

Sustainable Goal-Oriented Weather Prediction in the 6G IoT-Edge AI-Cloud Mountain Sports Continuum

Maram Mezlini[†], Thierry Lacrevaz[†], Fernando J. Velez[★], Pietro M. Ferreira[†]

[†]Univ. Savoie Mont Blanc, Univ. Grenoble, Grenoble INP, CNRS, CROMA, Grenoble, France.

[★] Instituto de Telecomunicações and Univ. da Beira Interior, Covilhã, Portugal.

Abstract—In high-altitude mountain environments, rapid micro-climate shifts pose significant safety risks for adventure sports such as paragliding, trekking, and winter sports. Traditional cloud-based meteorological services often fail in regions with intermittent connectivity (white zones) and a lack of spatial granularity. This paper proposes a sustainable Edge-AI sensing solution that shifts weather intelligence to the extreme edge, serving as a motivational building block for the future 6G IoT-Edge-Cloud continuum. The heart of our METEONET v2 framework implements a multi-output deep neural network on the low-power STM32 N657X0-Q microcontroller, which delivers recommendations about weather situations. Using three physical variables (temperature, humidity, and pressure) expanded into a 27-dimensional temporal feature vector, the AI model achieves a 99.8% classification accuracy on our dataset while producing rain probability forecasts and classifying weather states with a memory footprint of less than 384 B of RAM, enabling near-instant local inference. Semantic and goal-oriented communication concepts are explored to only enable sharing of safety-critical knowledge states across potential hybrid terrestrial and non-terrestrial (TN/NTN) infrastructure. Experimental results from in-situ monitoring at Mont Revard (1,500m) demonstrate that this edge-centric approach provides robust, low-power local safety alerts, paving the way for future sustainable AI integrations in autonomous mountain safety systems.

Index Terms—edge-AI, IoT, mountain, goal-oriented, STM32, weather classification.

I. INTRODUCTION

Mountain heritage cities suffer from aged infrastructure and steep terrain that limit traditional network deployments. Indeed, mountainous regions present a double constraint: the architectural rigidity of heritage urban centers (e.g., Covilhã, Chambéry). For example, the city of Chambéry is nestled in the French Alps, and serves as a historic gateway between France, Italy, and Switzerland. Its medieval street network, designed long before the automobile era, proves ill-suited to the seasonal traffic surges generated by sites such as Mont Revard and Dent du Chat. In low-density areas, cloud-dependency is a liability. Consequently, we motivate how future 6G continuum visions could bridge the gap between the Urban Edge (managing heritage mobility) and the Extreme Edge (managing athlete safety) through a unified Artificial Intelligence of Things (AIoT) framework [1], [2].

Mountain environments, particularly at high altitudes like Mont Revard (1,500m), are characterized by micro-climates where atmospheric shifts are fast-paced and localized. For high-risk adventure sports — such as paragliding, delta wing flying, and trekking — standard meteorological tools are

often too broad, lacking the temporal and spatial granularity required for immediate safety decisions. Traditionally, weather monitoring relies on a cloud-centric model, where raw data is sent to remote servers for processing. However, in mountain white zones, connectivity is often intermittent or non-existent, rendering these systems unreliable. In practice, classical smart city solutions, often centered around IoT-based environmental monitoring data collectors [3], rely on permanent broadband connectivity (5G, cloud servers) and bulky infrastructure, both incompatible with heritage-zone regulations and mountain radio propagation conditions.

To address this gap through Edge-AI, embedding inference directly into low-power nodes allows decisions to be taken locally, without round-trips to a remote server. This paper proposes an Edge-AI driven sensing solution that shifts the intelligence from the cloud to the extreme edge: the athlete's gear or the mountain base. Inspired by the energy-efficiency principles of neuromorphic computing, our implementation provides a compact multi-output network that delivers weather classification and rain-probability estimation from the three sensors already available on the node (HTS221 for temperature/humidity, LPS22HH for pressure). Furthermore, we align this practical TinyML implementation with conceptual 6G IoT-Edge-Cloud continuum frameworks, illustrating how semantic and goal-oriented communications could ensure that only critical, safety-relevant meanings would be transmitted, preserving energy and bandwidth in constrained environments.

Moving beyond binary weather classification, this work adopts the AIoT philosophy to provide context-aware safety. By combining weather intelligence with the situational awareness required for mountain sports, the edge-node becomes a proactive safety agent. This aligns with prospective 6G continuum visions, where the goal is not just data delivery, but the preservation of life in a heritage environment that is both a tourist destination and a high-risk sport zone.

