Journal of Agrometeorology
(A publication of Association of Agrometeorologists)
ISSN : 0972-1665 (print), 2583-2980 (online)
Vol. No. 27 (3) : 320-326 (September - 2025)
https://doi.org/10.54386/jam.v27i3.3081
https://journal.agrimetassociation.org/index.php/jam

Research Paper

AgroMetLLM: An evapotranspiration and agro-advisory system using localized large language models in resource-constrained edge

PARTHA PRATIM RAY* and MOHAN PRATAP PRADHAN

Department of Computer Applications, Sikkim University, Gangtok 737102, Sikkim
**Corresponding Author: ppray@cus.ac.in*

ABSTRACT

We introduce AgroMetLLM, an on-device agrometeorological advisory system that combines five validated evapotranspiration (ET) models with various quantized Large Language Models (LLMs) on a Raspberry Pi 4B. The users specify a location, 3-7-day horizon by using Gradio interface and LLM; Open-Meteo APIs then supply daily inputs T_{max}, T_{min}, T_{mean}, RH_{mean} precipitation, Food and Agriculture Organization (FAO) reference evapotranspiration (ET_0), and Rs for multiple Indian sites. Computed ET ranges (mm day^{-1}) across locations were: FAO ET_0 2.84-6.21; Hargreaves-Samani 6.28-13.74; Turc 0.17-0.21; Priestley-Taylor 5.64-9.06; Makkink 2.73-4.38. A few-shot prompting strategy, based on curated examples for 3-, 5-, and 7-day forecasts, is used to guide the Qwen LLM under Ollama to produce structured, five-point advisories in 1-2 s. One-way ANOVA (F = 3.30-6.71, p = 0.016-0.0002) and Kruskal-Wallis tests (χ^2 = 9.61-15.48, $p < 0.05$ except Turc p = 0.088) confirm significant ET differences among models and LLM sizes. All outputs and metadata persist in SQLite, and Matplotlib renders comparative bar charts in the dashboard. These results demonstrate that compact, quantized LLMs can reliably deliver actionable irrigation guidance-matching cloud-based accuracy-while operating offline, with minimal latency and energy use, thus empowering resource-constrained smallholder farmers.

Keyword: AgroMetLLM, Evapotranspiration, Raspberry Pi 4B, Quantized LLMs, Ollama Inference, Open-Meteo API

Accurate estimation of crop water requirements underpins sustainable irrigation management, especially in water-scarce regions where both over- and under-irrigation can detrimentally affect yields and deplete vital water resources. The evapotranspiration integrates soil evaporation and plant transpiration, serving as a primary driver for scheduling irrigation events (Alavi *et al.,* 2024; Bijlwan *et al.,* 2024; Mayani and Itagi 2024). Established ET models-ranging from the temperature-based Hargreaves-Samani estimator to energy-balance approaches such as Priestley-Taylor and Makkink-depend on reliable meteorological inputs that are frequently unavailable or intermittently recorded in rural agricultural landscapes (Naik *et al.,* 2025; Patel and Bunkar 2025; Rajput *et al.,* 2024). Concurrently, many agronomic advisory platforms rely on cloud-hosted machine-learning services to analyze weather data and dispense recommendations (Bakr *et al.,* 2025; Jan *et al.,* 2024); however, these solutions presuppose continuous internet connectivity and introduce concerns regarding latency, data privacy, and operational costs (Jiu *et al.,* 2024; Jhajharia, 2025; Shaloo *et al.,* 2024; Singh *et al.,* 2024).

Recent breakthroughs in model quantization and edge-AI inference have enabled deployment of compact, sub-billion-parameter LLMs on low-power devices, opening new avenues for offline, real-time decision support. Yet, these small LLMs often lack the zero-shot reasoning capacity required for domain-specific tasks unless guided by carefully engineered prompts. Furthermore, existing edge-based agrometeorological tools rarely offer an integrated suite of multiple ET algorithms paired with an interactive user interface, limiting their practicality for extension officers and smallholder farmers. There is therefore a pressing need for a cost-effective, easy-to-use platform that leverages both robust ET computation and lightweight, on-device LLM inference to empower rural growers with actionable irrigation guidance without reliance on cloud infrastructure.

