Mitigating Interleaved ADC Mismatches in Automatic Modulation Classification via Sampling Schemes and Layer Fusion

Alexander Gros
Electromagnetism and Telecommunications Department
Faculty of Engineering, UMONS
Mons, Belgium
alexander.gros@umons.ac.be

Véronique Moeyaert
Electromagnetism and Telecommunications Department
Faculty of Engineering, UMONS
Mons, Belgium
veronique.moeyaert@umons.ac.be

Patrice Mégret
Electromagnetism and Telecommunications Department
Faculty of Engineering, UMONS
Mons, Belgium
patrice.megret@umons.ac.be

***Abstract*—Time Interleaved Analog to Digital Converters (TI-ADCs) are widely used in wideband spectrum monitoring systems to achieve high sampling rates and large instantaneous bandwidths. However, as highlighted in the literature, accurately compensating for inter-channel mismatches such as gain, offset, and timing skew, is critical, as these impairments can distort the digitized signal. While traditional studies focus on spectral analysis under ideal sinusoidal inputs, this work examines the direct impact of TI-ADC mismatches on Automatic Modulation Classification (AMC) accuracy. Using a publicly available dataset of IQ-modulated signals, we introduce controlled mismatch impairments and evaluate classification performance across varying distortion levels, from relaxed to extreme. We propose a lightweight dyadic down-sampling scheme combined with a CNN-based fusion architecture that improves robustness of AMC to interleaving mismatches and calibration drift. Results show up to 6.7 percent improvement in classification accuracy under severe mismatch conditions, underscoring the method's consistent gains and robustness to hardware distortions. These findings are particularly relevant for field-deployed or long-duration monitoring applications where hardware degradation and limited recalibration are expected.**

***Index Terms*—AMC, AMR, automatic modulation recognition, classification, cognitive radios, CNN, dyadic, fusion, interleaved ADC, mismatches, skew.**

I. Introduction

Automatic Modulation Recognition (AMR) is the task of analyzing received signals to determine their modulation schemes. A key subset of this process, known as Automatic Modulation Classification (AMC), focuses on identifying the modulation type of an observed Radio Frequency (RF) signal based on its characteristics at a specific moment, frequency, and location. Typically, AMC takes place after the signal has been detected but before it undergoes demodulation. In modern wireless systems, while Software-Defined Radios (SDRs) are not typically used as core components in commercial infrastructure, the flexibility they represent has influenced the design of adaptive transceivers. These systems can adjust modulation schemes dynamically in response to channel conditions or system requirements to optimize performance and maintain Quality of Service (QoS). In this context, reliable modulation classification becomes important, allowing receivers, particularly in monitoring or cognitive radio scenarios, to detect these changes without relying on explicit signaling or metadata exchange. To accurately perform automatic modulation classification, the quality and fidelity of the digitized signal play a crucial role. However, as system requirements push towards higher sampling rates and wider bandwidths, practical constraints in analog-to-digital conversion become more pronounced. Time-interleaved ADC architectures have emerged as a popular solution to achieve these demanding specifications. However, mismatches introduced by time interleaving, including timing skew, gain mismatch, and offset, may degrade the quality of the digitized signal and impair modulation classification, which we examine in this work. Understanding these imperfections and their impact is therefore essential to design robust recognition systems. As demonstrated in our previous work [1], combining dyadic down-sampling with a CNN-based architecture that learns task-specific filters leads to significant improvements in modulation classification accuracy. In this study, we further show that this approach is inherently robust to interleaved ADC mismatches, calibration drift, and aging effects. This makes this process a lightweight and practical solution for near real-time, wideband signal monitoring with minimal computational overhead.

In this paper, Section II reviews the principles of interleaved ADC operation, the dataset used, and the adopted mismatch characterization methodology. Section III presents the dyadic

979-8-3315-0292-8/25/$31.00 ©2025 IEEE

down-sampling technique, outlines the overall dataflow, and details the proposed AI-based architecture. Finally, Section IV discusses the experimental results, evaluates the advantages and limitations of the proposed approach, and concludes the paper.

