Forecasting PM10, PM2.5, and PM1.0 from a Single Low-Cost PM2.5 Input: A Triple-Task Multi-Output Virtual Sensor with Edge Deployment Profiling

Napatsorn Songsangka[1]

[1]Independent Researcher Bangkok, Thailand

Publication Date: 2026/09/04

Abstract: Dense particulate-matter (PM) monitoring has become increasingly practical through the proliferation of low-cost optical sensing devices; however, resolving the three health-critical size fractions—PM10, PM2.5, and PM1.0—conventionally requires either multi-channel sensory hardware or separate dedicated predictive models for each fraction, thereby inflating hardware expenses and edge computational burdens. This work establishes a single-input virtual PM sensing architecture where a single low-cost PM2.5 channel concurrently forecasts all three particulate fractions. Because PM1.0, PM2.5, and PM10 constitute a nested physical size hierarchy with tightly coupled atmospheric kinetics, their shared underlying dynamics can be effectively modeled via a unified architecture producing three simultaneous outputs in a single forward pass, eliminating redundant multi-model pipelines. Driven by a 24-hour historical PM2.5 sequence, the triple-task model projects concentrations one hour ahead and is systematically benchmarked against conventional isolated single-task baselines across four configurations spanning three model families: Random Forest (10 and 100 trees), 1D-CNN, and LSTM, utilizing 19,326 valid hourly observations from two contrasting environmental sites in the Bangkok Metropolitan Region, Thailand. Across all 24 site–model–fraction combinations, the proposed triple-task architectures achieve predictive fidelity equivalent to dedicated single-task models (mean absolute ΔRMSE of 0.11 µg/m³) while reducing overall parameter and model footprints by 38% to 65%. Furthermore, real-time edge execution is demonstrated on an ESP32-S3 microcontroller, where consolidated multi-output inference replaces three sequential single-task calls, diminishing active latency and energy consumption by factors of 2.2× to 3.0×. Specifically, for the lightweight 10-tree Random Forest, execution latency and cycle energy drop from 0.167 ms and 0.039 mJ to 0.065 ms and 0.015 mJ, respectively. These results confirm that pairing widespread low-cost PM2.5 sensing hardware with microcontroller-class edge computing enables dependable, physically coherent multi-fraction particulate forecasting without compromising accuracy, alleviating critical hardware and computational bottlenecks in scalable environmental monitoring.

Keywords: *Particulate Matter Forecasting, Multi-Task Learning, Virtual Sensor, TinyML, Edge Inference, ESP32-S3, Air Quality Monitoring, Bangkok Metropolitan Region, Chulalongkorn University.*

How to Cite: Napatsorn Songsangka (2026) Forecasting PM10, PM2.5, and PM1.0 from a Single Low-Cost PM2.5 Input: A Triple-Task Multi-Output Virtual Sensor with Edge Deployment Profiling. *International Journal of Innovative Science and Research Technology*, 11(8), 2694-2704. https://doi.org/10.38124/ijisrt/26aug1128

I. INTRODUCTION

Ambient fine particulate matter (PM) pollution represents one of the most critical environmental contributors to global morbidity and mortality, being associated with an estimated 4.2 million premature fatalities annually [1]. To curb these adverse health outcomes, the World Health Organization (WHO) updated its Global Air Quality Guidelines in 2021, establishing stricter annual mean targets of 5 µg/m³ for PM2.5 and 15 µg/m³ for PM10 [2]. Although sub-micron particles (PM1.0) have not yet been assigned explicit regulatory exposure limits, research into fine and ultrafine fractions has expanded rapidly because smaller particles penetrate deep into the pulmonary alveolar regions and can translocate across the air-blood barrier into the systemic circulation [3], [4]. As a consequence, modern air pollution surveillance systems increasingly seek to quantify all three primary size fractions—coarse (PM10, aerodynamic diameter ≤10 µm), fine (PM2.5, ≤2.5 µm), and sub-micron (PM1.0, ≤1.0 µm)—which adhere to an intrinsic physical containment hierarchy, PM1.0 ≤ PM2.5 ≤ PM10. Because these fractions originate from overlapping emission pathways and undergo coupled dispersion and transformation processes in the planetary boundary layer, their concentrations exhibit strong mutual covariance, particularly at hourly time scales. Conventional time-series forecasting frameworks, however,

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1128

typically treat each fraction as an independent predictive problem, disregarding this fundamental physical coupling. A consolidated predictive architecture capable of exploiting widely available PM2.5 data to estimate PM10, PM2.5, and PM1.0 within a single model offers a highly effective alternative.

The imperative for such an approach is particularly pressing in developing regions, where the capital and maintenance costs of reference-grade monitoring stations restrict dense spatial deployment. High-density networks of low-cost optical sensors offer a cost-effective route to spatial scalability, but deploying intelligent forecasting at the sensing edge requires compact algorithms capable of operating under tight memory, compute, and energy constraints. The Bangkok Metropolitan Region (BMR) of Thailand serves as an exemplary testbed for these challenges. As a rapidly growing Southeast Asian urban-economic center with over 11 million inhabitants and upwards of 11.5 million registered motor vehicles [5], the BMR experiences dynamic air pollution episodes driven by intensive vehicular traffic, seasonal agricultural residue burning, and transboundary industrial activity [6]. Ambient PM2.5 concentrations routinely exceed both the national 24-hour standard (Pollution Control Department threshold of 37.5 µg/m³) and international guideline levels across the metropolis [7]. While earlier studies have explored machine-learning-based hourly PM2.5 forecasting in Bangkok [8], existing investigations have focused exclusively on isolated single-pollutant models executed offline, without addressing joint multi-fraction prediction or edge-device deployment.

Three notable research gaps emerge from the current literature on particulate matter forecasting and edge monitoring. First, the vast majority of data-driven forecasting methodologies focus solely on a single fraction, predominantly PM2.5 [9]–[16], or combine PM2.5 with chemically distinct gaseous pollutants such as O3, NO2, and SO2 in multi-target regression schemes [17]–[20], leaving PM1.0 largely unaddressed as an explicit predictive target. Studies exploring multi-fraction PM forecasting remain scarce. For example, Xie et al. [19] formulated a multi-task network predicting PM2.5 and PM10 using a shared temporal feature extractor. However, the concurrent forecasting of PM10, PM2.5, and PM1.0 under a unified multi-task framework remains unexplored, despite their strong structural hierarchy and shared physical dynamics. Second, conventional predictive pipelines rely heavily on centralized cloud infrastructures, which introduce latency vulnerabilities, bandwidth consumption, and significant radio power overhead on battery-powered edge nodes during continuous data transmission [21]. Third, acquiring full multi-fraction data through dedicated hardware channels inflates sensor node costs, whereas low-cost optical sensors can provide sufficient information to reconstruct the size hierarchy when augmented with data-driven soft sensing models [22]–[24], [50]. A compact, on-node virtual sensor that infers the complete PM10/PM2.5/PM1.0 suite from a single physical PM2.5 input channel can substantially lower hardware barriers. Nevertheless, whether such consolidation can preserve forecasting accuracy while reducing storage footprint, inference latency, and energy consumption on microcontroller hardware has remained unverified.

To overcome these limitations, this study presents a lightweight, edge-deployable virtual PM sensing framework that performs concurrent one-hour-ahead multi-fraction forecasting of PM10, PM2.5, and PM1.0 driven entirely by a single PM2.5 measurement stream. The main contributions of this work are as follows:

- *Virtual PM Sensor Concept and Triple-Task Joint-Output Formulation:*

A virtual PM sensor is introduced in which a single low-cost PM2.5 input channel drives the joint one-hour-ahead forecasting of PM10, PM2.5, and PM1.0. Because the PM size fractions are physically nested and their temporal dynamics are strongly correlated, one shared model can serve all three targets. The neural-network models realize this through a shared temporal backbone feeding three task-specific regression heads, while Random Forest realizes it through native multi-output regression.