Notwithstanding its offline, edgebased design, AgroMetLLM relies on periodic OpenMeteo API access for inputs, currently supports only five evapotranspiration models and one prompt paradigm on quantized LLMs, and remains to be validated

Article info - DOI: https://doi.org/10.54386/jam.v27i3.3081
Received: 14 June 2025; Accepted: 26 July 2025; Published online : 1 September 2025

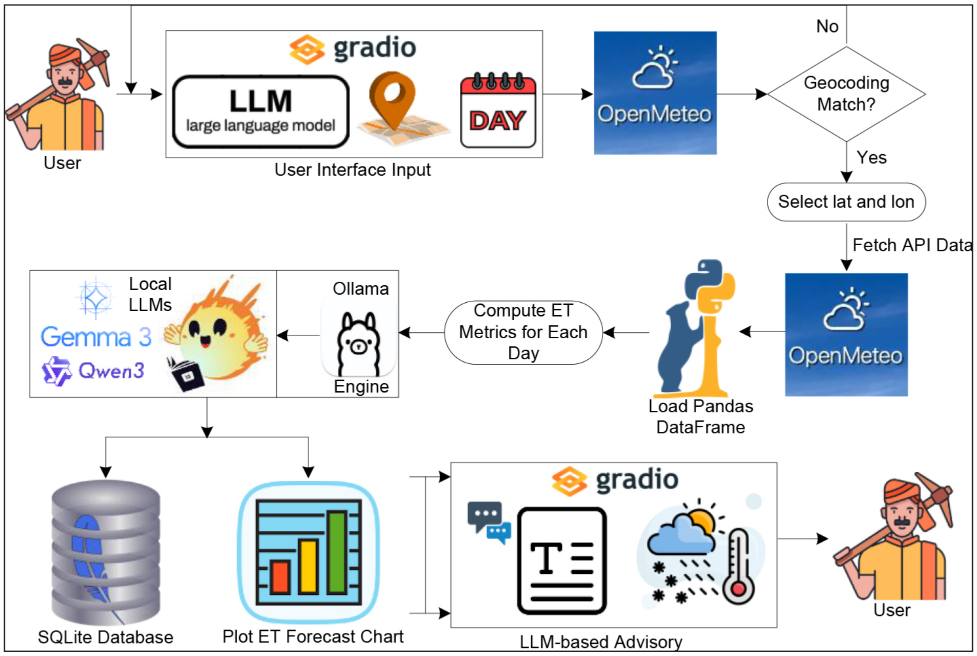

Fig. 1: AgroMetLLM system flow

in operational field trials-limitations that set realistic expectations for its immediate deployment in heterogeneous farming contexts.

In response, we present AgroMetLLM, an edge-computing system that consolidates five validated ET models with a 4-bit quantized LLM under the Ollama inference engine (Ollama, 2025), all hosted on a single Raspberry Pi 4B. Through a Gradio-based dashboard, users specify a location, forecast horizon, and preferred LLM; Open-Meteo APIs (Open-Meteo, 2025) supply geocoding and meteorological data, which are then processed in Python to compute ET metrics. A few-shot prompting strategy (Al Nazi *et al.,* 2025) -anchored by 3-, 5-, and 7-day exemplar pairs-enables the quantized model to generate structured, five-point agronomic advisories. The principal objectives of this work are: (i) to implement five standard ET algorithms on a low-cost, power-efficient platform, (ii) to harness compact, quantized LLMs with robust prompt engineering for domain-specific advisory generation, and (iii) to deliver a seamless, offline interface for comprehensive ET forecasting, visualization, and actionable recommendations tailored to resource-constrained, smallholder farming contexts.

MATERIALS AND METHODS

Software tools

The AgroMetLLM platform was deployed on a Raspberry Pi 4B-equipped with a quad-core ARM Cortex-A72 processor and 8GB of RAM-running Raspbian Bookworm and Python 3.12. All computations run on a single Raspberry Pi 4B, who's fullload power draw peaks at just 3-6 W. This ultralow energy footprint not only slashes operational costs compared to cloud or desktop servers, but also aligns with sustainable, offgrid deployment in energyscarce rural settings. We deployed five quantized LLMs—gemma3:1b (815 MB), granite3.1-moe1b (1.4 GB), qwen2.5:0.5b (398 MB), qwen2:0.5b (352 MB), and smollm2:360m (726 MB. Even after reserving ≈1 GB for the operating system and Gradio interface, the Raspberry Pi provided sufficient headroom to load the largest model (1.4 GB) and run inference, with peak memory usage remaining below the 4 GB limit. Key Python libraries and tools included: (i) Gradio to provide an interactive web interface accessible over the local network, (ii) Requests for RESTful API calls to both Open-Meteo and the local Ollama inference engine, (iii) Pandas and NumPy for efficient tabular data handling and numerical operations, (iv) Matplotlib to generate comparative evapotranspiration charts, (v) SQLite3 (built-in) for persistent storage of computed metrics and LLM advisories. Meteorological inputs were provided by the Open-Meteo geocoding API (to resolve place names to latitude/longitude) and forecast API (to retrieve daily: maximum, minimum, and mean 2 m air temperature; mean relative humidity; precipitation sum; FAO reference evapotranspiration; and total shortwave radiation). A locally-hosted, 4-bit quantized Qwen model ran under the Ollama inference engine (v0.9.0) on http://localhost:11434, ensuring low-latency, on-device advisory generation without external cloud dependencies.

