Visual Signatures of Suboptimal Sampling in Constellation-Based AMC with CNNs

Alexander Gros[1], Véronique Moeyaert[1], and Patrice Mégret[1]

[1]Electromagnetism and Telecommunication Department, Faculty of Engineering FPMs, UMONS, Mons, Belgium
Emails: {veronique.moeyaert,patrice.megret}@umons.ac.be
Correspondence: alexander.gros@umons.ac.be

***Abstract*—Automatic modulation classification (AMC) based on convolutional neural networks (CNNs) has shown strong performance using constellation diagram representations. However, the effect of symbol timing mismatch on these representations remains insufficiently studied. In this paper, we analyze how suboptimal sampling alters the structure of constellation diagrams and impacts the performance of learning-based classifiers. A peak-to-average power ratio (PAPR)-based approach is used to estimate the optimal sampling phase, and controlled timing offsets are introduced to systematically evaluate their effect. Both grayscale and RGB histogram-based constellations are considered and processed using lightweight CNN architectures. Results obtained on the CSPB.ML.2018 dataset demonstrate that even small timing offsets induce structured distortions and lead to significant performance degradation. These results highlight symbol timing as a critical factor in AMC and underline the importance of consistent sampling strategies for robust AI-based receivers.**

***Index Terms*—AMC, AMR, automatic modulation recognition, classification, cognitive radios, constellations, CNN, sampling mismatches, vector diagram.**

I. Introduction

Automatic modulation classification (AMC) is a key component in modern wireless communication systems, allowing receivers to determine the modulation scheme of an incoming signal without prior knowledge of the transmitter configuration. Such capability is particularly important in applications including cognitive radio, spectrum surveillance, interference management, and electronic warfare, where communication devices must operate in dynamic and potentially unknown environments. Accurate identification of the modulation format enables adaptive receivers to automatically adjust processing parameters, improving therefore spectral efficiency, and maintaining reliable communication in heterogeneous radio scenarios. Recent progress in machine learning, especially deep learning, has significantly advanced the performance of AMC systems. Instead of relying on manually designed features, deep neural networks can learn relevant signal characteristics directly from data. Among the various signal representations that have been investigated, constellation diagrams provide a compact and physically interpretable description of modulation behavior in the complex plane. When used as input to convolutional neural networks (CNNs), these representations allow efficient extraction of spatial patterns while keeping the computational complexity low compared to methods that operate directly on raw in-phase and quadrature samples. Despite the extensive body of work on AMC, most existing studies focus on robustness to channel impairments such as noise, carrier frequency offset, or phase rotation. In contrast, the impact of symbol timing mismatches on constellation-based representations remains largely unexplored, even though such mismatches naturally arise in practical receivers due to imperfect synchronization and sampling uncertainty. Since constellation diagrams are directly constructed from sampled IQ data, even small deviations from the optimal sampling instant can induce structured distortions that significantly affect the resulting representation. This paper addresses this gap by systematically analyzing the effect of sampling offsets on constellation-based AMC. In particular, we investigate how deviations from the optimal sampling instant alter the geometry and density distribution of constellation diagrams, and how these distortions impact the performance of convolutional neural network classifiers.

The main contributions of this paper are summarized as follows:

- A systematic and human-interpretable analysis of the impact of sampling offsets on IQ constellation diagrams, highlighting the structured distortions induced by suboptimal sampling.
- A comparative evaluation of histogram-based constellation representations, including grayscale and RGB encodings, with respect to their robustness to sampling mismatches.
- A simulation study demonstrating the effect of training and inference mismatch in sampling strategies, providing practical guidelines for the design of robust AI-based AMC systems.

This paper reviews the state of the art in AMC in Section II, with a particular emphasis on learning-based and constellation-driven approaches. The proposed methodology is presented in Section III, while the simulation setup is described in Sections IV, V and VI. Finally, the performance evaluation and discussion of results are provided in Section VII.