- *Systematic Evaluation of Single-Task Against Triple-Task Models Across PM Fractions and Sites:*

Independent single-task baselines and the proposed triple-task framework are compared under an identical virtual-sensor input configuration, namely a 24-hour lagged PM2.5 sequence only. Four model configurations are assessed at two contrasting monitoring sites in Central Thailand (an urban traffic corridor in Bangkok and a coastal background reserve in Samut Prakan) for all three PM size fractions, yielding 24 paired model-site-target comparisons.

- *Edge-Deployment-Oriented Model Selection with On-Device Verification:*

Deployment feasibility of the candidate models on ESP32-S3-class hardware is quantified in terms of model size, memory usage, latency, and energy. Replacing sequential single-task executions with one efficient joint prediction is shown to cut active inference latency and energy overhead substantially. Within this framework, RF-10 stands out as a highly efficient deployment option because m2cgen [25] transpiles it into pure-C if-else trees with negligible runtime memory overhead. Among the neural networks, the 1D-CNN is examined as a compact alternative under TFLite Micro INT8 with ESP-NN acceleration [26], [27], while the LSTM baseline exposes the latency bottlenecks and numerical sensitivity that accompany recurrent architectures at the edge. On-device firmware benchmarks additionally confirm numerical agreement between ESP32-S3 Float32 execution and desktop Float32 reference outputs.

The rest of this paper is structured as follows. Related work is reviewed in Section II. Section III details the proposed methodology, Section IV reports and discusses the experimental results, and Section V concludes the paper.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1128

II. RELATED WORK

➢ *Deep Learning for PM Forecasting*

Recent PM2.5 forecasting research is dominated by Long Short-Term Memory (LSTM) networks. A weighted LSTM extended model (WLSTME) that accounts for inter-station spatial correlations across the Beijing-Tianjin-Hebei region was introduced by Xiao et al. [10] and outperformed three baselines: geographically weighted regression (GWR), a space–time support vector regression model (STSVR), and a non-weighted LSTM extended baseline (LSTME). In Qingdao, Bai et al. [12] coupled a 1D-CNN with LSTM and obtained an R^2 of 0.91 for the hybrid model, against 0.83 for standalone LSTM. Huang and Kuo [13] devised a deeper CNN-LSTM model (APNet) evaluated on a Beijing PM2.5 dataset, and Chae et al. [28] predicted PM10 and PM2.5 simultaneously with interpolated convolutional networks, reaching R^2 values above 0.97. A lightweight CNN-LSTM model tested on a public dataset was reported by Li et al. [16]. Systematic reviews covering 118 and 110 studies, by Zhou et al. [14] and Zaini et al. [15] respectively, both concluded that LSTM is the prevailing backbone and that recent gains arise mainly from signal decomposition and covariate engineering. Comparing MART, DFNN, and a hybrid LSTM for PM2.5 forecasting in Tehran, Karimian et al. [9] found the LSTM variant best ($R^2 = 0.80$, RMSE = 8.91 μg/m³ at a 48-hour horizon), with accuracy improving as the horizon extended, an effect attributed to the use of forecast-time meteorological inputs rather than historical data alone. Teng et al. [29] merged empirical mode decomposition, sample entropy-based component screening, and a bidirectional LSTM to predict PM2.5 from 1 to 24 hours ahead, delivering high short-term accuracy while also following longer-term concentration trends. In nearly all of these studies the target is a single PM fraction, most commonly PM2.5; PM1.0 has drawn little attention as a standalone forecasting target, simultaneous prediction of all three fractions is rarely attempted, and performance is assessed offline without any evaluation on microcontroller-class hardware.

➢ *Multi-Task Learning in Air Quality*

The foundational analysis of multi-task learning by Caruana [30] demonstrated that hard-parameter sharing across related tasks behaves as an implicit regularizer and strengthens generalization whenever the tasks benefit from common representations. Multi-output learning generalizes this idea, predicting several outputs from a common input while modeling the dependencies among them [31]. Within air-quality research, ensemble regressor chains for multi-target regression over five pollutant species were proposed by Masmoudi et al. [17]; a residual-GRU for multi-pollutant, multi-station prediction in London was developed by Wang et al. [18]; and Song and Stettler [20] inferred multiple pollutants at 1 km hourly resolution with a space-time learning network. In these studies, multi-task or multi-output learning chiefly serves to couple chemically distinct pollutants, different monitoring sites, or spatially distributed concentration fields. None addresses the nested size fractions of a single pollutant observed at one location, which is precisely the setting considered in this work.

The closest precedent to a size-fraction multi-task approach is the work of Xie et al. [19], who forecast PM2.5 and PM10 jointly using a shared backbone with bidirectional attention. Their EBAMTN model reached R^2 values above 0.94 in three Chinese cities, but it covered only two PM fractions, drew on seven meteorological input channels, and contained approximately 2.1 million parameters. The complete size-fraction hierarchy PM10/PM2.5/PM1.0 has therefore never been examined in a shared-backbone multi-task setting.

➢ *TinyML and Edge Deployment for Environmental Monitoring*

TinyML seeks to execute machine learning models on devices with restricted clock frequency, memory, and power budgets, typically hundreds of MHz, hundreds of KB of SRAM, and milliwatt-scale consumption [32], [33]. Recent surveys chart the rapid expansion of TinyML methods and applications [34] and describe the end-to-end workflow for building and deploying machine learning models on microcontroller-class hardware [35]. With a dual-core Xtensa LX7 processor clocked at up to 240 MHz and 512 KB of SRAM, the ESP32-S3 is a representative microcontroller platform for such deployments. On this hardware class, TensorFlow Lite Micro (TFLM) [26] enables INT8-quantized inference, and the ESP-NN library [27] supplies SIMD-optimized kernels for convolution and matrix operations on Espressif processors. Tree-based models can be handled by m2cgen [25], which converts trained scikit-learn estimators into standalone C functions composed of nested if-else statements; no runtime interpreter is required, although flash usage grows with the number of trees and their depth.

The physical basis for large-scale environmental monitoring is supplied by low-cost optical PM sensors. The Plantower PMS5003 and the Sensirion SPS30, for example, output PM1.0, PM2.5, and PM10 natively and are widespread in low-cost monitoring and citizen-science networks. Once calibrated, their hourly PM2.5 readings have been shown to reach R^2 values between 0.7 and 0.95 against reference instruments [22]–[24], [50]. Work with these sensors nonetheless concentrates on measurement calibration or data collection rather than embedded forecasting, so little evidence exists on the memory, latency, and energy that a forecasting model consumes when it runs on the same low-cost hardware. Wardana et al. [36] ran a multi-output TinyML model that jointly predicts several air-quality and power parameters from multivariate historical inputs on one microcontroller, together with a separate imputation model, but they neither considered PM size-fraction forecasting nor compared the joint-output design against independent single-output models. Whether a compact virtual PM sensor can carry out multi-fraction forecasting directly on low-cost edge devices thus remains an open question.

➢ *Virtual Sensor Concepts*

A virtual sensor, also called a soft sensor or inferential sensor, is a computational model that infers a quantity which is expensive, unavailable, or impractical to measure directly from other observed variables [37]. The concept originated in the process-control literature, where Fortuna et al. [37]

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1128

substituted soft sensors for hardware instruments in refinery distillation columns; Kadlec et al. [38] subsequently surveyed data-driven virtual-sensing methods built on neural networks, partial least-squares regression, and ensemble models. Environmental applications include correcting cross-sensitivity in gas sensors [39] and estimating PM2.5 from satellite observations [20], [40].