Location resolution and data retrieval

User-supplied plain-text location strings (e.g., "Gangtok") are sent to Open-Meteo's geocoding endpoint with parameters for up to five matches and English-language labels. The response yields tuples (name, lat, lon, country) (name, lat, lon, country). If no result is returned, the interface displays a clear error ("Location not found. Please check spelling or try another location.") and halts further processing. Weather Fetch. Once a valid latitude/longitude pair is selected, a second API call requests daily aggregates between today's date and the user-specified horizon (3-7 days). The request parameters specify the following daily variables: (i) temperature_2m_max, temperature_2m_min, temperature_2m_mean (°C), (ii) relative_humidity_2m_mean (%), (iii) precipitation_sum (mm day^{-1}), (iv) et0_fao_evapotranspiration (mm day^{-1}), (v) shortwave_radiation_sum (MJ m^{-2} day^{-1}). These are returned in

JSON and loaded into a Pandas DataFrame with one row per forecast day. FAO Penman-Monteith (ET_0) is the globally recommended standard, it integrates solar radiation, air temperature, humidity and wind speed into a physically based model. Its comprehensive inputs make it highly accurate across diverse climates and the benchmark for crop water-requirement estimation. Fig. 1 presents the AgroMetLLM system flow.

Evapotranspiration calculations

Day-of-year n via pd.to_datetime (df["time"]).dt.day of year.

Mean daily air temperature: Computed locally from API values. This simple average reduces daily extremes into a representative value used by all subsequent models.

$T_{avg} = \frac{T_{max}+T_{min}}{2}$; where, T_{max},T_{min} obtained from the weather API and minimum 2 m air temperature (°C).

Extraterrestrial radiation Ra: Calculates solar radiation at the top of atmosphere for day-of-year and latitude ϕ.

Latitude in radians: $\phi = lat\_rad = \frac{\pi}{180} \times lat$; where, lat - latitude (°), from the geocoding API.

Inverse relative distance factor Earth-Sun: $d_r = 1 + 0.033 \times \cos(\frac{2\pi}{365} \times n)$

It accounts for the varying Earth-Sun distance over the year, where - day of year (1 - 365), computed via pandas from the API date.

Solar declination angle: $\delta = 0.409 \times \sin(\frac{2\pi}{365} \times n - 1.39)$; where, declination in radians of the sun, which varies seasonally.

Sunset hour angle: $\omega_s = arccos(-\tan\phi \tan\delta)$; where, ϕ - latitude in radians.

Extraterrestrial radiation Ra : $R_a = \frac{24\times60}{\pi} \times G_{sc} \times d_r[\omega_s \sin\phi \sin\delta + \omega_s \cos\phi \cos\delta \sin\omega_s]$

where, $G_{sc}=0.0820\ MJ\ m^{-2}\ min^{-1}$ is Solar constant, *Ra* in $MJ\ m^{-2}\ min^{-1}$.

Hargreaves-Samani method : It is an empirical, temperaturedriven formula needing only daily maxima and minima (and extraterrestrial radiation). Its minimal data needs suit arid or datascarce regions, though it may bias under high humidity variability. Locally computed, using (from API) and . Empirical estimate of reference ET_0 in .

$$ET_{HS} = 0.0023 \times (T_{avg} + 17.8)\sqrt{T_{max} - T_{min} \times R_a}$$

Turc method: Turc model is an energybalance approach combining mean temperature, relative humidity, and solar radiation. It performs well where radiation data are available but tends to underestimate evaporation in very dry environments. Locally computed if API provides mean temperature *T*, , mean *RH*, and shortwave radiation *Rs*. Energy-balance ET based on mean temperature, humidity, and radiation.

$$ET_{Truc} = 0.013 \times \frac{T + 15}{T + 15 + 50 \times \left(1 - \frac{RH}{100}\right)} \times R_s$$

Where, *T*- mean temperature (°C) received from API as in (°C), *RH*- relative humidity from API in (%), *Rs* - from API in *shortwave_radiation_sun* ($MJ\ m^{-2}\ min^{-1}$), *ET in mm* day^{-1}.

Slope of saturation vapor pressure curve : Locally computed from T_{avg} in kPa °C^{-1}.

$$\Delta = \frac{4098 \times 0.6108 \times \exp\left(\frac{17.27\, T_{avg}}{T_{avg} + 237.3}\right)}{\left(T_{avg} + 237.3\right)^2}$$

Psychrometric Constant **:** Locally computed from assumed atmospheric pressure.

$\gamma = 0.000665 \times P$; where, P in kPa uses P=**101.3** kPa, γ in kPa °C^{-1}.

Priestley-Taylor method**:** It is a semiempirical simplification of Penman-Monteith, it uses net radiation and a constant coefficient (α≈1.26) to approximate advection effects. It balances physical realism with reduced data demands, especially effective in humid conditions. Locally computed approximating net radiation by .

$ET_{PT} = \alpha_{PT} \times \frac{\Delta}{\Delta+\gamma} \times (R_n \times 0.408)$; where, $\alpha_{PT} = 1.26$, $R_n \approx$ API *shortwave_radiation_sun* ($MJ\ m^{-2}\ min^{-1}$).