II. State of the Art

Automatic modulation classification (AMC) is a well-established research area.In recent years, deep learning has renewed interest in AMC by shifting from handcrafted signal features to data-driven representations. Several surveys, including those by Zheng et al. [1] and Thakur et al. [2], highlight the growing role of artificial intelligence and report consistent performance gains, particularly with convolutional and recurrent neural networks. Within this paradigm, constellation diagram representations have attracted attention due to their intuitive visualization of modulation patterns. Doan et al. [3] map received signals to grayscale constellation images for CNN-based classification, although such single-channel representations may limit feature richness. To overcome this, multi-channel adaptations have been proposed. For instance, Leblebici et al. [4] generate three-channel images to leverage standard CNN architectures, but these approaches often rely on deep and computationally expensive models. Other works enhance constellation representations through additional preprocessing. Abdel-Moneim et al. [5] employ Gabor filtering and thresholding to produce RGB images, improving feature quality at the cost of increased preprocessing complexity, which may limit real-time or resource-constrained applications. In addition, [6] proposed a CNN-based transformation scheme to improve automatic modulation classification under severe phase imperfections, demonstrating enhanced robustness and accuracy in challenging wireless conditions. Overall, while deep learning-based AMC methods using constellation diagrams achieve strong performance, they often trade computational efficiency for accuracy through complex preprocessing or large network architectures. This motivates the search for approaches that preserve high accuracy while reducing computational overhead and system complexity. However, none of these works explicitly analyze the impact of symbol timing mismatches on modulation classification using constellation-based representations.

III. Dataset and Information flow

A. Dataset

The experiments are conducted using the CSPB.ML2018R2_NF dataset, which corresponds to the noise-free version of the CSPB.ML.2018 dataset [7]. This variant is used to isolate the intrinsic properties of the signal representations without the influence of additive noise. The dataset consists of synthetically generated baseband signals with controlled parameter variations. It includes a range of digital modulation schemes, such as BPSK, QPSK, 8-PSK, $\pi/4$-DQPSK, MSK, and several QAM formats. Each waveform contains 32,768 complex samples, and the dataset comprises 112,000 signal realizations. A 70/30 split is adopted for training and validation. The main generation parameters are summarized in Table I.

B. Data Generator and Pipeline

A custom data generator is used to dynamically produce training samples from the stored complex baseband signals of the dataset. For each sample, the raw IQ waveform is loaded and passed through a configurable preprocessing pipeline, illustrated in Fig. 1. The pipeline supports the following optional stages, which can be individually enabled or disabled:

1) carrier frequency offset (CFO) correction,
2) rational resampling to avoid fractional sps,
3) root raised cosine (RRC) matched filtering, and
4) symbol-rate sampling, where either the full waveform or a single symbol-spaced subsequence is retained.

In the latter case, the optimal sampling phase is selected based on a PAPR-based criterion, or a controlled offset is applied to simulate suboptimal sampling conditions (see Section VI-A). The processed IQ samples are then mapped onto a 2D histogram to form the constellation image, which is returned alongside its modulation label for training or evaluation. This modular design allows each processing stage to be systematically activated or deactivated, enabling controlled ablation of their individual contributions to classification performance.

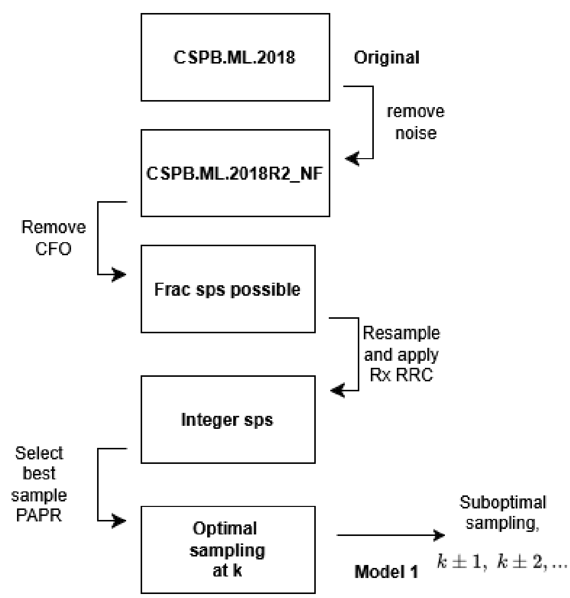

Fig. 1: Overall data-flow diagram of the preprocessing pipeline. This work focuses on Model 1.