II. Overview of High-Speed Analog-to-Digital Conversion

A. Interleaved ADC Architectures

To overcome the limitations imposed by individual ADC devices, system designers increasingly adopt interleaved ADC structures. These architectures employ multiple ADC cores operating in parallel to collectively increase the effective sampling rate or enhance resolution. The approach is particularly effective as improvements in converter miniaturization, power efficiency, and cost have made multi-channel configurations more feasible [2] [3]. Two primary interleaving strategies are commonly employed [4]: signal averaging, aimed at enhancing the signal-to-noise ratio (SNR), and time interleaving, which improves throughput. In this paper, we focus on time interleaving, a technique in which M ADC channels are configured to sample the input signal sequentially, with each channel phase-shifted in time as seen in Fig.1. This effectively multiplies the system's sampling rate by M, with each ADC operating at a M times reduced rate, thereby alleviating the limitations of individual devices. The appropriate clock phase offset for each channel can be computed using the following expression [5]:

$$\phi_m = 2\pi \left(\frac{m-1}{M} \right), \tag{1}$$

where m denotes the channel index (ranging from 1 to M), and M the total number of interleaved ADCs. The principal advantage of time interleaving lies in its ability to expand the Nyquist bandwidth of the converter system. In addition to bandwidth benefits, interleaved ADCs support the development of adaptable and scalable SDR platforms. These systems can accommodate new modulation schemes and frequency bands through software updates, rendering them highly adaptable to evolving standards. Despite these benefits, interleaved architectures introduce substantial design complexity. Even minor imperfections in clock distribution can introduce timing skew, degrading system performance, leading in the context of this paper to a lower AMC accuracy.

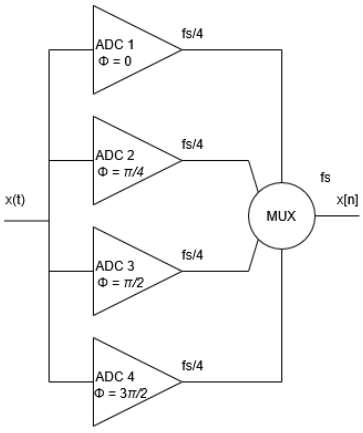

Fig. 1. Block diagram of a time-interleaved ADC system with 4 parallel sub-ADCs, each sampling at $\frac{f_s}{4}$, clock distribution not drawn

Fig. 2 shows the power spectral density (PSD) of a time-interleaved ADC system with $M = 4$ channels, each operating at one-fourth of the overall sampling rate. While interleaving increases the effective sampling frequency, it also introduces spurious spectral components due to gain mismatches among the channels (e.g., 1.000, 1.015, 1.010, 1.012). These mismatches produce image replicas at frequencies given by $f_{\text{spur}} = \pm f_{\text{in}} + \frac{k}{M} f_s$, where $k = 1, 2, 3, \ldots$ where f_{in} is the frequency of the input sinusoidal signal, and f_s is the total system sampling rate [6]. Frequencies in the spectrum are normalized with respect to a baseband carrier (i.e., carrier frequency set to 0 Hz). In the example of Fig. 2, this results in spurious components at the normalized frequencies 0.15, 0.35, and 0.40 Hz.

B. Dataset Description

Extensive validation, prototyping, and system-level testing are therefore essential to translate the theoretical benefits of interleaving into real-world performance gains. However, exploring the full range of mismatch effects in a controlled yet comprehensive manner calls for a synthetic dataset that can systematically vary key parameters. The dataset employed for both training and evaluation in this study is the CSPB.ML.2018 dataset [7], originally released for a machine learning challenge. It comprises synthetically generated IQ signals created in MATLAB, with the complete set of parameter ranges listed in Table I. Notably, this dataset incorporates often-overlooked parameters such as Carrier Frequency Offset (CFO) and the Root Raised Cosine roll-off factor (RRC). It comprises only digitally modulated signals, including BPSK, QPSK, 8PSK, $\frac{\pi}{4}$-DQPSK, 16QAM, 64QAM, 256QAM, and MSK. Apart from Additive White Gaussian Noise (AWGN), which introduces different SNR conditions, no channel impairments were applied. SNR levels span from -2 dB to 12.8 dB, with the highest concentration around 10 dB. Each signal is relatively long compared to other available databases (32,678 samples), with a total of 112,000 waveforms pro-