Makkink method: Makkink is another radiationbased, semiempirical model that scales incoming radiation by the slope of the saturation vaporpressure curve and a psychrometric constant. Its conservative estimates often serve as a lower bound in ensemble analyses. Locally computed similarly using .

$ET_{Makkink} = 0.61 \times \frac{\Delta}{\Delta+\gamma} \times (R_n \times 0.408)$; where, units and inputs as in Priestley-Taylor.

Algorithm: AgroMetLLM

The AgroMetLLM algorithm begins by accepting three user specified inputs: a location name (string), the forecast horizon in days (integer between three and seven), and the choice of a locally hosted, quantized LLM via Ollama. It first resolves the location through a geocoding API call-halting with an error message if no matches are found-and extracts the selected latitude and longitude. Using these coordinates, it retrieves daily meteorological data (maximum, minimum, and mean temperature; relative humidity; precipitation; FAO reference evapotranspiration; and shortwave radiation) for the specified period. The data are loaded into a DataFrame, and for each day the algorithm computes the day-of-year, extraterrestrial radiation *Ra* (via the solar declination, inverse Earth-Sun distance, and sunset hour angle), and mean air temperature T_{avg}. It then calculates four empirical and energy-balance estimates of potential evapotranspiration - Hargreaves-Samani, Turc (when mean temperature, humidity, and radiation are available), Priestley-Taylor, and Makkink-using the appropriate combinations of T_{avg}, shortwave radiation, and psychrometric constants. After appending these ET metrics to the DataFrame, it generates a concise daily summary string and submits it to the Ollama inference engine, which returns a structured agronomic advisory. Finally, both the computed metrics and advisory text are stored in a local SQLite database, and a comparative bar chart of all five ET estimates is rendered for the end user.

Advisory generation, logging, and visualization

Once all evapotranspiration calculations are complete, the system compiles a daily summary string for each forecasted date. This standardized representation ensures that the local LLM receives uniform, machine-readable input describing maximum and minimum temperature, precipitation, the FAO reference evapotranspiration, and the four computed model estimates. All computed ET metrics and LLMgenerated advisories are stored locally in an SQLite database, enabling users to review historical forecasts and recommendations without any network connection. Extension officers or farmers can query past entries via simple SQL or export tables to CSV for further analysis or recordkeeping. Simultaneously, Matplotlib bar charts and violin plots-annotated with exact values and clear legends-are embedded in the Gradio dashboard, so that nonspecialist users can instantly see which days demand higher irrigation, how different models compare on any given date, or spot trends over a multiday window. Because these graphics render entirely ondevice, they can be saved as images or printed, providing an offlinefriendly, visual decision aid in connectivitylimited settings.

Visualization and front-end update

To aid interpretation, Matplotlib produces a grouped bar chart displaying FAO ET_0 and the four model estimates side-by-side for each forecast day. Each bar is annotated with its numerical value, allowing for quick visual comparison of model behavior under varying weather conditions. Finally, the Gradio interface is updated to present: (i) the tabular ET metrics, (ii) the comparative chart, (iii) real-time model information, and (iv) the LLM-generated advisory text. This integrated dashboard provides users with an at-a-glance understanding of forecasted water requirements and actionable agronomic guidance.

RESULTS AND DISCUSSION

The AgroMetLLM web interface runs on a Raspberry Pi 4B. Across the top, the user selects a quantized local model (here “granite3.1-moe1b”) from a dropdown, enters a location in the text box, and adjusts the forecast horizon via a slider (set to 7 days). Immediately below, the model info bar reports family, parameter count (1.3 B), quantization level (Q8_0) and architecture. After clicking Search Location, the resolved match (“4. Pusa, India (25.9862, 85.6796)”) appears in a secondary dropdown, which in turn enables the *Generate ET Forecast & Local LLM Advisory* button.

Once executed, the “Evapotranspiration Metrics (mm/day)” table populates seven rows (2025-06-12 through 2025-06-18), showing for each date the geographic coordinates, maximum and minimum 2 m air temperatures (°C), precipitation (mm day^{-1}), FAO reference ET_0 (API), and the four computed ET estimates (Hargreaves-Samani, Turc, Priestley-Taylor, Makkink). Directly beneath, a grouped bar chart-titled “Comparative ET Models Forecasting - Pusa (25.9862, 85.6796)”-plots all five ET series side by side for each date, with values annotated above each bar and a legend indicating model colors.

Finally, the “Farmer Advisory (LLM Output)” textbox at the bottom presents five numbered recommendations (Irrigation plan, Crop/soil action, Livestock/labour, Pest & disease watch, Input-saving tip), generated by the local LLM using the standardized summary input. This integrated dashboard allows users to review numerical results, visualize multi-model ET comparisons, and immediately access tailored agronomic advice.

