Edge AI for Climate-Aware ET_0 Forecasting and Autonomous Precision Irrigation:

A Hardware-Software Co-Design for Autonomous Precision Irrigation

Horatio Harley Koffivi HODOUTO
UR-GAMMA3
Université de Technologie de Troyes
Troyes, France

Abel Cherouat
UR-GAMMA3
Université de Technologie de Troyes
Troyes, France

Milohum Mikesokpo Dzagli
LPMCS
Université de Lomé
Lomé, TOGO

David BASSIR
Smart Structural Health Monitoring and Control Laboratory
Dongguan University of Technology
Dongguan, CHINA

* Corresponding author: hodoutoh@gmail.com

Abstract— **Precision agriculture is hindered by its dependence on centralized cloud infrastructures, which prevents deployment of advanced deep learning (DL) in disconnected, resource-constrained rural environments. We propose a globally transferable Edge AI framework capable of executing reference evapotranspiration (ET_0) forecasting directly on low-power microcomputers. To overcome hardware limitations, we introduce a Distilled Noise-Invariance training paradigm: offline Empirical Mode Decomposition and Wavelet Denoising (EMD-WD) are applied only to target variables during training, enabling the network to learn an intrinsic auto-correction manifold without runtime filtering. We evaluate INT8-quantized Mamba-2 and Temporal Fusion Transformer (TFT) models enhanced with Köppen–Geiger climatological latent embeddings across 309 globally distributed sites. The optimized architectures preserve full predictive accuracy while achieving up to 72–80% reductions in memory and 68–75% reductions in latency, all under 5 W of power, enabling cloud-independent, noise-resilient precision irrigation.**

Keywords: Edge AI; TinyML; Quantized Neural Networks; Noise-Invariant Learning; Smart Irrigation

I. INTRODUCTION

The intensification of agriculture under worsening global water deficits has elevated irrigation optimization to a fundamental requirement for food security [1]. Central to this optimization is the accurate estimation of reference evapotranspiration (ET_0), which defines the thermodynamic baseline of atmospheric water demand [2]. However, the standard FAO-56 Penman–Monteith equation relies on high-fidelity meteorological data, and in real-world agricultural environments—open-field or greenhouse settings—these measurements are often corrupted by IoT sensor drift, device noise, and micro-scale turbulence, severely degrading the operational reliability of traditional physics-based models[3], [4].

Deep learning (DL) offers powerful non-linear representations that can mitigate these metrological uncertainties, and several recent works have demonstrated that sequential DL architectures can accurately forecast ET_0 from limited climate variables [5]. Yet their in-situ deployment in rural and off-grid zones remains constrained by several barriers. On one hand, spatial hyper-specialization limits model generalization beyond the specific basins or sites used for training[3]. On the other hand, the parameter density of state-of-the-art architectures, such as Transformers, entails heavy reliance on cloud infrastructures and power-hungry GPU servers, rendering continuous connectivity and high energy consumption a systemic bottleneck [6], [7]. In isolated agricultural regions, this dependency introduces critical failure points related to latency, bandwidth limitations, and data-privacy concerns.

Edge computing offers a promising alternative by relocating intelligence to low-power microcomputers (e.g., Raspberry Pi), enabling local decision-making without cloud connectivity [8]. However, it demands a tight reconciliation between advanced hydro-climatic modeling and stringent hardware constraints.

This study proposes a hardware-software co-design framework for fully off-grid Edge AI deployment, enabling precise ET_0 forecasting directly on resource-constrained devices. The core algorithmic contribution is Distilled Noise-Invariance: rather than executing real-time filtering on the Edge, Empirical Mode Decomposition combined with Wavelet

Denoising (EMD-WD) is applied exclusively to the target variables during offline training [9]. The neural network thus learns an intrinsic noise-cancellation manifold within its synaptic weights. At deployment, the model ingests raw sensor data and implicitly corrects measurement artifacts without any explicit preprocessing, effectively eliminating the need for battery-draining on-device signal processing.