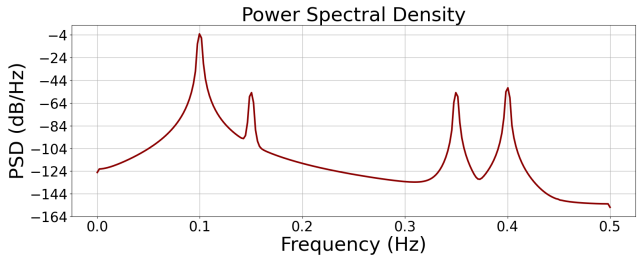

Fig. 2. Effect of gain mismatch on a sinusoidal input at 0.1 times the normalized sampling frequency and four interleaved ADCs

TABLE I
CSPB.ML.2018 DATASET PARAMETERS [7]

Parameter	Range / Values
Modulations	BPSK, QPSK, $\frac{\pi}{4}$-DQPSK, 8PSK, MSK, 16QAM, 64QAM, 256QAM
Base symbol period	1–15 samples
Carrier freq. offset (norm.)	$[-10^{-3}, 10^{-3}]$ Hz
Roll-off factor	0.1–1
SNR	–2 to 12.8 dB
Up/down sampling	(1,1), (3,2),(4,3)→(10,9)
Noise spec. density	0 dB
Signal length	32,678 samples
Number of waveforms	112,000

vided. In our work, 70% of the data is used for training, with the remaining 30% reserved for validation.

C. Post-Processing - applying mismatches to existing dataset

Unlike traditional methods that introduce mismatches during signal generation, we apply interleaved ADC mismatches as a post-processing step to the pre-generated in-phase and quadrature (IQ) waveforms from the previous section. This approach offers fine control over each mismatch parameter and enables independent evaluation on existing datasets. In radio receivers, the choice between direct RF sampling and heterodyne architectures significantly influences how mismatches in time-interleaved ADCs manifest. In direct sampling systems, the RF signal is sampled directly by the ADCs without prior down-conversion [8], making mismatch effects such as gain error, offset, and especially timing skew immediately visible in the sampled passband spectrum. These impairments introduce spurious tones at predictable frequency offsets, such as multiples of $\frac{f_s}{M}$, where f_s is the effective sampling rate and M the number of interleaved channels. In contrast, heterodyne receivers perform frequency down-conversion of the RF signal to an intermediate frequency (IF) prior to digitization. A related architecture, often called direct-conversion or homodyne, mixes the RF signal directly to baseband instead of an IF. Both approaches reduce the input frequency seen by the ADC compared to direct RF sampling, which changes how interleaving mismatches manifest in the digitized spectrum. In such architectures, mismatches that occur after the down-conversion stage, particularly within the IQ baseband paths, generally have a diminished effect on system performance. This is because, once the signal has been translated to a lower frequency, the resulting distortions are spectrally located closer to DC, where they can be more readily filtered or digitally compensated. Consequently, the impact of these mismatches on the overall spectral purity is typically limited, unless the mismatches are excessively large. Therefore, to accurately evaluate mismatch-induced distortion in time-interleaved ADC systems, simulations should apply the impairments in the passband domain prior to IQ extraction, preserving their true spectral behavior. This methodology provides a versatile tool for simulating a wide range of ADC architectures and mismatch scenarios. In this work, the simulation models a direct RF sampling architecture with digital IQ downconversion, a configuration commonly used in military receivers, wideband software defined radios, and spectrum monitoring systems, to evaluate the impact of time interleaving imperfections. The baseband I/Q signal is modulated to the desired RF passband, where interleaved ADC impairments including gain, offset, and timing skew are applied independently to each channel. The distorted passband waveform is then digitally downconverted to complex baseband, enabling accurate assessment of signal quality and spectral artifacts. A passband frequency of 10 MHz was chosen to represent a typical narrowband RF scenario while keeping computational requirements manageable. At this frequency, timing skew produces moderate phase errors, while gain and offset mismatches remain largely independent of frequency. By applying impairments in the passband before digital IQ extraction, the simulation preserves the true spectral behavior of mismatch induced distortions. Overall, the methodology provides realistic insights into interleaved ADC mismatches while remaining general enough to evaluate a wide range of sampling architectures and digital downconversion strategies. We consider an M-channel time-interleaved ADC. For each waveform, the original in-phase and quadrature components are denoted by $x_I[n]$ and $x_Q[n]$. Gain error, DC offset, and timing skew are applied independently to each sub-ADC $m \in 0, 1, \ldots, M-1$. The corrupted outputs are defined at positions $n = m + kM$ (with $k \in \mathbb{Z}$) as follows:

$$\begin{aligned} y_I[n] &= g_{I,m} \cdot x_I(n + \delta_{I,m}) + o_{I,m}, \\ y_Q[n] &= g_{Q,m} \cdot x_Q(n + \delta_{Q,m}) + o_{Q,m}, \end{aligned} \tag{2}$$

where:

- $g_{I,m}, g_{Q,m}$ are the gain mismatches for sub-ADC m,
- $o_{I,m}, o_{Q,m}$ are the offset mismatches,
- $\delta_{I,m}, \delta_{Q,m}$ are the sampling skews.

The expressions $x_I(n+\delta_{I,m})$ and $x_Q(n+\delta_{Q,m})$ are evaluated using cubic interpolation to approximate fractional shifts in time. After applying these transformations, the corrected sub-sequences are reassembled into a complete waveform. The proposed representation uses per channel gain, offset, and timing skew to model mismatches across a wide range of ADC architectures, including direct sampling, real IF down-conversion, and complex IQ down-conversion. This digital domain model captures the primary ADC impairments, while effects originating from the analog front end, such as mixer nonlinearities, RF filtering, or analog IQ imbalance, are not considered in this formulation. In particular, IQ imbalance between the I and Q paths of each interleaved ADC was tested in our simulations but did not produce a significant impact on the results for the chosen test configuration.

D. Selection of mismatch parameters

To realistically assess the performance of time interleaved ADC systems, we simulate gain, offset, and timing mismatch errors across M parallel channels. These mismatches are modeled as zero mean Gaussian random variables, with standard deviations chosen to reflect a wide range of system conditions. Table II summarizes representative values for four

levels of severity: relaxed, moderate, severe, and extreme [6] [8]. These categories span use cases from highly controlled laboratory setups to worst case, uncalibrated field deployments. As an illustrative example, the moderate configuration models a typical calibrated or semi-calibrated system with realistic imperfections. In this case, gain mismatch has a standard deviation of 0.5%, consistent with common post-fabrication calibration performance. Offset mismatch is simulated with a standard deviation of 1 least significant bit (LSB), assuming an 8-bit ADC and a normalized full-scale range of ± 1 V, yielding $\frac{2}{256}$. Timing skew is expressed as a percentage of the sampling period T_s, with a standard deviation of 0.3%, capturing moderate inter-channel misalignment. These values serve as a balanced reference point, though all severity levels listed in Table II are evaluated in our experiments.

TABLE II
MISMATCH STANDARD DEVIATIONS FOR DIFFERENT SYSTEM CONDITIONS

Mismatch	Relaxed	Moderate	Severe	Extreme
Gain (%)	0.10	0.50	1.50	15.0
Offset (LSB)	0.25	1.00	2.00	10.0
Timing Skew (% T_s)	0.05	0.30	1.00	20.0

E. Normalization effect

Normalization is applied per waveform as is standard practice in deep learning, each waveform is normalized during preprocessing, a step that consistently yields the best performance. Normalization rescales each waveform to a fixed amplitude range and typically removes its DC component. As a result, absolute offset values, such as those introduced by adding a constant bias to a sinusoidal signal, are eliminated. The waveform is re-centered, and its dynamic range is adjusted, effectively restoring the original signal shape without the absolute offset. However, normalization does not remove relative offset mismatches between channels in a time-interleaved ADC. Although compression reduces the signal's overall distortion magnitude, interleaving artifacts such as alternating steps from per-channel offset differences still persist in the normalized waveform, though at reduced intensity. While normalization helps mitigate some mismatch effects, further reduction is achieved through dyadic down-sampling. The next section introduces this approach.