3-Day average forecast ET outputs by location

The Fig. 2 illustrates the three-day average evapotranspiration (ET) forecasts from five models across seven Indian locations. Consistently, the Hargreaves-Samani method predicts the highest ET rates-ranging from 6.28 mm day^{-1} in Gangtok to 13.74 mm day^{-1} in Pusa-while the Turc model yields negligible values (~0.17-0.21 mm day^{-1}) at all sites. The FAO reference ET_0 falls in an intermediate range of approximately 2.84 mm day^{-1} (Gangtok) up to 6.21 mm day^{-1} (Bikaner), and the Priestley-Taylor estimates span 5.64 mm day^{-1} (Gangtok) to 9.06 mm day^{-1} (Pusa). Makkink results are lower than Priestley-Taylor but higher than FAO ET_0 (2.73 mm day^{-1} in Gangtok to 4.38 mm day^{-1} in Pusa). This pattern highlights the strong sensitivity of temperature-based algorithms (Hargreaves-Samani) compared to radiation- or energy-balance approaches, with arid locations (e.g., Bikaner, Pusa) showing markedly higher ET across all models.

Distribution of ET model data

Fig. 3 presents violin-plot distributions of the five ET model forecasts-Hargreaves-Samani, Turc, FAO ET_0, Priestley-Taylor, and Makkink-across the seven locations (Pusa, Muzaffarpur, Bikaner, Tiruchi, Cuttack, Karnal, and Gangtok). The Hargreaves-Samani estimates exhibit the highest medians and the greatest spread overall, with Pusa and Muzaffarpur showing the widest variability (≈11-15 mm day^{-1}) and Gangtok the lowest (≈6 mm day^{-1}). In contrast, the Turc method produces extremely low values at all sites (medians ~0.15-0.20 mm day^{-1}) with minimal dispersion, reflecting its sensitivity to humidity and radiation. FAO ET_0 forecasts are intermediate, with medians ranging from ~3 mm day^{-1} in Gangtok to ~8 mm day^{-1} in Bikaner and a moderate interquartile spread.

Priestley-Taylor shows a broad distribution similar to Hargreaves-Samani but shifted downward (≈4-11 mm day^{-1}), while Makkink remains the most conservative energy-balance model, with medians of ~2-4.5 mm day^{-1} and comparatively narrow violins. These patterns underscore both the systematic differences among ET formulations and the influence of local climate conditions on model variability.

LLM wise comparison of ET model data

Fig. 4 compares the five-model average ET forecasts produced by each of the five local LLMs (gemma3:1b, granite3.1-moe1b, qwen2.5:0.5b, qwen2:0.5b, and smollm2:360m). For the temperature-driven Hargreaves-Samani method, the three largest models cluster at high ET rates (~12.5-13.0 mm day^{-1}), while the smaller qwen2:0.5b and smollm2:360m yield lower averages of 9.14 mm day^{-1} and 10.26 mm day^{-1}, respectively. FAO reference ET_0 likewise declines with model size, from 5.38 mm day^{-1}

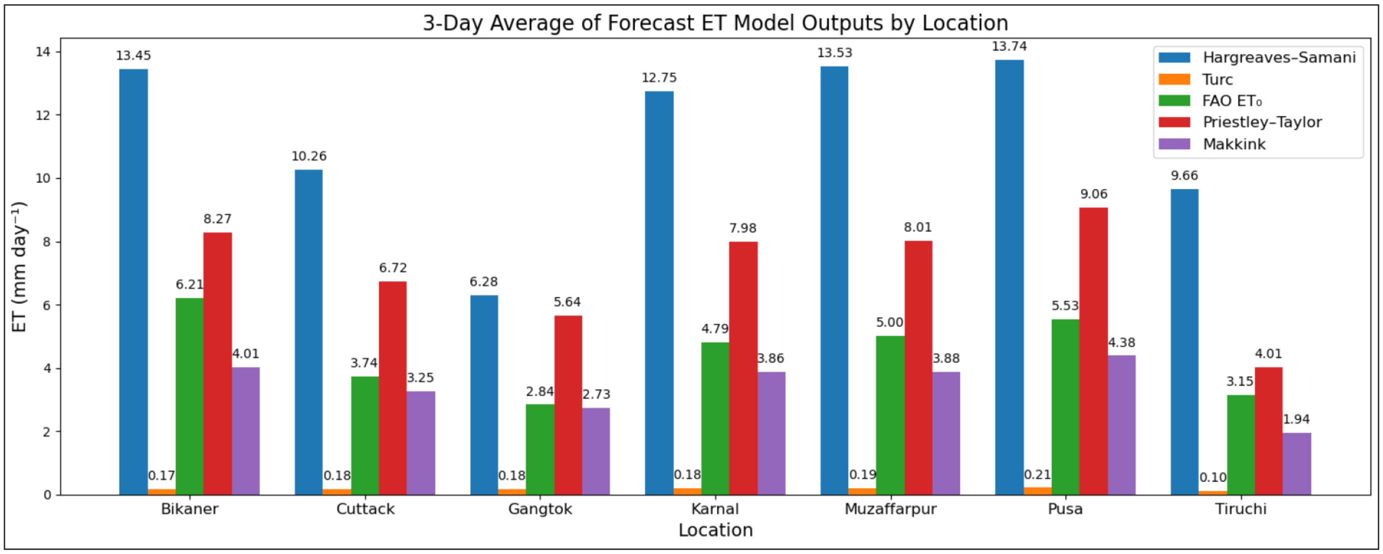

Fig. 2: Forecast of ET models by location.