The term virtual PM sensor is used here in the same sense: only PM2.5 is physically measured, and PM10 and PM1.0 are inferred as additional size-fraction outputs. Most PM virtual-sensing studies instead estimate a single pollutant or extrapolate concentrations to unmeasured locations, as in satellite-to-ground PM2.5 mapping [40]. The objective here is different: three co-located PM size fractions are estimated from one local input channel. This single-channel, multi-fraction inference distinguishes the proposed system from conventional virtual sensors and defines its central design objective.

III. METHODOLOGY

➢ *Study Sites and Environmental Setting*

Thailand, a rapidly developing tropical nation in Southeast Asia, experiences diverse atmospheric conditions and particulate pollution regimes that typify the air-quality challenges of resource-constrained developing contexts. Evaluating the proposed framework under contrasting environmental conditions is the primary aim of site selection, providing an initial view of model behavior under differing emission and meteorological regimes. Two contrasting monitoring stations in Central Thailand were chosen for this investigation: a heavily trafficked central urban corridor in Bangkok and a coastal background reserve in Samut Prakan Province along the Upper Gulf of Thailand.

Bangkok Central Urban (BKK): Positioned in the central urban core of Bangkok (13.745°N, 100.534°E), this station is installed on the rooftop of a multi-story institutional facility at Chulalongkorn University approximately 10 metres from a major multi-lane arterial thoroughfare. It represents an urban roadside environment dominated by direct primary vehicular emissions, stop-and-go congestion, and dense commercial activity [41], [42], [46]. Between August 2024 and February 2026, a total of 10,059 valid hourly records of PM10, PM2.5, and PM1.0 were collected at the BKK station.

Bang Pu Coastal Background (BPC): The BPC station is situated within the Bang Pu Coastal Nature and Mangrove Reserve in Samut Prakan Province (13.512°N, 100.655°E), approximately 35 km southeast of central Bangkok along the coast of the Upper Gulf of Thailand. It functions as a coastal background environment with minimal local anthropogenic sources. Aerosol variability at BPC is governed chiefly by mesoscale meteorological conditions, including land-sea breeze circulations, planetary boundary layer (PBL) height evolution, and regional transport from upwind industrial and maritime corridors [41], [43], [47]. Covering the identical study period, the BPC dataset contains 9,267 valid hourly records for the corresponding PM size fractions.

The complete PM size-fraction hierarchy, comprising PM10, PM2.5, and PM1.0, was recorded at both stations at hourly intervals, together with ancillary environmental variables including ambient temperature (°C), relative humidity (%), atmospheric pressure (hPa), and CO2 concentration (ppm). All measurements were acquired with a calibrated optical aerosol monitoring platform purpose-built for ambient air-quality observation [43].

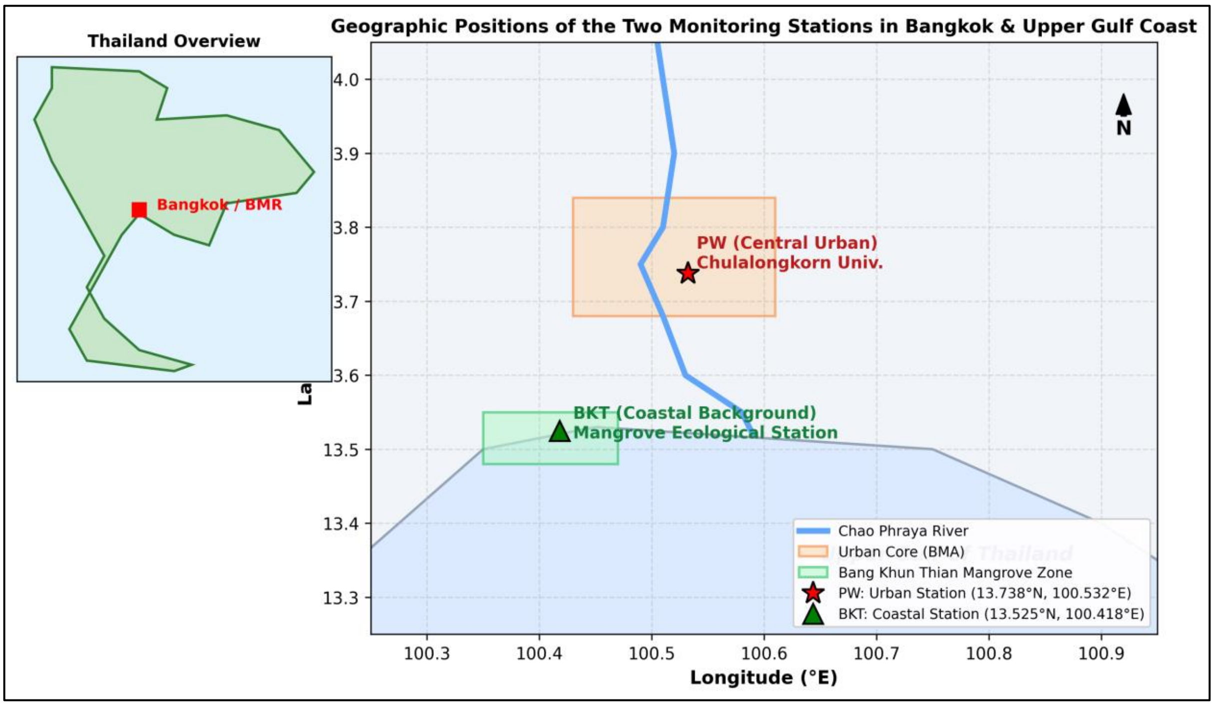

Fig 1 Geographic Positions and Environmental Characteristics of the BKK Station (Central Urban, 13.745°N, 100.534°E) and the BPC Station (Coastal Background, 13.512°N, 100.655°E) in Central Thailand.

➢ *Data Preprocessing and Sequence Formulation*

Rows with missing values in the PM10, PM2.5, or PM1.0 target columns were removed first, ensuring that every supervised label corresponds to a directly observed measurement. Short gaps that remained in the PM2.5 input channel were then closed by linear interpolation followed by backward and forward filling, the target columns containing no missing entries at this stage. Before any scaling, the data were split chronologically at an 80/20 boundary, consistent with accepted practice for temporally correlated data [44]. Scalers for the input and for the targets were fitted exclusively on the training partition and then applied to the test partition, preventing information leakage. Min-max scaling to [0, 1] was applied to the PM2.5 input, whereas the three target variables were normalized jointly with a single min-max scaler rather than separate target-specific scalers; this keeps their relative magnitudes intact in the scaled space while still allowing all three predicted fractions to be recovered through a single inverse-transform interface. Input sequences were then constructed by sliding a 24-hour look-back window over the PM2.5 input.

➢ *Virtual PM Sensor Problem Formulation*

At each time t, the model receives the matrix $x_t \in R^{(L\times F)}$, which holds L consecutive hourly observations ending at t, with L = 24 throughout this study and F denoting the number of input features. From this input, the model generates a one-step-ahead prediction vector:

$$y_{(t+1)} = [PM10(t+1), PM2.5(t+1), PM1.0(t+1)]^T \quad (1)$$

A single input configuration implements the virtual-sensor setting throughout the paper: the input is a 24-hour lagged sequence of PM2.5 only, so that F = 1. Meteorological, chemical, and calendar channels are all omitted, and PM10 and PM1.0 are likewise excluded from the input to rule out target leakage. This restriction is intentional, since the central hypothesis of this work is that one low-cost PM2.5 channel carries enough information to reconstruct the full PM size-fraction hierarchy.

