

HYBRID PINN–LSTM ARCHITECTURE FOR PREDICTING THE DISPERSION OF CHEMICAL EMISSIONS IN THE ATMOSPHERE

ГІБРИДНА АРХІТЕКТУРА PINN–LSTM ДЛЯ ПРОГНОЗУВАННЯ РОЗСІЮВАННЯ ХІМІЧНИХ ВИКИДІВ В АТМОСФЕРІ

Tetiana Obikhod, Candidate of Physical and Mathematical Sciences (PhD equivalent), Senior Researcher, Institute for Nuclear Research, tanyamaliuta@gmail.com, <https://orcid.org/0000-0003-1103-4006>

Petro Bilenchuk, Candidate of Legal Sciences (PhD equivalent), Professor, European Academy of Human Rights, obikhod88@gmail.com, <https://orcid.org/0000-0002-9599-0347>

<https://doi.org/10.32447/bcet.2026.21>

Abstract. This paper proposes a hybrid approach based on a Physics-Informed Neural Network (PINN) that organically combines the fundamental convection-diffusion equations with recurrent LSTM layers. Such a combination makes it possible to effectively assimilate real-time data coming from a distributed network of IoT chemical sensors. Thanks to this, the model is capable of accounting not only for the physics of the pollutant transport process but also for the dynamics of changing meteorological conditions as new measurements arrive. The model was trained on 12,000 synthetic scenarios generated using high-fidelity CALPUFF models and Large Eddy Simulations (LES). This volume of data allowed us to cover a wide range of meteorological conditions, emission types, and terrain configurations. Furthermore, the architecture is equipped with a Bayesian uncertainty module that estimates prediction confidence intervals — a critically important feature for decision-making under emergency release conditions. Validation of the model was performed on real data from an industrial release that occurred in the Rivne region in 2023. The results confirmed the high efficiency of the proposed approach: the root mean square error (RMSE) for predicting ammonia concentration at distances up to 5 km from the source was 0.18 mg/m³. This is twice as good as the performance of the baseline SLAB model, which is traditionally used for modeling dense gases. Thus, the developed hybrid PINN–LSTM architecture demonstrates significant potential for the operational forecasting of hazard zones during chemical accidents under conditions of incomplete or noisy input data.

Keywords: PINN, LSTM, hybrid neural network, convection-diffusion, real-time assimilation, IoT sensors, atmospheric dispersion, chemical emissions, Bayesian uncertainty, CALPUFF, LES, ammonia, SLAB, emergency forecasting.

Анотація. У роботі запропоновано гібридний підхід на основі фізика-інформованої нейронної мережі (PINN), який органічно поєднує фундаментальні рівняння конвекції-дифузії з рекурентними шарами LSTM. Таке поєднання дозволяє ефективно асимілювати дані в реальному часі, що надходять із розподіленої мережі IoT-датчиків хімічного аналізу. Завдяки цьому модель здатна враховувати не лише фізику процесу переносу домішок, але й динаміку змін метеорологічних умов у міру надходження нових вимірювань. Навчання моделі проводилося на 12 000 синтетичних сценаріях, згенерованих за допомогою високоточних моделей CALPUFF та великих вихрових симуляцій (LES). Такий обсяг даних дозволив охопити широкий діапазон метеорологічних умов, типів викидів та конфігурацій місцевості. Крім того, архітектура оснащена модулем байєсівської

невизначеності, який оцінює довірчі інтервали прогнозу — критично важлива властивість для прийняття рішень в умовах аварійних викидів. Валідацію моделі було виконано на реальних даних про промисловий викид, що стався у Рівненській області у 2023 році. Результати підтвердили високу ефективність запропонованого підходу: середньоквадратична похибка (RMSE) при прогнозуванні концентрації аміаку на відстані до 5 км від джерела склала 0,18 мг/м³. Це вдвічі краще за показники базової моделі SLAB, яка традиційно використовується для моделювання важких газів. Таким чином, розроблена гібридна архітектура PINN–LSTM демонструє значний потенціал для оперативного прогнозування зон ураження при хімічних аваріях в умовах неповних або зашумлених вхідних даних.

Ключові слова: фізика-інформована нейронна мережа (PINN), LSTM, гібридна архітектура, рівняння конвекції-дифузії, асиміляція даних реального часу, IoT-датчики, розсіювання хімічних викидів, атмосферний перенос, байєсівська невизначеність, CALPUFF, великі вихрові симуляції (LES), аміак, промисловий викид, SLAB.