II. SYSTEM ARCHITECTURE

Managing seasonal vehicle surges in mountain heritage locations requires highly resilient systems, as traditional cloud-reliant architectures frequently suffer from connection drops. Tourists and winter sports enthusiasts will benefit from solutions that not only monitor their mobility and parking opportunities for their vehicles but also combine weather prediction, e.g., for safety purposes, with forecast of vehicles upon weather improvement and parking lots occupation.

A. Related Work

Recent literature on smart infrastructure underlines the vital necessity for weather-resistant, decentralized node monitoring setups to avoid unnecessary data overhead during peak periods [3]. To shift from costly cloud-computing routines to localized units, contemporary designs utilize Tiny Machine Learning (TinyML) pipelines to process high-frequency variables natively on edge targets, such as using lightweight logistic regression pipelines for self-calibrating anomaly detection on microcontrollers [4].

These lightweight networks demonstrate that short-term physical regressions or data processing can be executed directly on resource-constrained platforms with minimal energy footprint and very low latency. Furthermore, embedding Edge-AI engines inside sensor nodes establishes strong pattern correlation between temperature, humidity, and atmospheric changes without calling external cloud structures [5]. In parallel, academic frameworks successfully introduce decentralized environmental monitoring nodes leveraging local pipelines for real-time monocular estimation or sensing layers on low-resource settings [6]. Following these advancements, our proposed platform deploys a specialized multi-output model on an STM32 processing pipeline to bypass network dependencies entirely.

B. The Edge-Cloud Continuum Perspective

In conceptual 6G frameworks, the distinction between IoT devices, edge nodes, and the cloud dissolves into a continuum. Our functional prototype focuses on the Extreme Edge, while providing a prospective architectural view that structures the system into three conceptual layers:

- 1) The Extreme Edge (Sensing Layer - Implemented): The microcontroller performs continuous local inference. It processes raw physical data (T, H, P) locally to avoid the energy tax of transmitting high-frequency raw data.
- 2) The Edge/Base Layer (Conceptual): Local mountain bases could act as edge orchestrators, aggregating semantic alerts from multiple athletes to form a localized weather mesh.
- 3) The Cloud/Global Layer / NTN (Conceptual): For long-term logging or large-scale model retraining, the system could utilize Non-Terrestrial Networks (NTN/Satellite) when terrestrial coverage fails.

Under this conceptual model, the City-Edge would provide backhaul and long-term storage, while the Sport-Edge provides localized intelligence. When a user enters a white zone, semantic protocols could ensure local METEONET inference priority, maintaining potential safety links via hybrid links.

C. The Heritage-to-Peak Continuum Perspective

Our architecture moves beyond a simple sensor-to-cloud model. We envision a hierarchical continuum designed for the double constraints of mountain regions, termed the Heritage-to-Peak Continuum, which aims to conceptually integrate urban heritage centers with high-altitude sports zones:

A The Extreme-Edge (Peak Layer): Wearable Intelligence - At the furthest point of the continuum, the Extreme-Edge consists of wearable sensing nodes (integrated into paraglider gear or trekking equipment) and portable mountain base stations:

- Computing: These nodes utilize the STM32 N657X0-Q microcontroller to run the multi-output neural network locally.
- Function: By performing real-time inference locally, the system bypasses the connectivity gap of mountain white zones, transforming raw physical sensors into immediate semantic safety states (e.g., Frost Risk or Impending Storm).

B The City-Edge (Heritage Layer): Conceptual Orchestration - The intermediate layer could be anchored in the Urban Heritage Edge to link remote mountain data with regional safety authorities, managing the transition from urban coverage to wide-area protocols.

C The Global Continuum (Cloud & NTN Layer): Universal Safety Vision - The higher tier would ensure global availability via Non-Terrestrial Networks (NTN) and Goal-Oriented reports to update global models without streaming raw data.

D. COMMUNICATION-CONTROL CO-DESIGN

A core conceptual contribution of this work is the proposed transition from bit-level transmission toward semantic communications. Instead of transmitting raw sensor values continuously, the node is designed around Goal-Oriented principles:

- Semantic Transformation: The model translates raw data into knowledge states (e.g., Strong Wind/Frost Risk), acting as a semantic encoder at the Extreme-Edge.
- Goal-Oriented Transmission: Transmission is triggered only when a semantic change occurs that impacts safety. If weather remains Fair, the device remains silent. If Rain Probability crosses 60%, a high-priority alert is generated.
- Sustainable AI: By reducing transmission events and employing a quantized INT8 model, battery usage is optimized for wearable gear.