On the hardware side, the framework builds upon the Mamba-2 selective state-space architecture, whose linear temporal complexity $\mathcal{O}(N)$ alleviates memory bottlenecks, combined with post-training asymmetric INT8 quantization for efficient inference [6]. To ensure cross-site transferability, climatological embeddings derived from the Köppen–Geiger classification are integrated into the model, providing a global inductive bias that reduces the need for local recalibration [10]. By unifying rigorous sequential modeling with frugal embedded engineering, this work advances precision agriculture and supports sustainable water-use practices[11].

Our main contributions can be summarized as follows:

1. A neural auto-correction paradigm (Distilled Noise-Invariance) that obviates the need for explicit signal-processing algorithms on constrained Edge devices.
2. The integration of macro-climatic inductive bias enabling global zero-shot or few-shot deployment, thereby minimizing the requirement for site-specific retraining.
3. Comprehensive hardware profiling demonstrates fast, robust hydro-climatic inference under a frugal energy envelope (below 5 W).

II. METHODS

A. Meteorological data

A globally scalable Edge AI model must avoid the need for site-specific retraining. To ensure spatial transferability, a large and heterogeneous dataset was compiled using a strict multi-scale stratification methodology based on high-resolution Köppen–Geiger (KG) climate classification maps, derived from 1991–2020 climatological normals. A total of 309 discrete geographical sites were selected to represent the full statistical diversity of the Earth's hydroclimate, covering all 30 Köppen–Geiger sub-classes. Continuous daily meteorological time series for these sites were extracted from the NASA POWER (Prediction of Worldwide Energy Resources) satellite database, spanning from 1 January 1984 to 30 September 2024. The multivariate input tensor comprises Air temperature at 2 m height (°C), Relative humidity (%), Wind speed at 2 m (m.s^{-1}), Downward shortwave solar irradiance (W.m^{-2}).

B. Target computation of reference evapotranspiration

To provide a rigorous thermodynamic ground truth for the supervised learning phase, daily ET_0 values were computed using the FAO-56 Penman-Monteith formulation. This approach combines the energy balance and aerodynamic mass transfer methods:

$$ET_0 = \frac{0.408\Delta(R_n - G) + \gamma \cdot \frac{900}{T+273} u_2 \cdot (e_s - e_a)}{\Delta + \gamma(1 + 0.34u_2)} \quad (1)$$

where ET_0 is the reference evapotranspiration ($\mathrm{mm} \cdot \mathrm{day}^{-1}$), Δ is the slope of the vapor pressure curve (kPa °C^{-1}), R_n is the net radiation at the crop surface, G is the soil heat flux density, γ is the psychrometric constant, T is the mean daily air temperature, u_2 is the wind speed, e_s is the saturation vapor pressure, and e_a is the actual vapor pressure. This derived vector $\boldsymbol{Y_{raw}}$ serves as the baseline target for the training phase.

C. Distilled Noise-Invariance via Offline EMD-WD Teacher-Forcing

A critical bottleneck in agricultural Edge Computing is the processing of heteroscedastic noise generated by low-cost field sensors. Conventional systems either suffer catastrophic accuracy degradation or burn vital CPU cycles running real-time signal processing algorithms like Fast Fourier Transforms (FFT).

We introduce a paradigm-shifting approach: Distilled Noise-Invariance. Based on the Information Bottleneck principle, we force the neural network to learn a latent manifold that is invariant to input noise. This is achieved by moving the computational burden entirely to the offline training phase. First, the raw target signal $\boldsymbol{Y_{raw}}$ is decomposed using Empirical Mode Decomposition (EMD) into n Intrinsic Mode Functions (IMFs) and a monotonic residual $r_n(t)$:

$$y(t) = \sum_{i=1}^{n} \mathrm{IMF}_i\,(t) + r_n(t) \quad (2)$$

Because EMD is sensitive to mode mixing, we apply Wavelet Denoising (WD) exclusively to the highest-frequency components ($\mathrm{IMF}_1 and$ IMF_2). We employ a Symlet 8 mother wavelet with a soft-thresholding function defined as:

$$\eta_s(w, \lambda) = \mathrm{sgn}(w) \max(|w| - \lambda, 0) \quad (3)$$

where the threshold λ is derived via the universal Threshold principle ($\lambda = \sigma\sqrt{2\ln N}$). The resultant signal $\boldsymbol{Y_{clean}}$ is a mathematically, thermodynamically coherent target. During backpropagation on the high-performance workstation, the neural network is trained to map the raw, noisy input tensor $\mathbf{X}_{\mathrm{noisy}}$ directly to $\mathbf{Y}_{\mathrm{clean}}$. By minimizing the robust Log-Cosh loss, the synaptic weights implicitly parameterize a non-linear noise filter. Once deployed on the Edge device, the model performs a standard forward pass on raw sensor data, auto-correcting parametric errors without any explicit on-device signal-processing code.

D. Hardware-Aware Sequence Modeling: Mamba-2 and TFT

Selecting an appropriate neural architecture for Edge deployment requires a careful trade-off between temporal receptive field and computational complexity (FLOPs). We benchmark two state-of-the-art architectures:

1. **Temporal Fusion Transformer (TFT):**

TFT employs specialized Variable Selection Networks (VSN) to filter irrelevant inputs and multi-head self-attention to capture long-range temporal dependencies. The attention mechanism computes the alignment score as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (4)$$

While TFT offers strong interpretability, the quadratic complexity $\mathcal{O}(L^2)$ of dot-product attention with respect to sequence length (L) causes severe memory-bandwidth bottlenecks and thermal throttling on ARM-based microprocessors.

2. **Mamba-2 (Selective State-Space Model):**

To circumvent the quadratic bottleneck, we implement Mamba-2, a continuous-time state-space model discretized for deep sequence learning. Mamba-2 maps a 1D sequence $x(t) \in \mathbb{R}$ through a hidden state $h(t) \in \mathbb{R}^N$ via the continuous dynamics:

$h'(t)=\mathbf{A}h(t)+\mathbf{B}x(t), y(t)=\mathbf{C}h(t).$

Unlike standard SSMs, Mamba-2 introduces selective, input-dependent parameters $(\mathbf{B}_t, \mathbf{C}_t, \Delta_t)$, allowing the model to selectively propagate relevant hydro-climatic memory and discard transient anomalies. The continuous matrices are discretized using a zero-order hold (ZOH) rule:

$$\bar{\mathbf{A}}_t = \exp(\Delta_t \mathbf{A}), \bar{\mathbf{B}}_t = (\Delta_t \mathbf{A})^{-1}(\exp(\Delta_t \mathbf{A}) - I)\Delta_t \mathbf{B}_t \quad (5)$$

The discretized forward pass becomes $h_t = \bar{\mathbf{A}}_t h_{t-1} + \bar{\mathbf{B}}_t x_t$. Crucially, Mamba-2 achieves strict $\mathcal{O}(N)$ linear-time scaling, making it exceptionally well suited for deployment on limited-RAM Edge architectures.

To ensure global transferability, the specific Köppen–Geiger climate class of the deployment site is injected as a static categorical embedding matrix $\mathbf{E} \in \mathbb{R}^{30 \times d_{\text{model}}}$. This embedding acts as a global context vector, shifting the latent distribution of the state-space matrices to adapt to local climatic regimes without requiring in-situ retraining.

E. Edge Optimization via Asymmetric Affine Quantization

Neural networks are conventionally trained using 32-bit floating-point (FP32) arithmetic, which is incompatible with the power constraints and ALU capabilities of Edge microcomputers. To deploy Mamba-2 on the Edge, we design a Post-Training Quantization (PTQ) pipeline. The FP32 weights W_{FP32} and activations are mapped to 8-bit integers (INT8) using an asymmetric affine transformation:

$$q = \text{clamp}\left(\left\lfloor\frac{r}{S}\right\rceil + Z,\ q_{min},\ q_{max}\right) \quad (6)$$

where r is the continuous FP32 value, S (Scale) is a positive FP32 scalar representing the quantization step size, and Z (Zero-point) is an INT8 value mapping the real zero. $q_{min} = -128$ *and* $q_{max} = 127$.