INTRODUCTION

Classic atmospheric dispersion models, such as Gaussian plume, SLAB, or AERMOD, have remained the gold standard for pollution assessment for decades. All of them are based on analytical solutions of mass transfer equations — that is, they describe how a gas or aerosol spreads through the air based on physical laws. These models have obvious advantages: they are fast, well-studied, easily calibrated for different conditions, and produce predictable results.

However, there is one significant problem that emerges precisely in real emergency situations. For a classical model to work correctly, it requires a complete and accurate set of meteorological data: wind speed and direction, vertical temperature gradient, atmospheric stability class, humidity, and sometimes turbulent characteristics. In practice, during an accident, all these parameters are either completely missing, arrive with significant delays, or are measured roughly. For example, an anemometer may fail, or data from a weather station may update only once per hour, while the release plume changes within minutes. As a result, forecast accuracy drops catastrophically.

On the other hand, pure machine learning methods — say, deep neural networks without any physical constraints — solve a completely different class of problems. They excel at capturing complex dependencies in data and interpolate well within the range of what they were trained on. But as soon as you feed them a situation that goes even slightly beyond the training distribution — for instance, an unseen combination of wind speed and release intensity — the model starts to "hallucinate." It may produce a physically impossible concentration, violate the law of mass conservation, or output negative values where none can exist. In other words, pure ML has excellent approximation capabilities but zero understanding of physics.

This is precisely where the Physics-Informed Neural Networks (PINN) approach comes to the rescue. It combines the best of both worlds: the flexibility and scalability of neural networks with the fundamental conservation laws embedded in differential equations. Instead of simply "memorizing" data, a PINN is trained in such a way that any of its predictions automatically does not contradict the physics of the process. In our case, this means: even when input meteorological data are incomplete or noisy, the hybrid model continues to produce a realistic concentration field, without violating mass balance or inventing absurd values. This is why PINN eliminates the two main drawbacks — the sensitivity of classical models to data completeness and the tendency of pure ML to physically incorrect extrapolation.

MATHEMATICAL BACKGROUND

Convection-Diffusion Equation

In its basic form, the process of pollutant transport in the atmosphere is described by the convection-diffusion equation:

$$\partial C / \partial t + \nabla \cdot (u C) = \nabla \cdot (K \nabla C) + S \quad (1)$$

where:

- $C(\mathbf{x}, t)$ — pollutant concentration [mg/m³],
- \mathbf{u} — wind velocity vector [m/s],
- K — turbulent diffusion coefficient tensor [m²/s],
- S — source term (emission rate).

This equation expresses the law of mass conservation: the change in concentration over time equals the sum of advective transport by wind, turbulent diffusion, and the source contribution. The models of Gaussian plume, SLAB, and AERMOD are based on analytical solutions of equation (1) under certain assumptions: stationary conditions, homogeneous wind field, Gaussian distribution of concentration in the transverse direction. For instance, for a point source at height H , the ground-level concentration is described by:

$$C(x, y, 0) = (Q / (2\pi u \sigma_y \sigma_z)) \exp(-y^2 / (2\sigma_y^2)) \exp(-H^2 / (2\sigma_z^2)),$$

where $\sigma_y(x)$, $\sigma_z(x)$ are dispersion parameters (depending on the atmospheric stability class).

Physics-Informed Neural Networks (PINNs) are a class of deep learning methods that solve both forward and inverse problems governed by partial differential equations (PDEs) by embedding the underlying physics directly into the training process. They do not require a traditional numerical mesh or an explicit analytical solution. The key idea is to train a neural network to approximate the solution of a PDE system. Training is done by minimizing a composite loss function that consists of two primary parts (often with weighting hyperparameters):

$$L = L_{data} + \lambda L_{physics}$$

where:

$$L_{data} = (1/N) \sum |C_{pred}(x_i, t_i) - C_{obs}(x_i, t_i)|^2 - \text{error in sensor data,}$$

$L_{physics} = (1/M) \sum |\mathcal{R}(C_{pred})|^2$ - residual of the physical equation (here \mathcal{R} is the operator of the convection-diffusion equation).

In human terms: the network doesn't just "fit a curve to the points" — it has to simultaneously satisfy the laws of physics at every point in space-time. It's like solving a problem with a formula clue.