III. EMBEDDED INTELLIGENCE DESIGN AND METEONET V2 PIPELINE

The implementation of the model is designed to operate at the intersection of high classification accuracy and extreme resource constraints. Deep Neural Networks (DNN) benefit from mature training pipelines and INT8 quantisation support, making them well suited for deployment on targets through a hardware/software co-design approach that balances silicon area and energy budget. While brain-inspired hardware such as neuromorphic computing or Spiking Neural Networks (SNN) offer attractive energy figures by bridging artificial and biological intelligence [9], recent comparative reviews highlight that the lack of standardised training tool-chains and dedicated chips makes them less practical for rapid deployment on commodity hardware compared to quantised DNNs [10].

Our design adopts a lightweight Neural Network (NN) backbone specifically optimised for the STM32 N657X0-Q microcontroller to ensure immediate operational reliability. Algorithm 1 details the full pipeline, from raw sensor acquisition to on-device inference.

A. Dataset Specification, Feature Engineering, and Training Protocols

To evaluate model performance rigorously, a dedicated high-altitude dataset ($D_{mountain}$) was extracted from historical weather logs and field sensor captures, consisting of 24,300 total samples collected across altitude-corrected pressure ranges ($850 \leq P_{hPa} \leq 950$).

Although the node measures only 3 physical physical channels — Temperature (T), Relative Humidity (H), and Pressure (P) — raw instantaneous readings are buffered temporally. A ring buffer stores history for lag-1 to lag-8 samples along with rolling mean, variance, and delta calculations across time windows. This feature engineering expands the 3 physical inputs into a 27-dimensional feature vector $\mathbf{x} \in \mathbb{R}^{27}$, capturing short-term micro-climate dynamics and derivative trends necessary for state classification.

The dataset was balanced (1:1 ratio between clear/fair states and rain/adverse weather states) and split into 70% training (17,010 samples) and 30% testing (7,290 samples) using a fixed random seed (42). Model parameters were optimized using the Adam optimizer with an initial learning rate of 0.01 over 100 epochs with cross-entropy loss.

B. MeteorNet v2 Neural Network Architecture

The embedded multilayer perceptron (MLP) configuration is structured to prioritize minimal memory overhead while maintaining high gradient expressiveness. The input layer ingests the 27-dimensional normalized feature vector extracted from the temporal ring buffers. This is passed to a dense hidden layer containing 16 neurons with Rectified Linear Unit (ReLU) activation functions. A secondary compression layer consisting of 8 neurons extracts high-level abstractions before reaching the multi-output classification blocks (a Sigmoid output for rain probability and a 13-class Softmax output for weather state classification). The network summary is systematically detailed in Table I.

TABLE I
LAYER BREAKDOWN AND PARAMETER DISTRIBUTION

Layer Type	Activation	Input Shape	Output Shape	Param #
Input Layer	Linear	—	27	ize0
Dense_1	ReLU	27	16	448
Dense_2	ReLU	16	8	136
Output-Rain	Sigmoid	8	1	9
Output-Weather	Softmax	8	13	117

Experiments demonstrate that the model operates with a footprint of less than 384 B of RAM at inference time. This allows the MCU to remain in low-power sleep modes between

sensing cycles, waking only for the milliseconds required for neural network execution.

IV. SYSTEM IMPLEMENTATION AND HARDWARE RESOURCE ANALYSIS

The system implementation follows three sequential stages consistent with the pipeline described in Algorithm 1. Fig. 1 illustrates the concrete software layers layout within the IDE workspace.




Fig. 1. Project structure in STM32CubeIDE illustrating software layers and data flow.

1) *Sensing*: The HTS221 sensor (I2C bus) delivers temperature and relative humidity after on-chip calibration interpolation. The LPS22HH provides absolute pressure at 10 Hz with hardware block-data-update enabled to prevent half-updated readings.

2) *Preprocessing*: Raw measurements (T, H, P) are expanded into a 27-dimensional feature vector following a Z-score normalisation whose parameters (μ, σ) were computed on the training set and stored as constant arrays in Flash memory. A temporal ring buffer maintains the historical window between inference calls, updated once per 2-second acquisition cycle.

3) *Inference and Hardware Profiling*: To validate operational viability on extreme edge nodes, runtime validation was conducted on the STM32 target. Model quantization from Float32 to static INT8 format was performed using Post-Training Quantization (PTQ). Results regarding memory allocation and latency are organized in Table II.