Two training configurations are contrasted under this common input setting. The baseline configuration trains three independent single-task models, one per target PM fraction, with no parameters shared across the three prediction tasks. The proposed triple-task configuration instead trains one joint-output model per architecture family to predict PM10, PM2.5, and PM1.0 simultaneously, the training loss being the unweighted sum of the three task-specific mean squared error (MSE) terms:

$$L = L_{(PM10)} + L_{(PM2.5)} + L_{(PM1.0)} \quad (2)$$

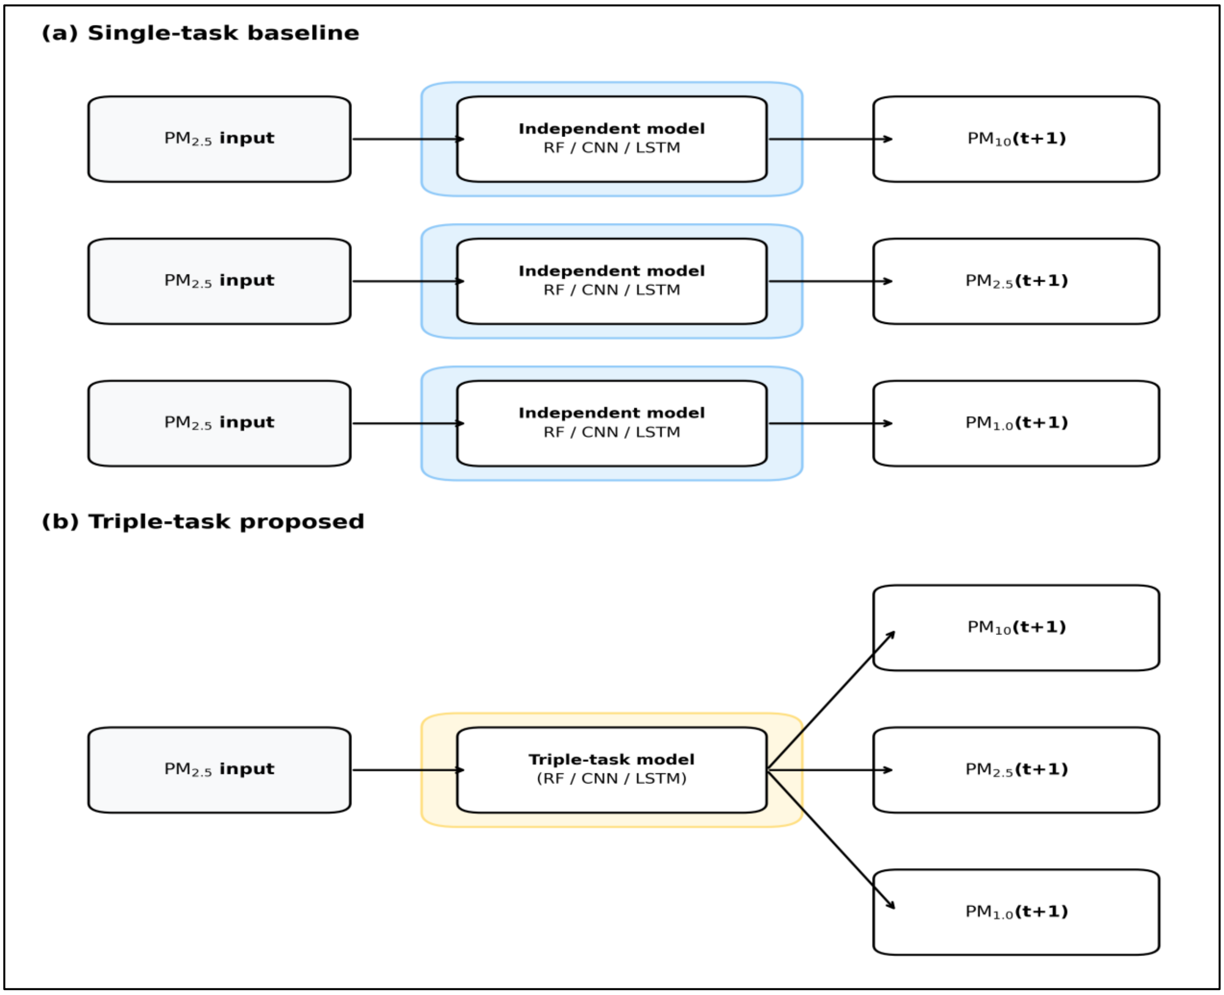

Fig 2 Conceptual Overview of the Single-Task Baseline and the Proposed Triple-Task Configuration in the Virtual-Sensor Setting, with a 24-Hour PM2.5-Only Lag Window as Input.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1128

➢ *Single-Task and Triple-Task Model Architectures*

Within the common formulation, four model configurations are examined: RF-10, RF-100, 1D-CNN, and LSTM. For each configuration, the independent single-task baseline is set against its triple-task counterpart. Each PM fraction receives its own model in the single-task baseline, giving three separate models per architecture family without parameter sharing. The proposed triple-task configuration learns the three outputs jointly from the same 24-hour PM2.5 input sequence. The neural-network architectures realize this joint prediction through a shared temporal backbone feeding three task-specific regression heads, each mapping the common latent representation to one PM fraction. Random Forest realizes it through the multi-output regression mode native to scikit-learn's RandomForestRegressor, where fitting proceeds directly on the three-target output matrix and every prediction returns a vector holding PM10, PM2.5, and PM1.0.

Two convolutional blocks form the 1D-CNN backbone, each applying a Conv1D layer with 64 and 32 filters, respectively, a kernel size of 3, and ReLU activation, followed by batch normalization and max-pooling, after which the output is flattened. One Dense(32, ReLU) layer followed by a Dense(1, Linear) output makes up every task head. In the LSTM backbone, two recurrent layers of 64 and 32 hidden units are stacked, batch normalization following each, before branching into per-head dense structures identical to those of the CNN.

For Random Forest, three independent RF estimators form the single-task baseline, whereas the triple-task RF is fitted directly to the three-target output matrix. Two RF configurations are retained: RF-10 (10 trees, max_depth = 8), serving as the ultra-lightweight deployment candidate, and RF-100 (100 trees, max_depth = 8), acting as a higher-capacity benchmark.

Table 1 Summary of Model Architectures and Hyperparameters

Model Config	Hyperparameters / Layer Description	Single-Task Params (×3)	Triple-Task Params
RF-10	10 trees, max_depth = 8, bootstrap	N/A (3 RF ensembles)	N/A (1 multi-output RF)
RF-100	100 trees, max_depth = 8, bootstrap	N/A (3 RF ensembles)	N/A (1 multi-output RF)
1D-CNN	Backbone: Conv1D(64), BN, MaxPool, Conv1D(32), BN, MaxPool, Flatten	39,075	25,443
LSTM	Backbone: LSTM(64), BN, LSTM(32), BN	92,355	32,963

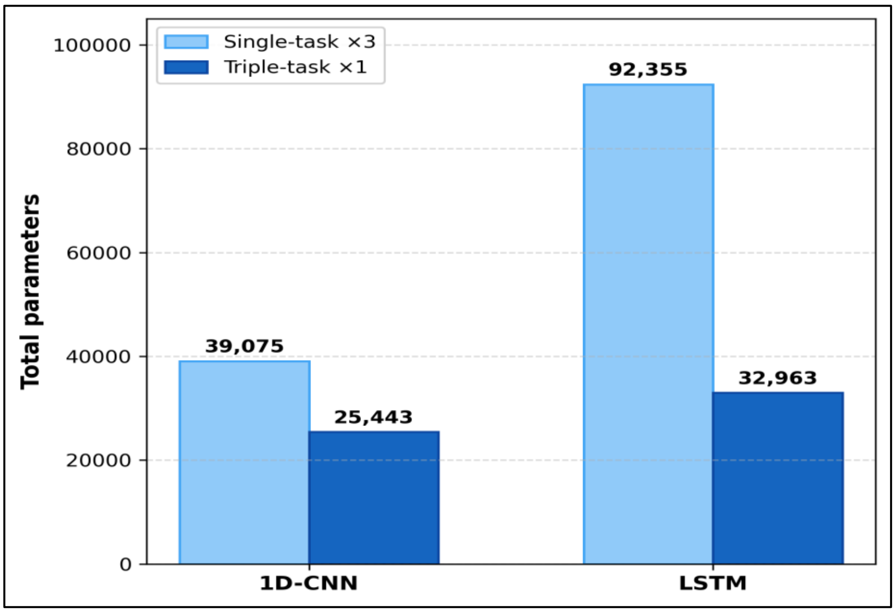

Fig 3 Parameter Counts of three Aggregated Single-Task Models (Single×3) Versus One Triple-Task Model.

