138495

[51] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, Aditya Vijaykumar, Alessandro Pietro Bardelli, Alex Rothberg, Andreas Hilboll, Andreas Kloeckner, Anthony Scopatz, Antony Lee, Ariel Rokem, C. Nathan Woods, Chad Fulton, Charles Masson, Christian Häggström, Clark Fitzgerald, David A. Nicholson, David R. Hagen, Dmitrii V. Pasechnik, Emanuele Olivetti, Eric Martin, Eric Wieser, Fabrice Silva, Felix Lenders, Florian Wilhelm, G. Young, Gavin A. Price, Gert-Ludwig Ingold, Gregory E. Allen, Gregory R. Lee, Hervé Audren, Irvin Probst, Jörg P. Dietrich, Jacob Silterra, James T. Webber, Janko Slavič, Joel Nothman, Johannes Buchner, Johannes Kulick, Johannes L. Schönberger, José Vinícius de Miranda Cardoso, Joscha Reimer, Joseph Harrington, Juan Luis Cano Rodríguez, Juan Nunez-Iglesias, Justin Kuczynski, Kevin Tritz, Martin Thoma, Matthew Newville, Matthias Kümmerer, Maximilian Bolingbroke, Michael Tartre, Mikhail Pak, Nathaniel J. Smith, Nikolai Nowaczyk, Nikolay Shebanov, Oleksandr Pavlyk, Per A. Brodtkorb, Perry Lee, Robert T. McGibbon, Roman Feldbauer, Sam Lewis, Sam Tygier, Scott Sievert, Sebastiano Vigna, Stefan Peterson, Surhud More, Tadeusz Pudlik, Takuya Oshima, Thomas J. Pingel, Thomas P. Robitaille, Thomas Spura, Thouis R. Jones, Tim Cera, Tim Leslie, Tiziano Zito, Tom Krauss, Utkarsh Upadhyay, Yaroslav O. Halchenko, Yoshiki Vázquez-Baeza, and SciPy 1.0 Contributors. 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods* 17 (mar 2020), 261–272. Issue 3. doi:10.1038/s41592-019-0686-2

[52] Xiong Wang, Linghe Kong, Liang He, and Guihai Chen. 2019. mLoRa: A Multi-Packet Reception Protocol in LoRa networks. In *2019 IEEE 27th International Conference on Network Protocols (ICNP)*. 1–11. doi:10.1109/ICNP.2019.8888038

[53] Xianjin Xia, Qianwu Chen, Ningning Hou, Yuanqing Zheng, and Mo Li. 2023. XCopy: Boosting Weak Links for Reliable LoRa Communication. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking* (Madrid, Spain) *(ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 14, 15 pages. doi:10.1145/3570361.3592516

[54] Xianjin Xia, Ningning Hou, Yuanqing Zheng, and Tao Gu. 2023. PCube: Scaling LoRa Concurrent Transmissions with Reception Diversities. *ACM Trans. Sen. Netw.* 18, 4, Article 66 (March 2023), 25 pages. doi:10.1145/3545571

[55] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2020. FTrack: Parallel Decoding for LoRa Transmissions. *IEEE/ACM Transactions on Networking* 28, 6 (Dec 2020), 2573–2586. doi:10.1109/TNET.2020.3018020

[56] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. 2020. FlipLoRa: Resolving Collisions with Up-Down Quasi-Orthogonality. In *2020 17th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. 1–9. doi:10.1109/SECON48991.2020.9158432

[57] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. 2023. From Demodulation to Decoding: Toward Complete LoRa PHY Understanding and Implementation. *ACM Trans. Sen. Netw.* 18, 4, Article 64 (Jan. 2023), 27 pages. doi:10.1145/3546869

[58] Zhenqiang Xu, Pengjin Xie, and Jiliang Wang. 2021. Pyramid: Real-Time LoRa Collision Decoding with Peak Tracking. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*. 1–9. doi:10.1109/INFOCOM42981.2021.9488695

[59] Zichuan Yu, Hongyu Ge, Xusheng Tang, and Lu Tang. 2025. A Frequency Offset Decoupled Radio Frequency Fingerprint Identification Scheme for LoRa Devices. *IEEE Internet of Things Journal* 12, 23 (Dec 2025), 51358–51371. doi:10.1109/JIOT.2025.3612500