III. DYADIC SAMPLING AND PROCESSING PIPELINE

A. The Role of Dyadic Down-Sampling in Minimizing ADC Mismatch Effects

Dyadic down-sampling refers to the process of reducing the sampling rate of a signal by factors of two, where the term "dyadic" indicates a relation to powers of two. Formally, for a 1D signal $x[n]$, the down-sampled version at dyadic level d is defined as:

$$x_d[n] = x[2^d n] \tag{3}$$

This method allows the representation of a signal at multiple resolutions, where each level retains every 2^d-th sample of the original. For 2D signals, such as images or IQ matrices, the down-sampling extends naturally:

$$x_d[m, n] = x[2^d m, 2^d n] \tag{4}$$

In this work, each waveform from the dataset described in II-B is represented as a 2-row matrix containing the in-phase (I) and quadrature (Q) components:

$$x_{dI}[n],\ x_{dQ}[n] = x_I[2^d n],\ x_Q[2^d n] \tag{5}$$

This process reduces signal dimensionality while preserving essential structure. It effectively suppresses high-frequency noise and redundancy, offering a more compact representation for neural network training [1]. Indeed, in high-speed acquisition systems such as those employing time-interleaved ADCs, signals are often over-sampled to capture a broad bandwidth. After band-selection and down-conversion, the resulting signal contains significant temporal redundancy, which can be effectively reduced via dyadic down-sampling to yield a more compact and informative representation for learning tasks. The resulting signal variants at multiple dyadic levels provide the AI model with complementary views of the same waveform, enhancing its ability to capture relevant patterns across scales in a computationally efficient manner.

A computational form of this operation is:

```
signal[:, :, ::df]
```  (6)

where `df` is the decimation factor and `df` $= 2^d$, with d denoting the dyadic level or scale. Here, the first index corresponds to waveform selection, the second index selects the in-phase (I) and quadrature (Q) components, and the third index performs sample selection according to the decimation factor. Careful selection of the down-sampling factor is essential: excessive reduction may discard critical signal features, whereas moderate decimation can enhance computational efficiency and model performance. Using the scales, we not only extract features at different resolutions but this process also implicitly selects samples from a subset of the interleaved ADCs, as illustrated in Fig. 3. After the first dyadic down-sampling step, the selected samples correspond to those within the marked squares, specifically from ADCs 1 and 3. In the second down-sampling stage, only ADC 1 remains, with its selected samples marked by circles. As a result, gain mismatches across ADCs become less apparent due to the progressive exclusion of certain channels.

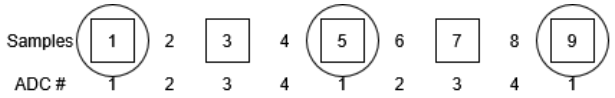

Fig. 3. Selection of dyadic samples in a TIADC system

B. Data flow

The complex IQ signal is extracted from the dataset (II-B) and then subjected to simulated ADC mismatches, including gain, offset, and timing errors. Following this, the signal undergoes dyadic down-sampling, decomposing it into multiple temporal scales. Each resulting complex sub-signal, along with the original input, is formatted as a two-row vector and fed into separate CNN branches of an AI architecture for modulation classification (see Fig. 4). The methodology is illustrated in Fig. 5. When using classical IQ reference system [9], only a single arm processes the original IQ samples. In Section IV-B1, the arms corresponding to the dyadic down-sampling scales are replaced with the individual ADC outputs. A data generator is used to stream waveform batches of size 32, minimizing RAM memory usage. When used, the generator performs dyadic down-sampling on the fly, just before feeding the data into the model, rather than relying on precomputed samples. Additionally, data is shuffled between epochs to promote generalization and reduce the risk of overfitting. In the tables, the labels "Gain," "Offset," "Skew," and "All" denote isolated mismatch types whereas "All" refers to all three applied simultaneously.