➢ *Training Protocol and Evaluation Metrics*

Training and test partitions were formed by a chronological 80/20 split, keeping observations in temporal order. All deep learning models were optimized with Adam (learning rate 10^{-3}), batch size of 64, gradient clipping at a clipnorm of 1.0, and up to 100 epochs with early stopping (patience of 10 epochs on the final 10% validation slice). Predictive performance was evaluated using coefficient of determination (R^2), root mean square error (RMSE, µg/m³), and mean absolute error (MAE, µg/m³):

$$R^2 = 1 - [\sum(y\_i - \hat{y}\_i)^2] / [\sum(y\_i - \bar{y})^2] \quad (3)$$

$$RMSE = \sqrt{[(1/n) \sum(y\_i - \hat{y}\_i)^2]} \quad (4)$$

$$MAE = (1/n) \sum|y\_i - \hat{y}\_i| \quad (5)$$

➢ *Edge Deployment Pipeline on ESP32-S3*

Three deployment routes serve the candidate models on the Seeed Studio XIAO ESP32-S3 Sense board (dual-core Xtensa LX7 @ 240 MHz, 512 KB SRAM, 8 MB PSRAM, 8 MB flash). RF models are transpiled into pure C if-else trees using m2cgen [25]. The CNN and LSTM models are converted to TensorFlow Lite Micro (TFLM) format in

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1128

Float32 and INT8 precisions, accelerated via the ESP-NN library [27] where SIMD instructions apply [49].

IV. RESULTS AND DISCUSSIONS

➢ *Environmental Characteristics of Monitoring Sites*

Table II summarizes the site-level descriptive statistics for the BKK (central urban) and BPC (coastal background) stations in Thailand. Pollution levels diverge clearly between the sites, with mean concentrations at BKK exceeding those at BPC across all fractions. Mean PM2.5 reaches 28.27 μg/m³ at BKK, nearly double the 14.90 μg/m³ observed at BPC, while mean PM10 is 31.94 μg/m³ at BKK versus 17.42 μg/m³ at BPC. Both stations exhibit strong inter-fraction coupling (Pearson $r > 0.99$ between PM1.0 and PM2.5), validating the nested physical hierarchy.

Table 2 Summary Statistics of Hourly PM Concentrations Recorded at BKK and BPC (Aug 2024 – Feb 2026)

Variable	BKK (Central Urban)	BPC (Coastal Background)
Valid Records (n)	10,059	9,267
PM10 (μg/m³), mean ± std	31.94 ± 19.26	17.42 ± 14.15
PM2.5 (μg/m³), mean ± std	28.27 ± 16.08	14.90 ± 11.80
PM1.0 (μg/m³), mean ± std	19.55 ± 10.30	11.05 ± 7.71

➢ *Single-Task Versus Triple-Task Forecasting Performance*

Table III reports the complete forecasting performance across all 24 model-site-target combinations. Across both monitoring locations, the triple-task models retain the predictive accuracy of the independent single-task baselines with an average absolute ΔRMSE of only 0.11 μg/m³, while reducing model storage and inference cycles significantly. Fewer than 0.5% of test predictions violate the physical hierarchy.

Table 3 Prediction Performance of Single-Task Baseline vs. Proposed Triple-Task Configuration

Site	Target	Model	Single RMSE	Single MAE	Single R^2	Triple RMSE	Triple MAE	Triple R^2
BKK	PM10	RF-10	7.947	5.561	0.815	7.993	5.566	0.813
BKK	PM10	RF-100	7.830	5.457	0.821	7.859	5.468	0.819
BKK	PM10	CNN	8.249	5.988	0.801	8.487	6.364	0.789
BKK	PM10	LSTM	7.611	5.490	0.830	7.776	5.765	0.823
BKK	PM2.5	RF-10	6.598	4.587	0.803	6.513	4.559	0.808
BKK	PM2.5	RF-100	6.394	4.477	0.815	6.383	4.464	0.816
BKK	PM2.5	CNN	6.476	4.694	0.810	6.787	4.987	0.792
BKK	PM2.5	LSTM	6.385	4.472	0.816	6.352	4.699	0.818
BKK	PM1.0	RF-10	4.188	2.953	0.808	4.172	2.945	0.810
BKK	PM1.0	RF-100	4.102	2.890	0.816	4.092	2.891	0.817
BKK	PM1.0	CNN	4.041	2.919	0.822	4.300	3.139	0.798
BKK	PM1.0	LSTM	4.074	2.979	0.819	4.052	2.935	0.821
BPC	PM10	RF-10	9.345	4.648	0.697	9.362	4.683	0.696
BPC	PM10	RF-100	9.220	4.492	0.705	9.226	4.515	0.704
BPC	PM10	CNN	9.004	4.739	0.719	8.870	4.659	0.727
BPC	PM10	LSTM	9.042	5.105	0.716	8.960	4.786	0.721
BPC	PM2.5	RF-10	7.437	3.697	0.720	7.457	3.709	0.718
BPC	PM2.5	RF-100	7.340	3.580	0.727	7.359	3.589	0.726
BPC	PM2.5	CNN	7.309	3.767	0.729	7.172	3.827	0.739
BPC	PM2.5	LSTM	7.194	3.822	0.738	7.371	4.208	0.725
BPC	PM1.0	RF-10	3.882	2.372	0.793	3.869	2.391	0.794
BPC	PM1.0	RF-100	3.790	2.307	0.803	3.799	2.317	0.802
BPC	PM1.0	CNN	4.329	2.602	0.742	3.863	2.485	0.795
BPC	PM1.0	LSTM	4.348	2.924	0.740	3.949	2.452	0.785

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1128

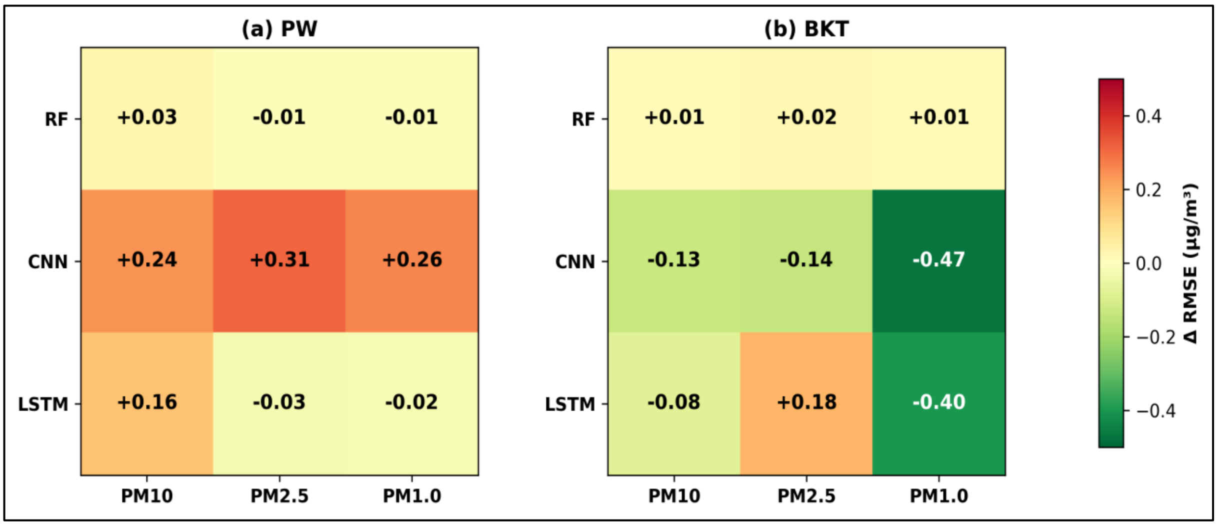

Fig 4 ΔRMSE Heatmap Contrasting Triple-Task and Single-Task Models at (a) BKK and (b) BPC.