TABLE I: CSPB.ML.2018 dataset parameter ranges

Parameters [8]	
Modulations	BPSK, QPSK, $\frac{\pi}{4}$-DQPSK, 8-PSK, MSK, 16-QAM, 64-QAM, 256-QAM
Base symbol period	1 to 15 samples
Carrier offset (normalized) [Hz]	$[-10^{-3}, 10^{-3}]$
Roll-off factor	[0.1, 1]
SNR[dB]	[-2, 12.8]
Up-Down sample factors	(1,1) and (3,2) to (10,9)
Noise spectral density [dB]	0
Signal length [samples]	32,768
Number of waveforms	112,000

IV. Constellations

Gray constellation diagrams correspond to two-dimensional histograms of the in-phase and quadrature components, where each bin stores the number of samples within its boundaries and the intensity is proportional to this count. The resulting histogram is represented in grayscale, with darker bins indicating higher sample densities, and is used as input to the neural network. Constellation diagrams are generated from IQ samples, with each waveform producing a single representation. In this work, two types of constellations are considered, both derived from an IQ histogram of the waveform samples. A density-based binning procedure maps the IQ plane into a discrete grid, yielding a two-dimensional intensity map where pixel values reflect local sample density, enabling the network to capture spatial patterns relevant for modulation classification. The histogram resolution is set to 128×128, resulting in images of the same size. This resolution was selected empirically based on preliminary experiments, although larger sizes have been explored in previous studies. RGB representations extend grayscale by encoding bin intensities using the Turbo colormap, mapping low densities to blue and high densities to red, with intermediate values represented by a continuous gradient. Each bin is thus represented by three components (Red, Green, Blue), forming a three-channel image suitable for convolutional neural networks. Examples of RGB constellation diagrams for BPSK, QPSK, 8-PSK, and 16-QAM signals at an SNR of approximately 8 dB are presented in Fig. 2.

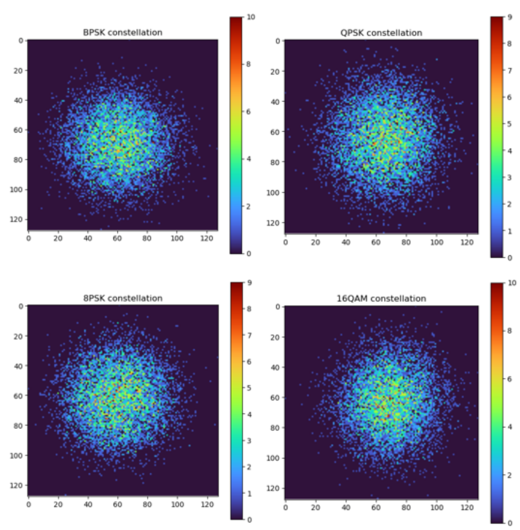

Fig. 2: RGB constellation visualization applied to the original CSPB.ML.2018 dataset.

V. Neural Network Models

Automatic modulation classification is formulated as a multiclass classification problem requiring robust feature extraction. Convolutional neural networks are used due to their effectiveness in learning spatial features from image-based IQ representations. All models share a common architecture consisting of convolutional feature extraction layers followed by fully connected classification layers. Variations are limited to the input representation and convolution type, enabling controlled comparison. Key architectural and training parameters are summarized in Table II. ReLU activations are used in hidden layers and softmax in the output layer. Weights are initialized using He normal and Glorot uniform initializers, and training is performed using the Adam optimizer with a learning rate of 0.001. Dropout is applied after convolutional layers to mitigate overfitting. Two CNN variants are considered. The baseline model processes single-channel grayscale constellation images obtained from direct IQ mapping, while the proposed model processes three-channel RGB tensors of size $3 \times 128 \times 128$ generated via histogram-based IQ mapping. Both architectures consist of two convolutional layers followed by flattening and a fully connected classifier. The grayscale CNN uses 40 filters of size 2×8 in the first convolutional layer and 10 filters of size 1×4 in the second, each followed by ReLU activation and dropout. The resulting feature maps are flattened and passed to a fully connected layer with 128 neurons and an 8-neuron softmax output layer. The RGB CNN mirrors this architecture, with differences confined to the input and convolution type. Separable convolutions are employed, decomposing standard convolution into depth- and pointwise operations to reduce the parameter count while maintaining expressive capacity. The extracted features are processed by the same fully connected and softmax layers, enabling direct comparison with the baseline.