This transformation reduces the model's memory footprint by a factor of 4× and allows the Raspberry Pi to execute the forward pass using ARM NEON SIMD (Single Instruction, Multiple Data) vector instructions, significantly accelerating matrix multiplications and reducing energy consumption per inference.

F. Actuation Logic: Dynamic Soil Water Balance

The ultimate objective of Edge AI is physical actuation. The output of the quantized Mamba-2 model is a multi-horizon forecast $\widehat{ET}_0 = [\widehat{ET_{t+1}}, \widehat{ET_{t+2}}, \widehat{ET_{t+3}}]$. This vector feeds a real-time root-zone soil water balance script running as a daemon on the Raspberry Pi:

$$S_{t+3} = S_t + \sum_{k=1}^{3}\left(P_{eff,t+k} + I_{t+k} - K_c \cdot \widehat{ET_{0,t+k}} - DP_{t+k}\right) \quad (7)$$

where S is the soil water depletion, P_{eff} is effective precipitation, I is irrigation, K_c is the crop coefficient, and DP is deep percolation. The Edge node evaluates S_{t+3} against the Readily Available Water (RAW) threshold. If $S_{t+3} <$ RAW, a high signal is sent to a switch, actuating the solenoid valves autonomously.

G. Experimental Setup and Hardware Architecture

1) Workstation vs. Edge Node Specifications

To rigorously evaluate the hardware–software co-design, two distinct computational environments were utilized. The centralized cloud/workstation (used for training and baseline evaluation) consists of an Intel Core i7-11850H CPU, 64 GB of DDR4 RAM, and an NVIDIA RTX A4000 GPU (8 GB GDDR6 VRAM), with execution performed using PyTorch 2.6.0 and CUDA 12.4. In contrast, the autonomous Edge node (target for inference) is based on a Raspberry Pi Zero WH, powered by a Broadcom BCM2835 (identical to Pi 1 but clocked at 1 GHz), equipped with 512 MB of RAM. Inference on the Edge is deployed using the TensorFlow Lite runtime (TFLite).

2) Evaluation Metrics

Agronomic predictive performance was evaluated using standard statistical metrics, including the Coefficient of Determination (R^2), Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Symmetric Mean Absolute Percentage Error (SMAPE). In parallel, hardware deployment viability was assessed through three key metrics: Memory Footprint (MB), representing storage and VRAM utilization during the forward pass; Inference Latency (ms), corresponding to the wall-clock time required to generate the 3-day multi-horizon forecast; and Energy Consumption (W), estimated as the peak power draw at the Edge node during inference.

III. RESULTS

A. Global Generalization and Climate-Aware Inductive Bias

The integration of the Köppen-Geiger categorical embeddings proved to be a transformative structural mechanism for spatial generalization. In ablation studies, the removal of the KG embedding caused the global model's performance to collapse in extreme climates. With the KG embeddings active, the FP32 TFT model improve by 13.2% of SMAPE. Mamba-2 followed closely with an improvement of 7.6% (Figure 1).

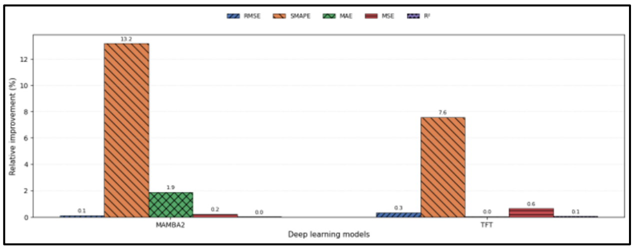

Figure 1: Performance gains by integrating KG

When cross-validated across the 309 sites, the single global model matched or outperformed site-specific models. Mamba-2 achieved $R^2 = 0.99$ (bias = 0.03) and TFT reached $R^2 = 1.00$ (bias = 0.00). While most sites show Global RMSE < Local RMSE, a few outliers exhibit slightly degraded performance, yet the overall R^2 near-perfect correlation validates that a unified model largely reduces the need for localized retraining (Figure 2).