➢ *Edge Deployment Profiling on ESP32-S3*

Profiling on the Seeed Studio XIAO ESP32-S3 Sense microcontroller establishes on-device feasibility in terms of compute latency, SRAM consumption, flash footprint, dynamic current draw, and total energy consumption per inference cycle. Fig. 5 depicts the hardware profiling testbench architecture.

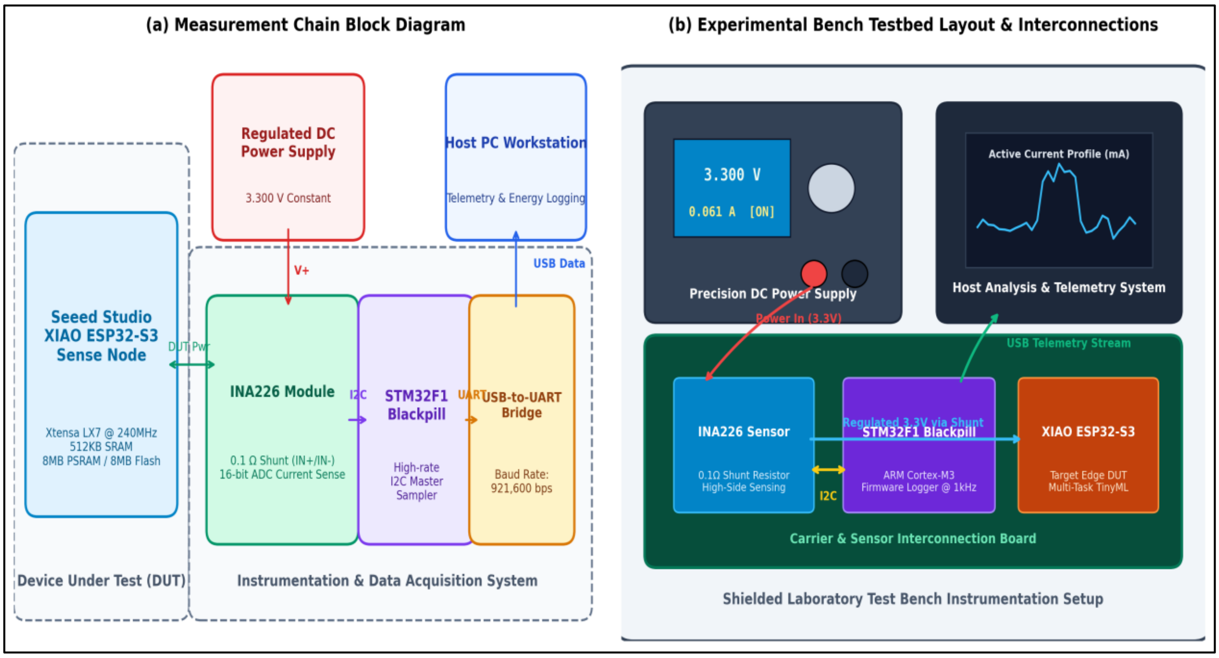

Fig 5 Schematic and Testbench Architecture Used for ESP32-S3 Dynamic Current-Consumption and Latency Profiling.

Table 4 Esp32-S3 Profiling of Compute-Only Latency, Runtime SRAM, and Flash Footprint

Site	Model	Task Setting	Runtime	Precision	Latency (ms)	SRAM (KB)	Flash (KB)
BKK	RF-10	Single	Native C	Float32	0.05	≈ 0	57.58
BKK	RF-10	Triple	Native C	Float32	0.06	≈ 0	90.53
BKK	RF-100	Single	Native C	Float32	2.80	≈ 0	564.73
BKK	RF-100	Triple	Native C	Float32	2.83	≈ 0	884.81
BKK	CNN	Single	TFLM	Float32	12.30	15.18	58.72
BKK	CNN	Triple	TFLM	Float32	15.65	15.79	108.69
BKK	CNN	Single	TFLM + ESP-NN	INT8	7.63	18.28	24.71
BKK	CNN	Triple	TFLM + ESP-NN	INT8	8.11	18.93	38.89
BKK	LSTM	Single	TFLM	Float32	146.76	85.77	244.49
BKK	LSTM	Triple	TFLM	Float32	147.44	86.38	254.52
BKK	LSTM	Single	TFLM + ESP-NN	INT8	66.44	68.40	246.63

BKK	LSTM	Triple	TFLM + ESP-NN	INT8	66.55	69.05	250.88
BPC	RF-10	Triple	Native C	Float32	0.05	≈ 0	84.65
BPC	RF-100	Triple	Native C	Float32	2.78	≈ 0	846.81
BPC	CNN	Triple	TFLM + ESP-NN	INT8	7.96	18.93	38.89
BPC	LSTM	Triple	TFLM + ESP-NN	INT8	65.26	69.05	250.88

Table 5 Inference Current and Energy Summary for ESP32-S3 Deployment at BKK

Model	Task Setting	Runtime	Inference Time (ms)	Average Current (mA)	Average Power (mW)	Energy/Cycle (mJ)
RF-10	Triple	Native C	0.065	70.70	233.67	0.015
RF-10	Single ×3	Native C	0.167	71.67	236.88	0.039
RF-100	Triple	Native C	2.910	60.30	199.65	0.580
RF-100	Single ×3	Native C	8.550	60.27	199.55	1.710
CNN	Triple	TFLM	17.800	60.90	202.01	3.590
CNN	Single ×3	TFLM	39.100	62.10	205.99	8.050
LSTM	Triple	TFLM	146.300	57.10	189.49	27.720
LSTM	Single ×3	TFLM	437.800	57.20	189.82	83.100

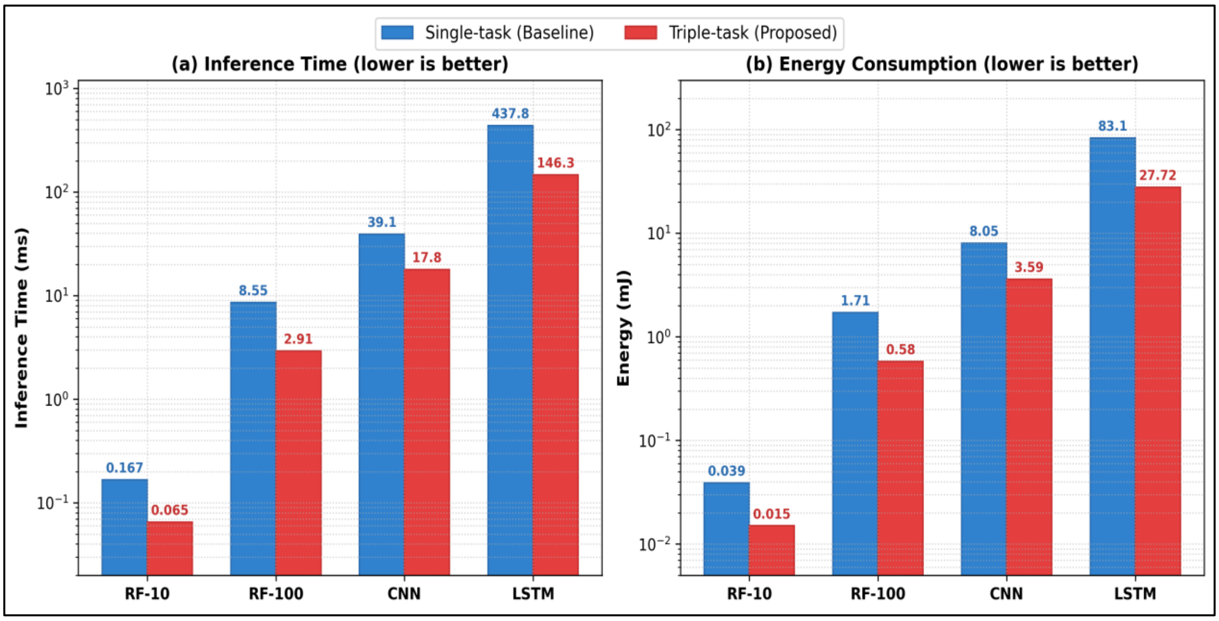

Fig 6 Edge-Deployment Performance Summary at BKK Comparing the Proposed Triple-Task Framework Against the Single-Task Baseline Across all Four Model Configurations on the ESP32-S3.