C. Automatic Modulation Recognition

Automatic Modulation Classification (AMC) involves identifying the modulation scheme of a received signal, typically formulated as a multi-class classification problem. In this study, we employ a deep learning approach based on Convolutional Neural Networks (CNNs), which are well-suited for learning local patterns in time-series data [10]. A convolutional neural network (CNN) layer can be interpreted as a learned generalization of a finite impulse response (FIR) filter. Like an FIR filter, a CNN layer performs a convolution operation over an input signal, computing a weighted sum of local input samples using a fixed-size kernel. However, whereas FIR filter coefficients are typically designed to meet specific frequency-domain criteria (e.g., low-pass characteristics), the CNN's filter weights are learned directly from data through gradient-based optimization. This allows CNNs to automatically discover task-specific filters that extract meaningful features from input signals. In essence, a CNN layer functions can be interpreted as a bank of adaptive FIR-like filters, typically followed by nonlinearities, enabling the modeling of complex patterns in time-series or spatial data. In our approach, we down-sample first and rely on the CNN to perform the filtering instead of using a traditional anti-aliasing filter. The learned kernels exhibit pass band characteristics, a common trait in CNNs, which means aliasing is present but is effectively taken into account by the network during training. The architecture consists of several key components that can be found in Table III: convolutional layers that extract hierarchical features using learnable filters, followed by dense (fully connected) layers that integrate these features for final decision-making. The convolutional layers utilize filter sizes of 2×8 and 1×4 in the first and second layers, respectively. These are followed by non-linear ReLU (Rectified Linear Unit) activations, dropout layers for regularization, and flattening operations to transition from convolutional to dense layers. The output layer is a dense layer of size 8, corresponding to the number of modulation classes present in the dataset, and concludes with a softmax activation function to generate class probabilities. Having outlined the system architecture and processing flow, we now present and analyze the experimental results that validate the effectiveness of the proposed approach.

TABLE III
CNN ARCHITECTURE SUMMARY

Layer Type	Kernel / Units	Activation
Input	(batch size, samples, 1)	–
Conv2D	(2, 8), 40	ReLU
Conv2D	(1, 4), 10	ReLU
Flatten	–	–
Dropout	50%	
Dense	256	ReLU
Dense (output)	8	Softmax

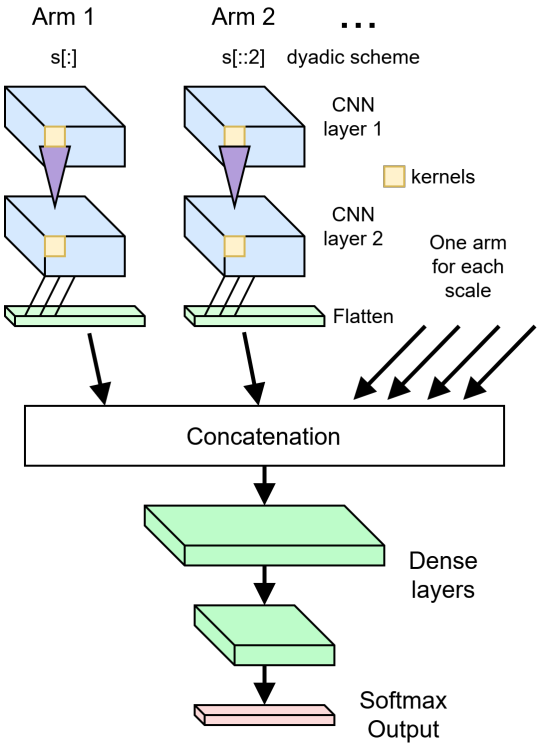

Fig. 4. Global overview of the CNN layer fusion architecture

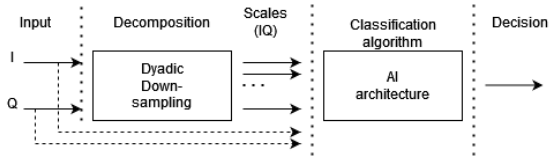

Fig. 5. Overall data-flow diagram

IV. Results-Discussion-Conclusion

A. Results

A comparative analysis between the original IQ method and dyadic down-sampling at $M = 2$ and $M = 4$ demonstrates consistent improvements across all evaluated metrics with dyadic down-sampling, as shown in Tables IV and V. Relative improvements range from approximately 4.6% to 8.5%, with the most notable gains in category *GAIN* and *ALL* scores. Absolute increases in accuracy range from 4.12% to 6.59%. To assess the impact of ADC mismatches on system performance, accuracy can be evaluated as a function of the input signal-to-noise ratio (SNR). At low SNR levels, the performance is primarily limited by noise, which masks the effects of mismatches such as gain or offset errors. However, as SNR increases, these mismatch-induced distortions become more prominent relative to the noise floor, potentially degrading classification accuracy. This approach effectively distinguishes noise-limited from mismatch-limited regimes and highlights the importance of calibration in high-SNR scenarios. The code, additional results, and supplementary tests related to this study will be made available on GitHub, as referenced in [11].