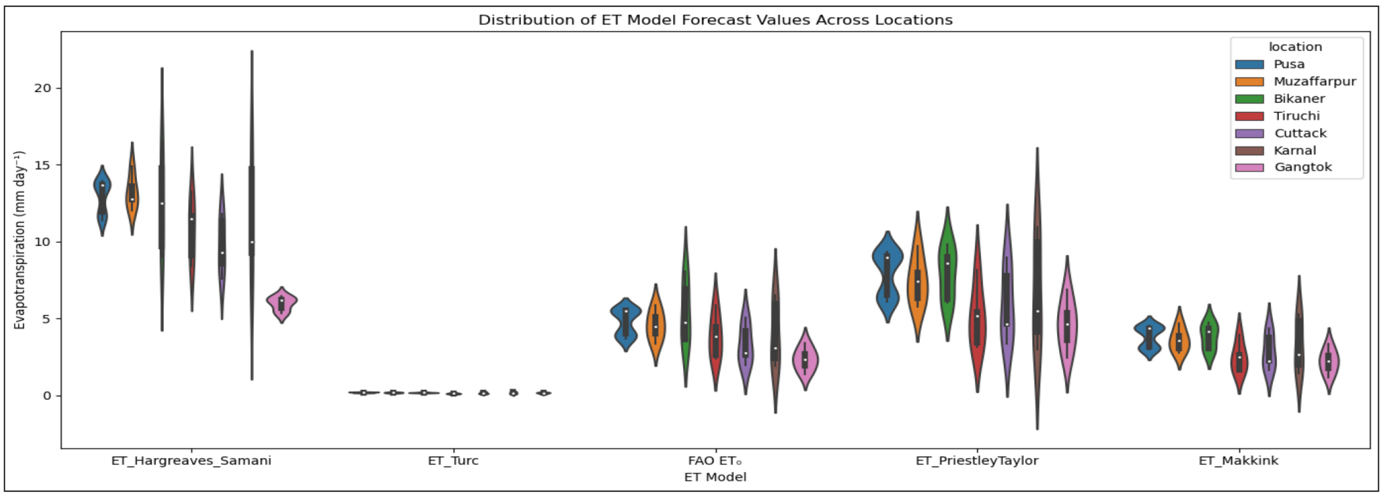

Fig. 3: Distribution of forecast data of ET models by location.