TABLE II: Architecture parameters

Parameter	Value	Remarks
CNN kernels layer 1	40	activation="ReLU"
Dropout	50%	
CNN kernels layer 2	10	activation="ReLU"
Dropout	50%	
Dense nodes layer 1	128	activation="ReLU"
Aggregation layer nodes	8	activation='softmax'

VI. Model and Simulations

A. Symbol Sampling Model

After demodulation and matched filtering, the received baseband signal (including root raised cosine filtering) is represented as the discrete-time complex sequence

$$x_{\mathrm{rrc}}[n] = I[n] + jQ[n], \tag{1}$$

where $I[n]$ and $Q[n]$ denote the in-phase and quadrature components.

Let S denote the number of samples per symbol. In the ideal case, symbols are recovered by sampling at the optimal instants, yielding

$$x_s[k] = x_{\mathrm{rrc}}[kS], \tag{2}$$

where k is the symbol index.

In practice, the optimal sampling phase is unknown and must be estimated. To obtain a reference sampling phase, we employ a peak-to-average power ratio (PAPR) criterion. The signal $x_{\mathrm{rrc}}[n]$ is partitioned into S subsequences corresponding to offsets $\ell \in \{0, \ldots, S-1\}$:

$$x_\ell[m] = x_{\mathrm{rrc}}[\ell + mS].$$

where m indexes the samples within the subsequence associated with offset ℓ, with $m = 0, \ldots, M-1$, such that consecutive values of m are separated by one symbol period (i.e., S samples).

The average power associated with each offset is computed as

$$P_\ell = \frac{1}{M} \sum_{m=0}^{M-1} \left| x_{\mathrm{rrc}}[\ell + mS] \right|^2,$$

where M is the number of samples in each subsequence. To ensure comparability across offsets, the powers are normalized as

$$\tilde{P}_\ell = \frac{P_\ell}{\frac{1}{S}\sum_{i=0}^{S-1} P_i}.$$

The reference sampling phase is then defined as

$$k^\star = \arg\max_{\ell} \tilde{P}_\ell,$$

which corresponds to the offset best aligned with the symbol peaks.

Rather than relying solely on this estimated optimal phase, we introduce controlled sampling offsets to study the effect of timing mismatch. The sampled symbol sequence is therefore defined as

$$x_s[k] = x_{\mathrm{rrc}}[kS + k^\star + \Delta], \tag{3}$$

where $\Delta \in \{0, \pm 1, \pm 2, \ldots\}$ denotes an integer sampling offset relative to the estimated optimal phase.

When $\Delta = 0$, sampling aligns with the estimated optimal instant, and the constellation points are tightly concentrated around their ideal positions. For $\Delta \neq 0$, the sampling instant falls between two consecutive symbol periods, capturing the signal during its transition from one constellation point to the next. This introduces structured trajectory artifacts in the constellation diagram, as the sampled values reflect interpolated states between adjacent symbols rather than their ideal locations, producing the point splitting and geometric distortions visible in Figs. 3–4. The sampled sequence $x_s[k]$ is subsequently used to construct constellation diagrams or density-based histogram representations that serve as inputs to the neural network classifier.