B. Discussion

1) Exploring ADC Output as an Alternative to Dyadic Down-sampling: Instead of using a downsampled representation as in the dyadic approach, we retain the same fusion architecture but modify the input strategy. Specifically, the model receives both the full-resolution input waveform and the individual signals from each ADC channel. The motivation is that, by exposing the network to the raw per-channel signals, any mismatch effects are explicitly preserved rather than hidden in the composite waveform, making them easier for the model to learn and correct during training. To assess overall performance differences between the dyadic and ADC-based methods, we computed the average percentage improvement (API) across all four metrics (GAIN, OFFSET, SKEW, and ALL) using:

$$\text{API} = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{\text{Dyadic}_i - \text{ADC}_i}{\text{ADC}_i} \times 100\right) \tag{7}$$

where $N = 4$ is the number of encompassed metrics. As summarized in Table VI, the dyadic method outperformed the ADC-based approach by an average API of 4.65% at $M = 2$ and 3.48% at $M = 4$. These results suggest that dyadic down-sampling provides a consistent advantage across metrics. Moreover, the training process and overall architecture of the ADC-based method must be adapted to the number of ADC channels in the system, which limits its flexibility. In contrast, the dyadic methodology is inherently more adaptable across different system configurations.

TABLE IV
COMPARISON OF AVERAGE (10 TRIALS) CLASSIFICATION SCORES BETWEEN ORIGINAL IQ AND DYADIC DOWN-SAMPLING AT $M = 2$ AND $M = 4$, SEVERE CASE.

Metric	Original IQ (M=2)	Dyadic (M=2)	Original IQ (M=4)	Dyadic (M=4)
NONE	0.9049	0.9464 (+4.15%)	0.9049	0.9464 (+4.15%)
GAIN	0.9029	0.9458 (+4.29%)	0.9035	0.9462 (+4.27%)
OFFSET	0.9048	0.9463 (+4.15%)	0.9049	0.9466 (+4.17%)
SKEW	0.9049	0.9465 (+4.16%)	0.9049	0.9466 (+4.17%)
ALL	0.9028	0.9455 (+4.27%)	0.9035	0.9461 (+4.26%)

TABLE V
COMPARISON OF AVERAGE (10 TRIALS) CLASSIFICATION SCORES BETWEEN ORIGINAL IQ AND DYADIC DOWN-SAMPLING AT $M = 2$ AND $M = 4$, EXTREME CASE.

Metric	Original IQ (M=2)	Dyadic (M=2)	Original IQ (M=4)	Dyadic (M=4)
NONE	0.9049	0.9464 (+4.15%)	0.9049	0.9464 (+4.15%)
GAIN	0.7837	0.8491 (+6.54%)	0.8411	0.9069 (+6.58%)
OFFSET	0.9021	0.9433 (+4.12%)	0.9034	0.9447 (+4.13%)
SKEW	0.9015	0.9477 (+4.62%)	0.8980	0.9439 (+4.59%)
ALL	0.7768	0.8427 (+6.59%)	0.8337	0.9009 (+6.72%)

TABLE VI
COMPARISON OF AVERAGE CLASSIFICATION SCORES BETWEEN ORIGINAL IQ AND ADC-BASED SCALING AT $M = 2$ AND $M = 4$, EXTREME CASE.