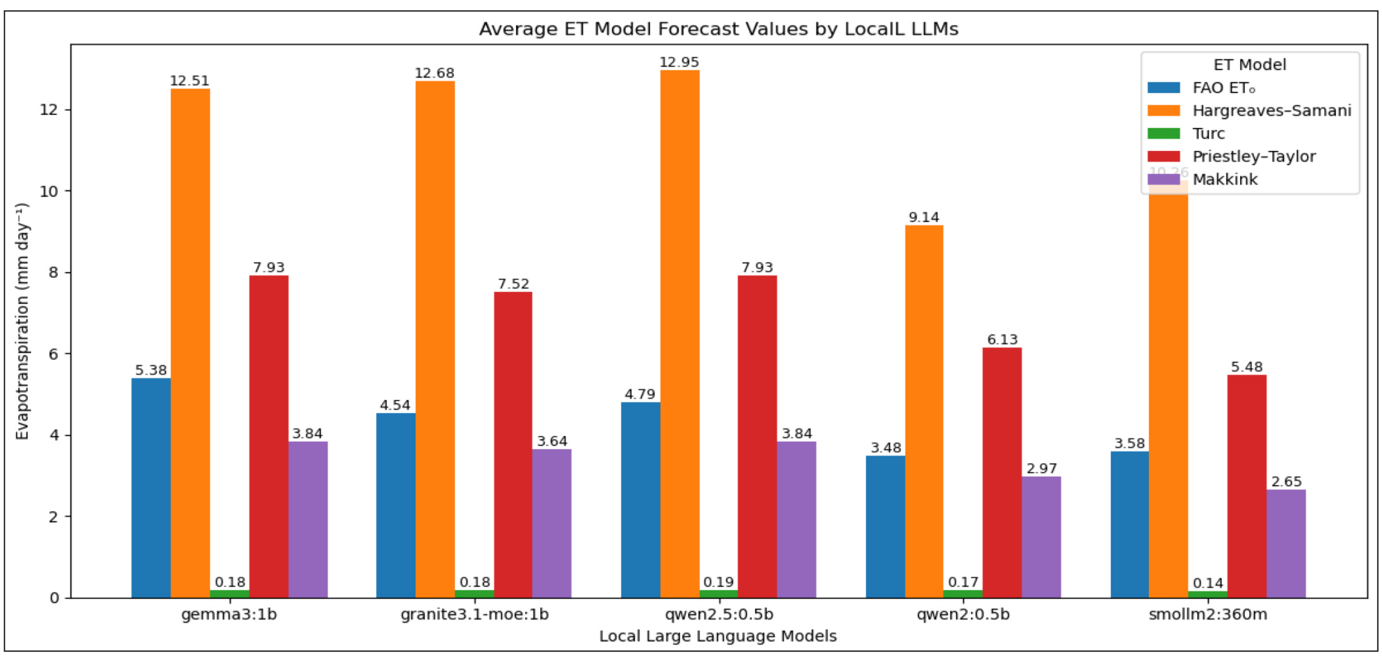

Fig. 4: Average ET model wise forecast data by LLMs.

Table 1: ANOVA and Kruskal-Wallis test results of ET models across LLMs

ET Models	F_statistic	p_value	Significance	Kruskal_statistic	p_value	Significance
FAO ET0	3.638907	0.010145	*	9.611525	0.047505	*
Hargreaves-Samani	6.708787	0.000157	*	15.4817	0.0038	*
Turc	3.304287	0.016345	*	8.09704	0.088087	-
Priestley-Taylor	3.705408	0.009231	*	11.84606	0.018533	*
Makkink	3.705408	0.009231	*	11.84606	0.018533	*

(gemma3) down to 3.58 mm day^{-1} (smollm2). The Turc estimates remain negligible (0.14-0.19 mm day^{-1}) across all models, reflecting its sensitivity to humidity and radiation inputs. Energy-balance methods show a similar size-dependent pattern: Priestley-Taylor averages range from ~7.9 mm day^{-1} in the larger LLMs to 5.48 mm day^{-1} in smollm2, and Makkink from ~3.8 mm day^{-1} down to 2.65 mm day^{-1}. Overall, larger, higher-parameter LLMs systematically produce higher ET estimates for temperature- and radiation-based models, whereas all models converge on near-zero Turc values.

Statistical analysis

The one-way ANOVA (Table 1) revealed that mean evapotranspiration (ET) estimates differ significantly across the four tested quantized LLMs for every metric: FAO ET_0 ($F = 3.64$, $p = 0.010$), Hargreaves-Samani ($F = 6.71$, $p < 0.001$), Turc ($F = 3.30$, $p = 0.016$), Priestley-Taylor ($F = 3.71$, $p = 0.009$), and Makkink ($F = 3.71$, $p = 0.009$). In all cases, the null hypothesis of equal means was rejected at = 0.05, indicating that at least one model's daily ET output differs from the others for each method. Because some ET distributions departed from normality or exhibited unequal variances, we also applied the nonparametric Kruskal-Wallis test (Table 1). Consistent with the ANOVA, we found significant differences in distribution for FAO ET_0 (= 9.61, $p = 0.048$), Hargreaves-Samani (= 15.48, $p = 0.004$), Priestley-Taylor (= 11.85, $p = 0.019$), and Makkink (= 11.85, $p = 0.019$). The Turc method, however, did not reach significance (= 8.10, $p = 0.088$), suggesting broadly similar Turc outputs across the four models. The Turc model consistently produced very low ET estimates (≈0.14-0.21 mm day^{-1}) with minimal spread across all seven locations and five quantized LLMs. Because Turc's formulation heavily weights incoming radiation and mean humidity-and our test sites exhibited similar radiation-humidity regimes-its output remains nearly constant regardless of slight variations in the summary inputs or the LLMdriven advisory context.

CONCLUSION

AgroMetLLM successfully demonstrates that a Raspberry Pi 4B running open-source tools and quantized LLMs can deliver end-to-end ET forecasting and agronomic advice without cloud dependencies. By combining five validated ET algorithms with multi-shot prompt engineering, the system produces consistent, structured recommendations tailored to local meteorology. The significant differences observed-both among ET formulations and across LLM sizes-highlight the importance of model selection for field applications. Persistence in SQLite and real-time visualization via Gradio ensure transparency and ease of interpretation for end users. Future work will extend the platform to incorporate additional climate variables (e.g., wind speed, soil moisture), support multiple languages, and validate advisory accuracy in operational farm trials. Ultimately, AgroMetLLM paves the way for scalable, LLM-driven decision support in precision agriculture under resource constraints. Despite these strengths, the current system still depends on intermittent Open-Meteo API access for inputs, supports only five ET algorithms and a single few-shot prompt paradigm on quantized LLMs (e.g., Phi-2, TinyLLaMA, Mistral-instruct), and remains to be validated in operational field trials-limitations that may constrain its robustness and generalizability across diverse agroclimatic zones.