B. Histogram Normalization

Constellation-like representations are obtained by mapping I/Q samples onto two-dimensional histograms. The choice of normalization is critical as it directly affects the preservation of signal structure and discriminative features, especially under low-SNR conditions. Two normalization strategies are considered. The first applies a logarithmic compression followed by max normalization. The $\log(1+x)$ or $log1p$ transform (natural logarithm) reduces the dynamic range of histogram counts while preserving relative differences between bins, and the result is scaled by the maximum value to constrain it to $[0, 1]$. This approach enhances low-intensity regions while avoiding domination by high-density bins. The second strategy is used for RGB representations. The histogram is first max-normalized to the range $[0, 1]$, without enforcing a probabilistic interpretation. A fixed colormap is then applied to generate a three-channel RGB image, enabling the network to exploit color-coded intensity variations. This transformation can be interpreted as a structured nonlinear embedding of density values into a higher-dimensional feature space, improving separability for CNN-based learning and enabling efficient channel-wise feature extraction [9].

Let $H \in \mathbb{R}^{N \times N}$ denote the histogram with bin values $H_{i,j}$, and let $\epsilon > 0$ ensure numerical stability. The logarithmic normalization is defined as

$$H_{\log} = \frac{\log(1 + H)}{\max_{i,j} \log(1 + H_{i,j}) + \epsilon}. \tag{4}$$

For the RGB representation, max normalization is given by

$$H_{\mathrm{rgb}} = \frac{H}{\max_{i,j} H_{i,j} + \epsilon}, \tag{5}$$

after which a colormap is applied to obtain the RGB image.

VII. Visual Signatures and Results

Figures 3–4 illustrate IQ constellation diagrams for different modulation schemes, obtained at the optimal sampling instant as well as with sampling delays of +1, -1, and +2 samples. In order of appearance, the samples for BPSK are distributed along the in-phase axis, while QPSK forms its characteristic diamond-shaped constellation. The 8-PSK modulation produces circular patterns, whereas DQPSK results in square-like structures. Finally, MSK shows convergence of its ideal sampling points, while QAM constellations exhibit a dispersion of points similar to additive white Gaussian noise (AWGN). Across all modulation schemes, deviations from the optimal sampling phase introduce structured distortions consistent with inter-symbol interference (ISI), including constellation spreading, point splitting, and trajectory deformation. These effects directly alter the density distribution of histogram representations, leading to the observed degradation in classification accuracy. Table III quantifies this effect for two training configurations: training using only the optimal sampling instant and training using the full waveform.

These experiments are conducted for both grayscale representations with log1p normalization and RGB based representations. In Table III, 'Best sample' refers to models trained using only the optimal sampling instant, while 'All samples' corresponds to training performed on the full waveform. The test cases $k^\star$, $k^\star + 1$, $k^\star + 2$, etc., denote evaluation at the optimal sampling phase and at increasing sampling offsets, respectively. The results show that classification accuracy degrades as the sampling offset increases. Furthermore, a mismatch between training and evaluation conditions leads to a significant performance drop. Validation is performed