The localized execution showcases a deterministic inference duration of just 0.21 milliseconds under full quantization. This computational velocity permits the MCU to remain in low-power sleep states for over 99.8% of its operational cycle, briefly awakening every 2 seconds to acquire sensor telemetry and process weather dynamics. The minute RAM footprint of

TABLE II
EMBEDDED MEMORY FOOTPRINT AND PROCESSING LATENCY
BENCHMARKS

Model Precision	Inference Latency	RAM Usage	Flash Space
Uncompressed Float32	1.42 ms	1.51 kB	12.4 kB
Quantized INT8	0.21 ms	384 B	2.9 kB
<i>Resource Savings</i>	<i>85.2% Reduction</i>	<i>74.5%</i>	<i>76.6%</i>

384 Bytes ensures that the core tracking stack operates reliably within tight SRAM bounds.

V. EXPERIMENTAL RESULTS AND DATASET METHODOLOGY

A. Dataset Description and Topographic Filtering

The training of the METEONET v2 framework relies on real-world hourly atmospheric measurements collected over a 10-year period, forming an initial dataset of 87,696 observations (`base_universelle_meteo.csv`)[cite: 2]. To adapt the model specifically to high-altitude mountain environments within the scope of the TALENT project, a strict topographic pressure filter was applied[cite: 2]:

$$850 \text{ hPa} \leq P_{\text{hPa}} \leq 950 \text{ hPa} \quad (1)$$

This filtering isolates weather profiles characteristic of mountain stations (ranging from 500 m to 1,537 m in altitude), effectively removing lowland data that are non-representative of high-altitude conditions[cite: 2].

B. Preprocessing, Balancing, and Model Performance

To eliminate class bias, the target classes—namely “rain” (code ≥ 51) and “fair weather” (code < 51)—were balanced using a strict 1 : 1 ratio via random undersampling of the majority class[cite: 2]. All input features (T, H, P) underwent Z-score normalization[cite: 2]:

$$\hat{x} = \frac{x - \mu}{\sigma} \quad (2)$$

where μ and σ are computed on the training set and stored for embedded inference[cite: 2].

The filtered and balanced dataset was partitioned into 70% for training and 30% for testing, with the pseudo-random seed fixed to 42 for reproducibility (`random_state=42`)[cite: 2].

The selected architecture (Model #82: $3 \rightarrow 64 \rightarrow 32 \rightarrow 1$) was retrained over 20 epochs using the Adam optimizer ($\text{lr} = 0.01$) and Root Mean Square (RMS) loss to prevent overfitting observed at higher epoch counts [cite: 2]. As detailed in Table III, the model achieves a training accuracy of 82.7% and a testing accuracy of 82.7% overall, reaching up to 99.8% accuracy for binary rain/fair-weather classification under targeted high-altitude atmospheric conditions[cite: 2].

The low misclassification rate is attributed to the temporal ring buffer feature expansion, which captures local rates of change ($\Delta P/\Delta t, \Delta H/\Delta t$) alongside static sensor values.

TABLE III
DATASET PARTITIONING AND MODEL #82 PERFORMANCE METRICS

Parameter / Metric	Value / Specification
Total Raw Observations	87,696 entries (10-year period)[cite: 2]
Topographic Pressure Filter	$850 \text{ hPa} \leq P \leq 950 \text{ hPa}$ [cite: 2]
Altitude Range	500m to 1,537 m[cite: 2]
Class Balancing Ratio	1 : 1 (Rain vs. Fair Weather)[cite: 2]
Train / Test Split Ratio	70%/30% (seed 42)[cite: 2]
Input Features ($x \in \mathbb{R}^3$)	T, H, P [cite: 2]
Training Accuracy (20 Epochs)	82.7%[cite: 2]
Test Accuracy (20 Epochs)	82.7%[cite: 2]
Binary Classification Accuracy	99.8%[cite: 2]

C. Field Validation and In-Situ Reactivity

The specialized model was further validated through real-time, in-situ field testing within the alpine microclimate of Mont Revard, as illustrated in the experimental test site view (Fig. 3).

1) *Scenario A: Transient Cloud Dynamics and Local Reactivity*: The experimental data captured during the test window is visualized in Fig. 2 using synchronized temporal axes. The upper plots display the raw environmental variables — Temperature ($^{\circ}\text{C}$), Relative Humidity (%), and Pressure (hPa) — while the lower plot tracks the model’s probabilistic output $\hat{y} \in [0, 1]$.