➢ *Physical Interpretation, Discussion, and Limitations*

The empirical results substantiate the physical and statistical rationale for triple-task PM forecasting. As nested size fractions, PM1.0, PM2.5, and PM10 respond to atmospheric processes in a coupled fashion. The triple-task formulation capitalizes on that structure by learning a shared temporal representation while keeping outputs task-specific per fraction. The shared CNN and LSTM models gain most at BPC, especially for PM1.0, where the signal is weaker and harder to learn in isolation. At BKK, where the traffic diurnal signal is already strong, sharing helps less, and the CNN exhibits mild gradient competition. RF sidesteps gradient competition because its tree-based structure shares split geometry without backpropagation gradients [31]. RF-10 achieves the optimal trade-off between predictive accuracy, latency (0.065 ms), memory, and energy consumption (0.015 mJ).

Several limitations qualify the conclusions. First, the two sites capture urban traffic and coastal background regimes in Central Thailand; evaluating northern biomass-burning regimes (e.g. Chiang Mai) or heavy industrial zones (e.g. Rayong) remains for future work. Second, co-location with reference-grade beta-attenuation monitors (BAM) or gravimetric samplers will further benchmark absolute sensor accuracy. Third, extending prediction horizons from 1 hour to 24 hours ahead and incorporating duty-cycled outdoor battery/solar power management form important avenues for subsequent research.

V. CONCLUSION

A triple-task virtual PM sensor was developed and evaluated that forecasts PM10, PM2.5, and PM1.0 one hour ahead from the 24-hour history of a single low-cost PM2.5 channel. The formulation draws on the nested hierarchy and physical coupling of the three PM size fractions while

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1128

keeping the sensing configuration deliberately minimal. At two contrasting monitoring sites in Central Thailand (Bangkok urban corridor and Bang Pu coastal mangrove reserve), the joint models stayed on par with their independent single-task baselines, with a mean absolute ΔRMSE of 0.11 μg/m³ over the 24 combinations of model, site, and target, and hierarchy violations below 0.5%.

Parameter sharing cut the total parameter count of the neural architectures by 34.9% for CNN and 64.3% for LSTM relative to independent single-task models. Among the evaluated triple-task configurations, RF-10 delivered the lowest on-device latency (0.065 ms) and energy consumption (0.015 mJ) on the ESP32-S3 microcontroller, while keeping its RMSE within 0.14 μg/m³ of RF-100 across both sites. Substituting one triple-task call for three sequential single-task calls cut inference time and energy 2.2- to 3.0-fold across all evaluated model configurations. Overall, pairing a single PM2.5 sensing channel with an ESP32-S3 microcontroller offers a practical, energy-efficient solution for dense, multi-fraction PM monitoring in developing regions.

ACKNOWLEDGMENT

The author expresses sincere gratitude to the faculty and research staff at the Department of Computer Engineering, Chulalongkorn University, as well as field technical teams who assisted in the continuous operation and maintenance of the ambient air quality monitoring stations in Bangkok and Samut Prakan, Thailand.

REFERENCES

[1]. World Health Organization, "Ambient (outdoor) air pollution," WHO Fact Sheet, Geneva, Switzerland, 2024.
[2]. World Health Organization, "WHO Global Air Quality Guidelines: Particulate Matter, Ozone, Nitrogen Dioxide, Sulfur Dioxide and Carbon Monoxide," World Health Organization, Geneva, Switzerland, 2021.
[3]. C. Terzano, F. Di Stefano, V. Conti, E. Graziani, and A. Petroianni, "Air pollution ultrafine particles: toxicity beyond the lung," Eur. Rev. Med. Pharmacol. Sci., vol. 14, pp. 809–821, 2010.
[4]. D. E. Schraufnagel, "The health effects of ultrafine particles," Exp. Mol. Med., vol. 52, pp. 311–317, 2020.
[5]. Pollution Control Department (PCD), Thailand Ministry of Natural Resources and Environment, "Air Quality and Noise Situation in Thailand," PCD Annual Environmental Report, Bangkok, Thailand, 2024.
[6]. W. Phairuang, M. Hata, and M. Furuuchi, "Influence of agricultural burning and traffic emissions on fine and ultrafine particles in Bangkok, Thailand," Atmos. Environ., vol. 217, p. 116949, 2019.
[7]. S. ChooChuay, S. Pongpiachan, et al., "Long-term monitoring of PM2.5 and associated polycyclic aromatic hydrocarbons in the Bangkok Metropolitan Region," Aerosol Air Qual. Res., vol. 20, no. 6, pp. 1312–1324, 2020.
[8]. P. Srimuruganandam and K. Srichandr, "Machine learning approaches for urban PM2.5 concentration forecasting: A case study of Bangkok Metropolitan Area," Urban Clim., vol. 45, p. 101260, 2022.
[9]. H. Karimian, Q. Li, C. Wu, Y. Qi, Y. Mo, G. Chen, X. Zhang, and S. Sachdeva, "Evaluation of different machine learning approaches to forecasting PM2.5 mass concentrations," Aerosol Air Qual. Res., vol. 19, pp. 1400–1410, 2019.
[10]. F. Xiao, M. Yang, H. Fan, G. Fan, and M. A. A. Al-qaness, "An improved deep learning model for predicting daily PM2.5 concentration," Sci. Rep., vol. 10, p. 20988, 2020.
[11]. X. Li, L. Peng, Y. Hu, J. Shao, and T. Chi, "Deep learning architecture for air quality predictions," Environ. Sci. Pollut. Res., vol. 23, pp. 22408–22417, 2016.
[12]. X. Bai, N. Zhang, X. Cao, and W. Chen, "Prediction of PM2.5 concentration based on a CNN-LSTM neural network algorithm," PeerJ, vol. 12, p. e17811, 2024.
[13]. C. J. Huang and P. H. Kuo, "A deep cnn-lstm model for particulate matter (Pm2.5) forecasting in smart cities," Sensors, vol. 18, p. 2220, 2018.
[14]. S. Zhou, W. Wang, L. Zhu, Q. Qiao, and Y. Kang, "Deep-learning architecture for PM2.5 concentration prediction: A review," Environ. Sci. Ecotechnol., vol. 21, p. 100400, 2024.
[15]. N. Zaini, L. W. Ean, A. N. Ahmed, and M. A. Malek, "A systematic literature review of deep learning neural network for time series air quality forecasting," Environ. Sci. Pollut. Res., vol. 29, 2022.
[16]. T. Li, M. Hua, and X. Wu, "A Hybrid CNN-LSTM Model for Forecasting Particulate Matter (PM2.5)," IEEE Access, vol. 8, pp. 26933–26940, 2020.
[17]. S. Masmoudi, H. Elghazel, D. Taieb, O. Yazar, and A. Kallel, "A machine-learning framework for predicting multiple air pollutants' concentrations via multi-target regression and feature selection," Sci. Total Environ., vol. 715, 2020.
[18]. B. Wang, Z. Yan, J. Lu, G. Zhang, and T. Li, "Deep Multi-task Learning for Air Quality Prediction," in Proc. 25th Int. Conf. Neural Information Processing (ICONIP), 2018, pp. 93–103.
[19]. Z. A. Xie, C. O. Chow, J. H. Chuah, and W. J. K. Raymond, "Multi-pollutant air quality forecasting using bidirectional attention and multi-scale temporal networks," Front. Environ. Sci., vol. 13, 2025.
[20]. J. Song and M. E. J. Stettler, "A novel multi-pollutant space-time learning network for air pollution inference," Sci. Total Environ., vol. 811, 2022.
[21]. Z. Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang, "Edge Intelligence: Paving the Last Mile of Artificial Intelligence with Edge Computing," Proc. IEEE, vol. 107, pp. 1738–1762, 2019.
[22]. G. H. Hong, T. C. Le, J. W. Tu, C. Wang, S. C. Chang, J. Y. Yu, G. Y. Lin, S. G. Aggarwal, and C. J. Tsai, "Long-term evaluation and calibration of three types of low-cost PM2.5 sensors at different air quality monitoring stations," J. Aerosol Sci., vol. 157, 2021.
[23]. F. M. J. Bulot, S. J. Johnston, P. J. Basford, N. H. C. Easton, M. Apetroaie-Cristea, G. L. Foster, A. K. R.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1128