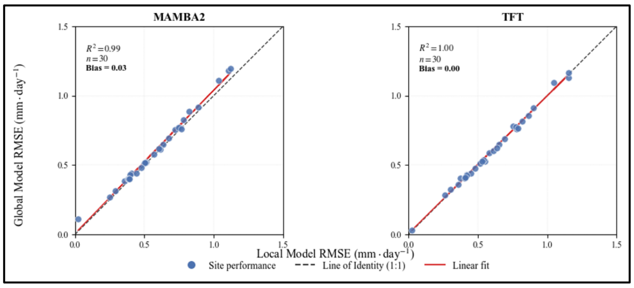

Figure 2: global vs site-specific performance

B. Validation of Distilled Noise-Invariance

To evaluate the impact of the Distilled Noise-Invariance methodology, we compared the performance of models trained with and without the EMD-WD target denoising. The results, expressed as relative improvement (%), demonstrate that the distillation paradigm significantly enhances prediction stability. The most substantial gain is observed in the Mean Squared Error (MSE), with a 12.0% improvement for the TFT architecture and 10.0% for Mamba-2. The reduction in quadratic errors (RMSE of 6.2% and 5.1% respectively) confirms that forcing the network to learn a noise-invariant latent manifold during training effectively compensates for input corruption. This eliminates the need for any explicit signal-processing filters on the Edge device, as the synaptic weights themselves perform the necessary auto-correction (Figure 3).

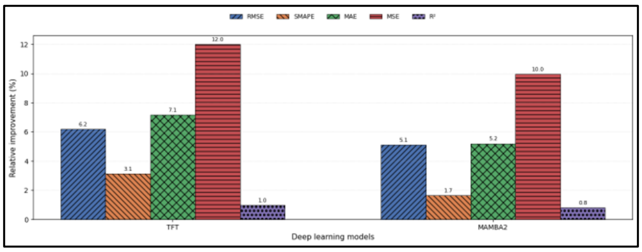

Figure 3: Performance gains by integrating KG

C. Hardware Profiling: The Superiority of INT8 Edge Deployment

The INT8 Edge deployment demonstrates significant efficiency gains compared to the centralized workstation baseline. The Raspberry Pi Zero WH reduces peak memory usage by 88.4% (113.30 MB vs. 978.71 MB), highlighting its suitability for resource-constrained environments.

Inference performance evaluated over 103,380 samples shows that the Edge-based model achieves multi-horizon forecasting in 86.41 ms per inference, compared to 0.0087 ms on the workstation (~9,930× slower). However, this latency remains below one second and is therefore compatible with real-time irrigation requirements.

Despite reduced numerical precision, the INT8 model maintains identical predictive performance across platforms ($R^2 = 0.9204, MAE = 0.4150mm$), corresponding to a 0% loss in accuracy.

In terms of energy consumption, processing the full dataset requires 10,719.28 J on the Edge device and 90.06 J on the workstation. This corresponds to approximately 0.10 J per inference (Edge) and 8.7×10^{-4} J per inference (workstation), reflecting the difference in execution time (0.0087 ms vs. 86.41 ms).

However, in practical deployment, inference is executed sequentially rather than in batch mode. Under such conditions, energy consumption is governed by the system's power profile and duty cycle rather than isolated per-inference efficiency. The Edge device operates within a low-power envelope, while the workstation requires sustained high power (~90 W), limiting its suitability for autonomous or off-grid operation.

Therefore, the Edge system is better aligned with real-world deployment scenarios involving intermittent inference and constrained energy availability, such as off-grid agricultural environments.