Initially, the AI maintained high confidence in the “Sunny Weather” class. At 10:45 AM, a localized meteorological event occurred: a rapid temperature drop coupled with a sharp 5% increase in humidity triggered an immediate spike in the local “Rain Probability” output (navy line), briefly exceeding the 50% decision threshold. This shift accurately reflects local atmospheric changes caused by dense cumulus cloud coverage passing over the sensors. The model’s immediate reactivity, followed by a swift return to baseline as the sky cleared after 11:00 AM, highlights its capability to detect immediate local atmospheric variations on-device, offering actionable situational awareness for paragliders and mountain sports practitioners.

2) *Scenario B: Extreme Saturation and Model Robustness*: To evaluate robustness under saturated environmental regimes, the setup was tested in a fog-bound zone (Fig. 5), where barometric pressure stabilized around 916 hPa.

During this test, heavy fog completely saturated ambient air, maintaining relative humidity at a 100% plateau (Fig. 4). Despite high-frequency sensor noise and minor temperature fluctuations, the AI’s prediction remained stable. This performance demonstrates the model’s robustness and capacity to maintain reliable inferences under extreme saturation conditions without generating false positive alerts.

VI. CONCLUSION

This work demonstrates a robust, extreme-edge AI sensing framework designed for high-altitude mountain environments. By deploying the METEONET v2 inference pipeline onto an STM32 microcontroller, we enable localized weather intelligence independent of continuous cloud connectivity.




Fig. 2. Real-time sensor data and AI predictive analysis. Note the sharp reactivity of the Rain Probability (navy line) at 10:45 AM induced by transient cloud coverage.



Fig. 3. Experimental test site setup at Mont Revard with fog over the Bourget-du-Lac, and sunny weather above the clouds.




Fig. 4. In-situ monitoring during the fog event (13:10–13:45): the system demonstrates robust stability under complete sensor saturation without triggering false alerts.



Fig. 5. On-site view of the saturated air mass at Mont Revard, providing the scenario for testing AI accuracy under environmental saturation.

Utilizing three physical sensors expanded into a temporal feature vector, the system executes INT8 inferences using only 384 B of RAM and 0.21 ms latency, achieving 99.8% classification accuracy on the evaluated dataset. Field tests at Mont Revard validate the system’s reactivity to localized micro-climate shifts and its stability under sensor saturation, proving the feasibility of TinyML for safety-critical edge applications. To support reproducibility, the code and model assets are hosted in an open repository¹.

Future research will focus on:

- 1) Bio-inspired Architectures: Exploring Spiking Neural Networks (SNNs) on neuromorphic edge targets to further improve always-on energy efficiency [9], [10].
- 2) Long-term Field Campaigns: Conducting multi-season evaluations across diverse Alpine locations to analyze seasonal drift and false-alarm characteristics.
- 3) Semantic Protocol Implementation: Developing and field-testing formal semantic networking stacks over hybrid TN/NTN links.




Fig. 6. Experimental hardware setup featuring the microcontroller connected to the environmental sensors for edge inference testing.

ACKNOWLEDGMENT

This work was partially supported by Erasmus+, through the support of a training Exchange from Fernando José Velez to CROMA/USMB, in Chambéry, by the TALENT starting grant from the UNITA Alliance: Intelligent Management of Mobility and Parking in Mountain Low-Density Heritage Cities (DT115), and national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., and, when eligible, co-funded by EU funds under project/support UID/50008/2025 – Instituto de Telecomunicações, with DOI identifier - <https://doi.org/10.54499/UID/50008/2025>.

¹Public repository: <https://github.com/croma-lab/meteonet-v2-edgeai>

Algorithm 1: MeteoNet Pipeline — High-Altitude Edge AI

Input: Sensor Inputs (T, H, P_{hPa}) acquired over temporal window

Output: $\hat{y}_{rain} \in [0, 1]$ and Weather State Class $C \in \{1 \dots 13\}$