Morris, S. J. Cox, and M. Loxham, "Long-term field comparison of multiple low-cost particulate matter sensors in an outdoor urban environment," Sci. Rep., vol. 9, 2019.

[24]. F. M. J. Bulot, S. J. Ossont, A. K. R. Morris, P. J. Basford, N. H. C. Easton, H. L. Mitchell, G. L. Foster, S. J. Cox, and M. Loxham, "Characterisation and calibration of low-cost PM sensors at high temporal resolution to reference-grade performance," Heliyon, vol. 9, 2023.

[25]. BayesWitnesses, "m2cgen: Machine learning models to native code transpiler," GitHub Repository, 2023. [Online]. Available: https://github.com/BayesWitnesses/m2cgen

[26]. R. David, J. Duke, A. Jain, V. J. Reddi, N. Jeffries, J. Li, N. Kreeger, I. Nappier, M. Natraj, S. Regev, R. Rhodes, T. Wang, and P. Warden, "TensorFlow Lite Micro: Embedded machine learning on TinyML systems," in Proc. 4th MLSys Conf., San Jose, CA, 2021.

[27]. Espressif Systems, "ESP-NN: Optimized neural network functions for Espressif chipsets," GitHub Repository, 2023. [Online]. Available: https://github.com/espressif/esp-nn

[28]. S. Chae, J. Shin, S. Kwon, S. Lee, S. Kang, and D. Lee, "PM10 and PM2.5 real-time prediction models using an interpolated convolutional neural network," Sci. Rep., vol. 11, 2021.

[29]. M. Teng, S. Li, J. Xing, G. Song, J. Yang, J. Dong, X. Zeng, and Y. Qin, "24-Hour prediction of PM2.5 concentrations by combining empirical mode decomposition and bidirectional long short-term memory neural network," Sci. Total Environ., vol. 821, 2022.

[30]. R. Caruana, "Multitask Learning," Mach. Learn., vol. 28, pp. 41–75, 1997.

[31]. D. Xu, Y. Shi, I. W. Tsang, Y.-S. Ong, C. Gong, and X. Shen, "A Survey on Multi-output Learning," IEEE Trans. Neural Netw. Learn. Syst., vol. 31, pp. 2409–2429, 2020.

[32]. V. J. Reddi, B. Plancher, S. Kennedy, L. Moroney, P. Warden, A. Agarwal, et al., "Widening Access to Applied Machine Learning with TinyML," Harv. Data Sci. Rev., vol. 4, 2022.

[33]. P. Warden and D. Situnayake, TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers, 1st ed., O'Reilly Media, Sebastopol, CA, 2019.

[34]. Y. Abadade, A. Temouden, H. Bamoumen, N. Benamar, Y. Chtouki, and A. S. Hafid, "A Comprehensive Survey on TinyML," IEEE Access, vol. 11, pp. 96892–96922, 2023.

[35]. S. S. Saha, S. S. Sandha, and M. Srivastava, "Machine Learning for Microcontroller-Class Hardware: A Review," IEEE Sens. J., vol. 22, pp. 21362–21390, 2022.

[36]. I. N. K. Wardana, S. A. Fahmy, and J. W. Gardner, "TinyML Models for a Low-Cost Air Quality Monitoring Device," IEEE Sens. Lett., vol. 7, 2023.

[37]. L. Fortuna, S. Graziani, A. Rizzo, and M. G. Xibilia, Soft Sensors for Monitoring and Control of Industrial Processes, 1st ed., Springer Science & Business Media, 2007.

[38]. P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven Soft Sensors in the process industry," Comput. Chem. Eng., vol. 33, pp. 795–814, 2009.

[39]. S. De Vito, E. Esposito, M. Salvato, O. Popoola, F. Formisano, R. Jones, and G. Di Francia, "Calibrating chemical multisensory devices for real world applications: An in-depth comparison of quantitative machine learning approaches," Sens. Actuators B Chem., vol. 255, pp. 1191–1210, 2018.

[40]. J. Sun, J. Gong, and J. Zhou, "Estimating hourly PM2.5 concentrations in Beijing with satellite aerosol optical depth and a random forest approach," Sci. Total Environ., vol. 762, 2021.

[41]. S. Uttamang and V. P. Aneja, "Investigation of criteria air pollutants and aerosol characterization in Bangkok, Thailand," Environ. Sci. Atmos., vol. 3, pp. 245–258, 2023.

[42]. N. Thongyen and P. Sompongchaiyakul, "Spatial distribution and health risk assessment of particulate-bound toxic metals in the Bangkok urban atmosphere," Atmos. Pollut. Res., vol. 14, p. 101780, 2023.

[43]. K. Thepnuan, B. Yabueng, et al., "Atmospheric transport and regional dispersion of particulate matter and trace elements along the coastal zones of the Upper Gulf of Thailand," Bull. Environ. Contam. Toxicol., vol. 110, pp. 45–56, 2023.

[44]. C. Bergmeir and J. M. Benítez, "On the use of cross-validation for time series predictor evaluation," Inf. Sci., vol. 191, pp. 192–213, 2012.

[45]. P. Dejchanchaiwong and P. Tekasakul, "Characterization of carbonaceous aerosols and regional smoke haze transport in Central and Southern Thailand," Atmos. Pollut. Res., vol. 15, p. 101950, 2024.

[46]. W. Jinsart and K. Tamura, "Seasonal variation of PM10 and PM2.5 in Bangkok traffic corridors: Chemical composition and diurnal patterns," Chemosphere, vol. 280, p. 130640, 2021.

[47]. S. Wetchayont and S. Pengcheun, "Diurnal and seasonal dynamics of boundary layer meteorology and air quality in the coastal plain of Thailand," Atmos. Environ., vol. 310, p. 120010, 2024.

[48]. A. Kendall, Y. Gal, and R. Cipolla, "Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2018, pp. 7482–7491.

[49]. S. Shanmugavelu, M. Taillefumier, C. Culver, O. Hernandez, M. Coletti, and A. Sedova, "Impacts of floating-point non-associativity on reproducibility for HPC and deep learning applications," in Proc. SC 2024-W: Workshops Int. Conf. High Perform. Comput. Netw. Storage Anal., Atlanta, GA, USA, 2024, pp. 170–179.

[50]. P. Adong, E. Bainomugisha, D. Okure, and R. Sserunjogi, "Applying machine learning for large scale field calibration of low-cost PM2.5 and PM10 air pollution sensors," Appl. AI Lett., vol. 3, 2022.