D. Autonomous Actuation Simulation

In a simulated edge-actuation environment utilizing historical data from Riyadh, Saudi Arabia (Arid, BWh), the Edge node faced an extreme heatwave scenario. The INT8 Mamba-2 model projected an escalating atmospheric demand, calculating a severe depletion in soil water storage at day t+3 ($S_{t+3} = -67.22mm$, crossing the -50 mm RAW alert threshold). Operating entirely offline, the Python daemon instantly executed the actuation logic, sending a signal to simulate the application of 42.22 mm of irrigation. In a contrasting simulation for Lomé, Togo (Tropical, Aw), the system recognized the low evaporative demand and inhibited the relays, successfully maximizing the utilization of effective precipitation and conserving groundwater (Table 4).

Table 1: Cross-climatic validation of the Edge AI irrigation decision-making system.

Site	**Climate**	(S_j) $(mm$	$\left(k_c \sum ET_0\right)$ $(3days, mm)$	(S_{j+3}) (mm)	**Decision**
Lomé	Tropical (Aw)	-35	11,51	-46,51	No irrigation (0)
Longyearbyen	Polar (ET)	-35	4,46	-39,46	No irrigation (0)
Moscow	Continental (Dfb)	-35	15,50	-50,50	Irrigate (25,50 mm)
Paris	Temperate (Cfb)	-35	14,04	-49,04	No irrigation (0)
Riyadh	Arid (Bsh)	-35	32,22	-67,22	Irrigation (42,22 mm)

IV. Discussion

A. The Paradigm Shift from Cloud-Tethered to Autonomous Edge

The results of this study suggest a shift toward Edge-based deployment in agricultural cyber-physical systems. Cloud-based inference (e.g., AWS, Google Cloud) introduces dependencies related to connectivity, latency variability, operational costs, and data privacy, particularly in rural environments.

In this work, a globally trained Temporal Fusion Transformer (TFT) model is deployed on a Raspberry Pi Zero WH using INT8 quantization. The Edge device requires 113.30 MB of peak RAM versus 978.71 MB on the workstation, corresponding to an 88.4% reduction, while maintaining identical predictive performance. This demonstrates that local inference is feasible under constrained hardware conditions, with a trade-off in execution speed.

B. The Epistemology of "Real-Time" in Agronomy

The workstation achieves a mean inference time of 0.0087 ms per sample, compared to 86.41 ms on the Edge device. Despite this difference, both remain negligible relative to agronomic process timescales (minutes to days). Therefore, "real-time" performance should be interpreted with respect to physical system dynamics rather than computational latency alone. In this context, the observed Edge delay remains compatible with irrigation control requirements.

C. The Elegance of Distilled Noise-Invariance

The Distilled Noise-Invariance paradigm resolves one of the main bottlenecks in IoT deployments: the computational cost of data cleaning. Traditional Edge systems expend battery-draining CPU cycles running moving averages, Kalman filters, or FFTs on noisy sensor inputs before inference. By moving EMD-WD preprocessing entirely to the offline target-generation phase, we exploit the neural network as a universal function approximator that internalizes an inverse-noise mapping. This zero-cost filtering mechanism is a novel contribution with broad applicability.

D. Future Trajectories: Towards TinyML and Parameter-Efficient Fine-Tuning

While the Raspberry Pi Zero already demonstrates the feasibility of intelligent irrigation inference at the edge, future work may further explore more resource-constrained deployments. This includes investigating model compilation through TinyML frameworks (e.g., TensorFlow Lite for Microcontrollers) for deployment on microcontrollers such as

ESP32 or STM32, which operate under significantly lower power budgets compared to single-board computers.

Additionally, Parameter-Efficient Fine-Tuning methods such as Low-Rank Adaptation (LoRA) could be explored to enable lightweight on-device adaptation. Integrating feedback from in-situ capacitive soil moisture sensors may allow periodic or incremental model updates at the edge, supporting adaptive calibration under changing environmental conditions.