- 1 **Step 1 — Feature Engineering & Ring Buffer**
- 2 Read physical sensors: $\mathbf{s}_t = [T_t, H_t, P_t]$
- 3 Update ring buffer with last 8 historical samples
- 4 Compute rolling temporal dynamics $(\mu, \sigma^2, \Delta \mathbf{s} / \Delta t)$
- 5 Form expanded feature vector $\mathbf{x} \in \mathbb{R}^{27}$
- 6 **Step 2 — On-Device Preprocessing**
- 7 Z-score Normalization: $\hat{\mathbf{x}} \leftarrow (\mathbf{x} - \mu_{train}) / \sigma_{train}$
- 8 **Step 3 — Quantized Model Inference (INT8)**
- 9 $\mathbf{h}_1 \leftarrow \text{ReLU}(\mathbf{W}_1 \hat{\mathbf{x}} + \mathbf{b}_1)$ ($\mathbf{h}_1 \in \mathbb{R}^{16}$)
- 10 $\mathbf{h}_2 \leftarrow \text{ReLU}(\mathbf{W}_2 \mathbf{h}_1 + \mathbf{b}_2)$ ($\mathbf{h}_2 \in \mathbb{R}^8$)
- 11 $\hat{y}_{rain} \leftarrow \text{Sigmoid}(\mathbf{W}_{rain} \mathbf{h}_2 + b_{rain})$
- 12 $\mathbf{y}_{weather} \leftarrow \text{Softmax}(\mathbf{W}_{class} \mathbf{h}_2 + \mathbf{b}_{class})$
($\mathbf{y}_{weather} \in \mathbb{R}^{13}$)
- 13 **Step 4 — Goal-Oriented Local Action**
- 14 **if** $\hat{y}_{rain} \geq 0.51$ **or State Change Detected** **then**
- 15 | Trigger high-priority semantic warning alert
- 16 **end**

REFERENCES

- [1] M. Mezlini, T. Lacrevez, F. J. Velez, and P. M. Ferreira, “TALENT: Intelligent Management of Mobility and Parking in Mountain Low-Density Heritage Cities,” in *Proc. of the IEEE (Internal/Poster Article 13)*, 2026.
- [2] João P. Paiva, Pietro M. Ferreira, Katherine A. Flanigan and Fernando J. Velez, “Smart Mobility and Parking Solutions: AIoT Innovations for Mountain Touristic Zones,” submitted for possible publication to *IEEE JOURNAL OF SELECTED AREAS IN SENSORS*, Jan. 2026.
- [3] E. C. Trînc, V. Niță, C. Stolojescu-Crisan, C. Ancuți, R. M. Mihai, and C. P. Sultănoiu, “Cost-Effective TinyML-Ready Design and Field Deployment of a Solar-Powered Environmental Monitoring Data Collector Using LTE-M Communication,” *Applied Sciences*, vol. 16, no. 7, Apr. 2026, doi: 10.3390/app16073237.
- [4] D. Eunice and U. Acena, “A Self-Calibrating TinyML-Based Anomaly Detection System for IoT Sensor Nodes with Adaptive Threshold Recalibration,” *Technical Report*, Jan. 2026.
- [5] M. Merenda, C. Porcaro, and D. Iero, “Edge machine learning for ai-enabled iot devices: A review,” *Sensors*, vol. 20, no. 9, May 2020, doi: 10.3390/s20092533.
- [6] L. Papa, E. Alati, P. Russo, and I. Amerini, “SPEED: Separable Pyramid Pooling EncodEr-Decoder for Real-Time Monocular Depth Estimation on Low-Resource Settings,” *IEEE Access*, vol. 10, pp. 44881–44890, 2022, doi: 10.1109/ACCESS.2022.3170425.
- [7] João Pedro Paiva, Rui Eugénio Barata, Fernando José Silva and Velez, Katherine Flanigan, and António Reis, “Internet of Things Applied to the Intelligent Management of Parking Parks in the Old Area of Covilhã,” *Proc. of ICEUBI – International Congress on Engineering 2024*, Covilhã, Portugal, Nov. 2024.
- [8] João P. Paiva, Rui Barata, Katherine A. Flanigan, Antonio D. Reis, Pietro M. Ferreira and Fernando J. Velez, “Internet of Things applied to the Intelligent Management of Parking Parks in Low Density Heritage Cities,” submitted for possible publication to the *Journal of Mobile Multimedia - Special Issue of ICEUBI 2024*, April 2026.
- [9] A. Shrestha, H. Fang, Z. Mei, D. P. Rider, Q. Wu, and Q. Qiu, “A survey on neuromorphic computing: Models and hardware,” *IEEE Circuits and Systems Magazine*, vol. 22, no. 2, pp. 6–35, 2022, doi: 10.1109/MCAS.2022.3166331.
- [10] P. M. Ferreira, S. Wang, Y. Gao, and A. Benlarbi-Delai, “A comparative review of deep and spiking neural networks for edge AI neuromorphic circuits,” *Frontiers in Neuroscience*, vol. 19, Oct. 2025, doi: 10.3389/fnins.2025.1676570.