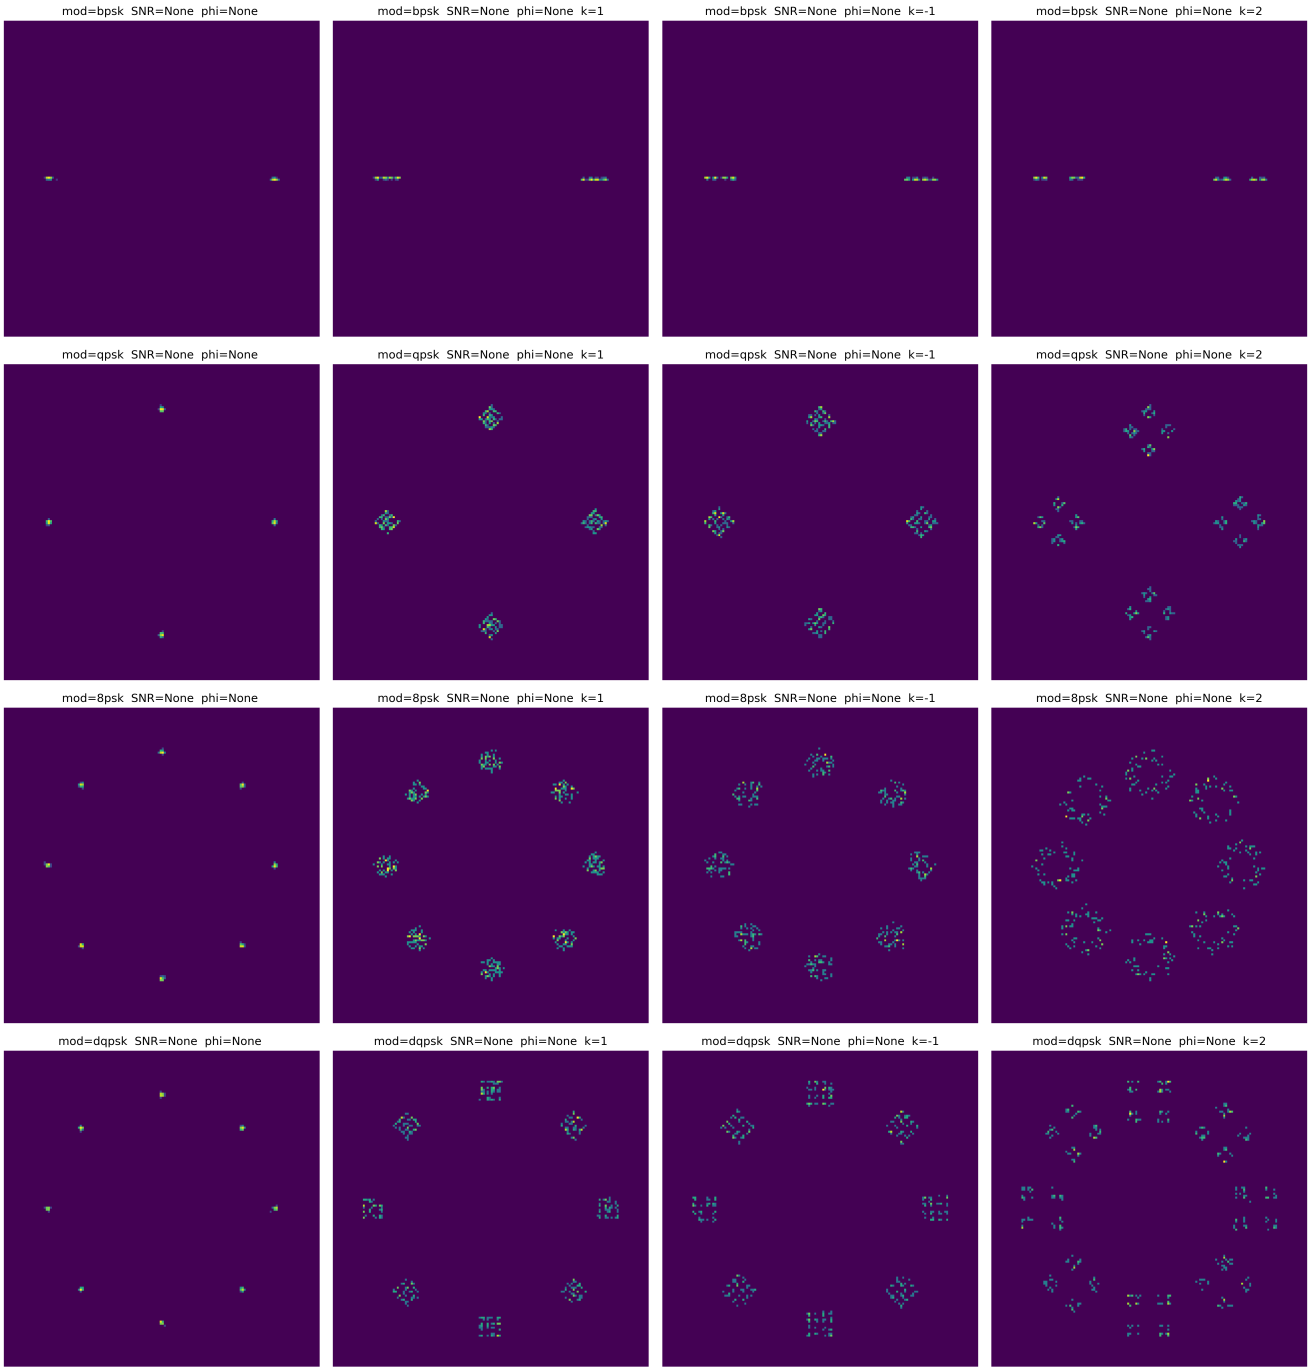

Fig. 3: Constellation diagrams of BPSK, QPSK, 8-PSK, DQPSK (rows) for four sampling offsets $k^\star = 0, 1, -1, 2$ (columns)

using both the optimal sampling instant and sampling delays ranging from -1 to 4 samples. For reference, when both training and validation are conducted on the full waveform, the classification accuracies reach 99.7% and 99.5% for log1p and RGB normalization, respectively. From these results, several conclusions can be drawn. First, in applications where communication parameters are known or can be accurately estimated, high classification performance can be achieved even with lightweight models, as opposed to the more complex