ACKNOWLEDGEMENT

The authors acknowledge the support received from the Department of Computer Applications, Sikkim University for providing necessary computation and o4-mini-high for polishing the article.

Conflict of Interests: The authors declare that there is no conflict of interest related to this article.

Funding: None.

Data availability: Data was generated in real-time by using the code in Python virtual environment in 2025. Researchers can use the code to generate their data in real-time.

Code availability: https://github.com/ParthaPRay/agrometerology_ollama_raspberrypi_forecasting

Authors contribution: **P.P. Ray:** Conceptualization, Methodology, Visualization, Writing-review and editing; **M.P. Pradhan:** Supervision

Authors certificate: We hereby certify that this manuscript, "AgroMetLLM: An Evapotranspiration and Agro-Advisory System using Localized Large Language Models in Resource-Constrained Edge," is our original work. We declare no conflicts of interest and confirm adherence to ethical and publication standards.

REFERENCES

Alavi, M., Albaji, M., Golabi, M., Naseri, A.A. and Homayouni, S. (2024). Estimation of sugarcane evapotranspiration from remote sensing and limited meteorological variables using machine learning models. *J. Hydrol.,* 629: 130605.

Al Nazi, Z., Hossain, M.R. and Al Mamun, F. (2025). Evaluation of open and closed-source LLMs for low-resource language with zero-shot, few-shot, and chain-of-thought prompting. *Nat. Lang. Process. J.,* 10: 100124.

Bijlwan, Amit, Shweta Pokhriyal, Rajeev Ranjan, R.K. Singh and Ankita Jha. (2024). Machine learning methods for estimating reference evapotranspiration. *J. Agrometeorol.*, 26(1): 63-68. https://doi.org/10.54386/jam.v26i1.2462

Bakr, D.I., Al-Khalidi, J. and Abed, H.N. (2025). Estimation of climatological parameters using ANN and WEKA models in Diyala Governorate, *Iraq. J. Agrometeorol.*, 27(2): 210-215. https://doi.org/10.54386/jam.v27i2.2877

Jan, S., Khan, U., Khalil, A., Khan, A.A., Jan, H.A. and Ullah, I. (2024). Use of machine learning techniques in predicting inflow in Tarbela reservoir of Upper Indus Basin. *J. Agrometeorol.,* 26(4): 501-504. https://doi.org/10.54386/jam.v26i4.2676

Jiu, T.A., Wei, W.A., Min, X.U., Ning, H.U., Shang, C.H., Xu, J., Hu, X., Lv, H., Z.H. Zi-han and Yu-jing, L.A. (2024). Simulation of evapotranspiration in a well-facilitated paddy field based on machine learning algorithms and energy balance closure. *Chin. J. Agrometeorol*., 45(10): 1109.

Jhajharia, K. (2025). Wheat yield prediction of Rajasthan using climatic and satellite data and machine learning techniques. *J. Agrometeorol.,* 27(1): 63-66. https://doi.org/10.54386/jam.v27i1.2807

Mayani, M.B. and Itagi, R.L. (2024). Sugarcane acreage estimation using satellite imagery and machine learning. *J. Agrometeorol.,* 26(4): 505-508. https://doi.org/10.54386/jam.v26i4.2669

Naik, R., Majhi, B. and Naidu, D. (2025). Prediction of pan evaporation in Chhattisgarh using machine learning techniques. *J. Agrometeorol.,* 27(1): 86-91. https://doi.org/10.54386/jam.v27i1.2731

Ollama (2025). Ollama: Run LLMs locally. https://ollama.com/ (Accessed: June 14, 2025)

Open-Meteo (2025). Open-Meteo: Free weather forecasts via API. https://open-meteo.com/ (Accessed: June 14, 2025)

Patel, R.M. and Bunkar, K. (2025). Soybean yield prediction leveraging advanced ensemble machine learning models. *J. Agrometeorol.,* 27(2): 227-229. https://doi.org/10.54386/jam.v27i2.2971

Rajput, J., Singh, M., Lal, K., Khanna, M., Sarangi, A., Mukherjee, J. and Singh, S. (2024). Data-driven reference evapotranspiration (ET_o) estimation: A comparative study of regression and machine learning techniques. *Environ. Dev. Sustain.,* 26(5): 12679-12706.

Shaloo, Kumar, B., Bisht, H., Rajput, J., Mishra, A.K., T.M., K.K. and Brahmanand, P.S. (2024). Reference evapotranspiration prediction using machine learning models: An empirical study from minimal climate data. *Agron. J.,* 116(3): 956-972.

Singh, Y.P., Singh, P.K. and Tomar, A.S. (2024). Performance comparison of linear regression and ANN models in estimating monthly reference evapotranspiration (ET_o). *J. Agrometeorol.,* 26(3): 383-386. https://doi.org/10.54386/jam.v26i3.2641