V. Conclusion

This paper establishes a rigorous, globally scalable, and hardware-optimized framework for deploying deep sequence learning in autonomous precision irrigation. We show that the severe hardware constraints of Edge computing can be overcome through hardware–software co-design. By introducing a Distilled Noise-Invariance methodology—where target signals are denoised offline via EMD-WD—we eliminate real-time signal processing, endowing the model with intrinsic auto-correction of sensor errors. Köppen–Geiger climatological embeddings provide structural inductive bias, eliminating the need for localized retraining. The TFT architecture, optimized via INT8 quantization, maintains high predictive performance on the Edge device, achieving an $R^2 = 0.9204$ with no measurable loss in accuracy. On the Raspberry Pi Zero WH, the model requires 113.30 MB of peak memory and achieves an inference time of 86.41 ms per inference. This configuration enables fully local execution without reliance on cloud-based inference, making it suitable for offline deployment in resource-constrained environments. The resulting framework provides a viable approach for edge-based precision irrigation under limited connectivity conditions.

Acknowledgment

This work was supported by the Agence Française de Développement (AFD) through the PEA-IMPACT programme (Partenariat avec l'Enseignement Supérieur Africain – Ingénierie en Mathématiques Appliquées: Programme d'Accompagnement et de Coopération avec le Togo), implemented by the Agence Nationale de la Recherche (ANR).

References

[1] B. R. Scanlon *et al.*, « Global water resources and the role of groundwater in a resilient water future », *Nat. Rev. Earth Environ.*, vol. 4, nº 2, p. 87-101, févr. 2023, doi: 10.1038/s43017-022-00378-6.

[2] R. G. Allen, L. S. Pereira, D. Raes, et M. Smith, « Crop Evapotranspiration – Guidelines for Computing Crop Water Requirements », FAO, 56, 1998.

[3] G. Papacharalampous, H. Tyralis, Y. Markonis, P. Máca, et M. Hanel, « Features of the Earth's seasonal hydroclimate: characterizations and comparisons across the Köppen–Geiger climates and across continents », *Prog. Earth Planet. Sci.*, vol. 10, nº 1, p. 46, août 2023, doi: 10.1186/s40645-023-00574-y.

[4] M. Taheri *et al.*, « An Overview of Evapotranspiration Estimation Models Utilizing Artificial Intelligence », *Water*, vol. 17, nº 9, mai 2025, doi: 10.3390/w17091384.

[5] S. Zhao, J. Wu, R. Qiu, T. Zhang, Y. Luo, et W. Hu, « Time-frequency insights: Uncovering the drivers of reference evapotranspiration across China », *Agric. Water Manag.*, vol. 310, p. 109367, avr. 2025, doi: 10.1016/j.agwat.2025.109367.

[6] A. Gu et T. Dao, « Mamba: Linear-Time Sequence Modeling with Selective State Spaces », 31 mai 2024, *arXiv*: arXiv:2312.00752. doi: 10.48550/arXiv.2312.00752.

[7] B. Lim, S. Ö. Arık, N. Loeff, et T. Pfister, « Temporal Fusion Transformers for interpretable multi-horizon time series forecasting », *Int. J. Forecast.*, vol. 37, nº 4, p. 1748-1764, oct. 2021, doi: 10.1016/j.ijforecast.2021.03.012.

[8] P. Warden et D. Situnayake, *Tinyml: Machine learning with tensorflow lite on arduino and ultra-low-power microcontrollers*. O'Reilly Media, 2019.

[9] N. E. Huang *et al.*, « The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis », *Proc. R. Soc. Math. Phys. Eng. Sci.*, vol. 454, nº 1971, p. 903-995, mars 1998, doi: 10.1098/rspa.1998.0193.

[10] H. E. Beck *et al.*, « High-resolution (1 km) Köppen-Geiger maps for 1901–2099 based on constrained CMIP6 projections », *Sci. Data*, vol. 10, nº 1, p. 724, oct. 2023, doi: 10.1038/s41597-023-02549-6.

[11] B. S. Kumar, S. K, et S. M. Nair, « Optimizing micro irrigation efficiency in water-scarce agroecosystems using AI and remote sensing technologies », *Int. J. Biol. Sci.*, vol. 7, nº 4, p. 33-37, 2025, doi: 10.33545/26649926.2025.v7.i4a.324.