Metric	Original IQ (M=2)	ADC (M=2)	Original IQ (M=4)	ADC (M=4)
GAIN	0.7837	0.8006 (+2.16%)	0.8411	0.8925 (+6.11%)
OFFSET	0.9021	0.9358 (+3.74%)	0.9034	0.9373 (+3.75%)
SKEW	0.9015	0.9385 (+4.10%)	0.8980	0.9356 (+4.19%)
ALL	0.7768	0.7964 (+2.52%)	0.8337	0.8845 (+6.09%)

C. Conclusion

For wideband spectrum monitoring, particularly in the context of modulation classification, our simulations yield the following insights. In well-calibrated interleaved ADC systems with minimal mismatches, the impact on classification accuracy is negligible under both relaxed and moderate cases, particularly when using a database with difficult channel conditions that tend to mask residual mismatch effects. This is an encouraging result, as it indicates that advanced monitoring systems, such as those employing high-performance single-chip interleaved ADCs, can fully leverage their hardware capabilities without significant degradation in classification performance. Conversely, scenarios characterized by substantial mismatches due to device aging, prolonged operation without recalibration, or hardware variability in cost-constrained systems pose significant challenges. These conditions often lead to increased spurious artifacts and degraded signal integrity due to the lack of precise calibration and synchronization. Consequently, effective wideband classification in such environments demands more sophisticated digital signal processing techniques. Furthermore, our results demonstrate that the proposed dyadic down-sampling technique, when combined with a CNN-based fusion architecture, is particularly effective in mitigating the impact of gain mismatches on modulation classification performance in time-interleaved ADC systems. Moreover, when additional IQ imbalance is introduced, I has mismatches, Q has not, the inherent structure of the AI architecture helps to mitigate its impact. This is primarily due to the use of line wise convolutional kernels instead of chunk wise processing, which makes the imbalance effect negligible

in terms of classification accuracy. The modular nature of the proposed framework allows straightforward extension to other impairments such as IQ imbalance, jitter, or analog front-end drift, making it suitable for robust classification in practical wideband receivers.

ACKNOWLEDGMENT

This work is supported by the Walloon Region research project "CyberExcellence", n° 2110186. Computational resources have been provided by the "Consortium des Equipements de Calcul Intensif (CECI)", funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S.- FNRS) under Grant n° 2.5020.11, and by the Walloon Region.

REFERENCES

[1] A. Gros, V. Moeyaert, and P. Megret, "Enhancing modulation classification through lightweight dyadic down-sampling schemes and cnn layer fusion," in *2025 IEEE International Black Sea Conference on Communications and Networking (BlackSeaCom)*, 2025.

[2] Analog Devices, "Ad9625: 12-bit, 2.6 gsps jesd204b, dual analog-to-digital converter," https://www.analog.com/en/products/ad9625.html, 2020, datasheet.

[3] Texas Instruments, "ADC12DJ4000RF: Dual 12-Bit, 4 GSPS, RF Sampling ADC," https://www.ti.com/lit/ds/symlink/adc12dj4000rf.pdf, 2020, accessed: 2025-05-28.

[4] S. Louwsma, A. van Tuijl, and B. Nauta, *Time-Interleaved Analog to Digital Converters*. Germany: Springer, Oct. 2010, 10.1007/978-90-481-9716-3.

[5] J. Harris, "The abcs of interleaved adcs," *Analog Devices*, 2012. [Online]. Available: https://www.analog.com/en/resources/technical-articles/the-abcs-of-interleaved-adcs.html

[6] N. Kurosawa, H. Kobayashi, K. Maruyama, H. Sugawara, and K. Kobayashi, "Explicit analysis of channel mismatch effects in time-interleaved adc systems," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 48, no. 3, pp. 261–271, 2001.

[7] C. Spooner, "Dataset for the machine-learning challenge [cspb.ml.2018]," https://cyclostationary.blog/2019/02/15/data-set-for-the-machine-learning-challenge/, accessed: 07.06.2024.

[8] T. Instruments, "Interleaving basics: Sample rate vs. conversion rate," Texas Instruments, Tech. Rep. SLAA824, 2017. [Online]. Available: https://www.ti.com/lit/an/slaa824/slaa824.pdf

[9] T. J. O'Shea, T. Roy, and T. C. Clancy, "Over-the-air deep learning based radio signal classification," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 168–179, 2018.

[10] Z. Wang, W. Yan, and T. Oates, "Time series classification from scratch with deep neural networks: A strong baseline," 2016.

[11] A. Gros, "Tiadc-amc," https://github.com/AlexanderGros/TIADC-AMC, 2025, [Source code].