architectures typically employed in the literature. However, classification accuracy degrades as the sampling offset increases or when the optimal sampling instant is incorrectly selected. The results further highlight the importance of consistency between training and inference conditions. Models trained on full waveforms do not generalize effectively when evaluated on single-sample representations, and conversely, training on optimally sampled data does not ensure robustness to sampling offsets. **Overall, this work demonstrates that**

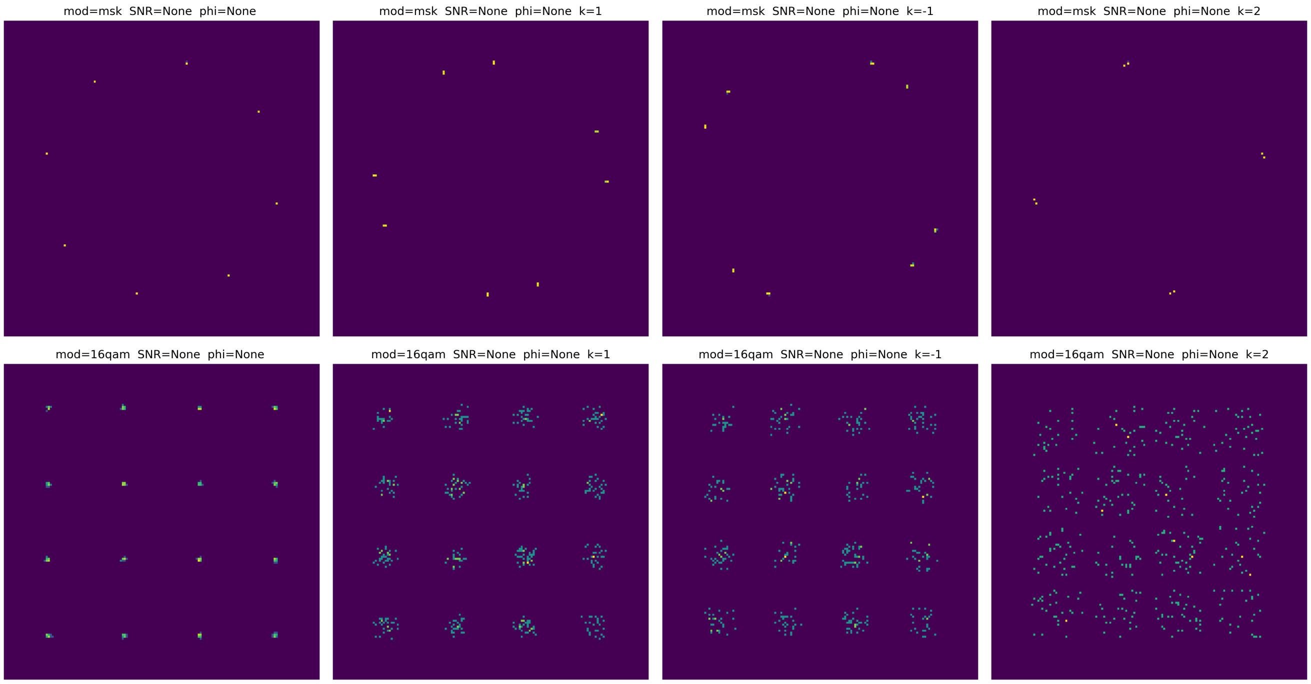

Fig. 4: Constellation diagrams of MSK, 16-QAM (rows) for four sampling offsets $k^\star = 0, 1, -1, 2$ (columns)

TABLE III: Validation accuracy (%) under different training strategies and sampling offsets. Here, $k^\star$ denotes the optimal sampling instant, and $k^\star + \Delta$ represents a sampling offset of Δ samples. Gray denotes grayscale histogram representations with log1p normalization.

Trained on	Best sample		All samples	
Test case	Gray + Log1p	RGB	Gray + Log1p	RGB
$\Delta = 0$	96.0	95.8	77.9	74.8
$\Delta = -1$	73.9	78.1	66.0	67.1
$\Delta = 1$	83.2	85.5	70.3	71.2
$\Delta = 2$	65.7	68.6	63.8	64.6
$\Delta = 3$	55.6	57.4	61.1	60.5
$\Delta = 4$	58.3	59	62.5	59.5

symbol sampling is not merely a preprocessing detail but a key factor influencing the reliability of learning-based modulation classifiers. From a practical perspective, the results advocate for the consistent use of all available samples during both training and inference, thereby leveraging oversampling not only as a classical digital signal processing technique but also as an effective strategy to improve robustness in AI-based receivers.

Acknowledgment

This work is supported by the Walloon Region research project CyberExcellence, n°2110186. The computations were performed using the resources of "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] Q. Zheng, X. Tian, L. Yu, A. Elhanashi, and S. Saponara, "Recent advances in automatic modulation classification technology: Methods, results, and prospects," *International Journal of Intelligent Systems*, vol. 2025, no. 1, p. 4067323, 2025. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1155/int/4067323

[2] A. S. Thakur and M. Imtiaz, "Deep learning for automatic modulation classification: A review," *Electronics*, vol. 15, no. 10, 2026. [Online]. Available: https://www.mdpi.com/2079-9292/15/10/2163

[3] V.-S. Doan, T. Huynh-The, C.-H. Hua, Q.-V. Pham, and D.-S. Kim, "Learning constellation map with deep cnn for accurate modulation recognition," in *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*, 2020, pp. 1–6.

[4] M. Leblebici, A. Çalhan, and M. Cicioğlu, "Deep learning-based modulation recognition with constellation diagram: A case study," *Physical Communication*, vol. 63, p. 102285, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S187449072400003X

[5] M. A. Abdel-Moneim, R. M. Al-Makhlasawy, N. Abdel-Salam Bauomy, E.-S. M. El-Rabaie, W. El-Shafai, A. E. A. Farghal, and F. E. Abd El-Samie, "An efficient modulation classification method using signal constellation diagrams with convolutional neural networks, gabor filtering, and thresholding," *Transactions on Emerging Telecommunications Technologies*, vol. 33, no. 5, p. e4459, 2022. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/ett.4459

[6] T. K. Oikonomou, N. G. Evgenidis, D. G. Nixarlidis, D. Tyrovolas, S. A. Tegos, P. D. Diamantoulakis, P. G. Sarigiannidis, and G. K. Karagiannidis, "Cnn-based automatic modulation classification under phase imperfections," *IEEE Wireless Communications Letters*, vol. 13, no. 5, pp. 1508–1512, 2024.

[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] C. M. Spooner, W. A. Brown, and G. K. Yeung, "Automatic radio-frequency environment analysis," in *Conference Record of the Thirty-Fourth Asilomar Conference on Signals, Systems and Computers (Cat. No.00CH37154)*, vol. 2, 2000, pp. 1181–1186 vol.2.

[9] A. Gros, V. Moeyaert, and P. Mégret, "Signal surface augmentation for artificial intelligence-based automatic modulation classification," *Electronics*, vol. 14, no. 23, 2025.