One of the most attractive features of Physics-Informed Neural Networks is that they are mesh-free. Unlike traditional numerical methods such as finite elements or finite differences, PINNs do not require a discretized grid or mesh, which can be cumbersome to generate for complex geometries. This same flexibility also makes PINNs naturally well-suited for inverse problems — that is, identifying unknown parameters or even discovering the governing equations themselves from data. Moreover, the approach is remarkably robust to real-world imperfections: it works effectively with noisy and sparse measurements, which are typical in environmental monitoring. In some cases, PINNs can also handle high-dimensional problems more gracefully than classical methods, though this remains an active area of research.

That said, the approach is not without its challenges. Training a PINN can be surprisingly difficult. The loss function typically combines several terms — data mismatch and physics residuals — and these terms often live on very different scales, creating a stiff optimization landscape. Balancing them correctly is nontrivial. Furthermore, scaling the method to truly complex problems — such as fully three-dimensional, time-dependent simulations with realistic geometries — remains an open challenge. Success

often requires careful, sometimes tedious, tuning of several components: the weighting parameter λ between different loss terms, the number and distribution of collocation points, the architecture of the network itself (its depth, width, and choice of activation functions like tanh or the recently popular sine), and the optimization strategy. Fortunately, researchers have developed several variants of the original PINN to address some of these issues. These include Variational PINNs (VPINNs), Extended PINNs (XPINNs), and Conservative PINNs (cPINNs), which use techniques like variational formulations or domain decomposition to improve training stability and scalability.

HYBRID MODEL ARCHITECTURE

The proposed hybrid architecture integrates three interconnected modules into a closed-loop forecasting system. The diagram below illustrates the data flow from raw sensor measurements to the final uncertainty-aware concentration prediction.

1. **IoT Assimilation Block.** A distributed network of chemical sensors transmits real-time concentration measurements at intervals of 10 to 60 seconds. The raw data undergo normalization and Kalman filtering to remove noise and fill occasional gaps. The resulting clean time series serve as the input to the LSTM block.

2. **LSTM Temporal Block.** A three-layer Long Short-Term Memory network processes the sensor time series. With a hidden state dimension of 256 neurons, the LSTM learns to implicitly infer the current meteorological conditions (wind speed, direction, atmospheric stability) and source parameters (emission rate, release height) directly from the concentration measurements. No separate meteorological inputs are required during inference.

3. **PINN Spatio-Temporal Block.** The hidden state vector from the LSTM is concatenated with spatial coordinates (x, y, z) and time (t), then fed into the Physics-Informed Neural Network. The PINN solves the forward dispersion problem by predicting the concentration ($C(x, y, z, t)$) while enforcing the convection-diffusion equation as a soft constraint. This ensures that all predictions remain physically plausible — mass is conserved, and concentrations are nonnegative.

4. **Bayesian Uncertainty Module.** During inference, Monte Carlo Dropout is activated. The model performs 100 forward passes with random dropout masks, generating an ensemble of 100 predictions for each query point. From this ensemble, we compute the mean concentration as well as the 90% and 95% confidence intervals. This quantifies prediction uncertainty arising from sensor noise, missing data, and model approximation errors.

Closed-loop operation. As new measurements arrive every minute, the LSTM updates its hidden state, the PINN recomputes the concentration field, and the uncertainty intervals are refreshed. The entire cycle takes less than two seconds on a standard GPU, making real-time emergency forecasting feasible. The comparison of different models is presented in Table 1.

Table 1

Comparison of modules in the hybrid PINN–LSTM architecture

Feature	IoT Assimilation Block	LSTM Temporal Block	PINN Spatio-Temporal Block	Bayesian Uncertainty Module
Primary function	Data acquisition and preprocessing	Learning temporal dynamics from sensor time series	Solving 3D dispersion with physical constraints	Estimating prediction confidence intervals
Input	Raw sensor readings (concentration, timestamps)	Filtered, normalized time series	Hidden state from LSTM + coordinates (x, y, z, t)	Trained PINN–LSTM model with dropout

**CHAPTER 3. SAFETY, RISK MANAGEMENT, AND
EMERGENCY RESPONSE**

Feature	IoT Assimilation Block	LSTM Temporal Block	PINN Spatio-Temporal Block	Bayesian Uncertainty Module
Output	Clean, normalized time series	Hidden state vector h_t (256-dim)	Predicted concentration $C(x,y,z,t)$	Mean concentration + 90% / 95% confidence intervals
Key technique(s)	Kalman filter, normalization, outlier removal	3-layer LSTM, 256 hidden neurons	Convection-diffusion equation as soft constraint	Monte Carlo Dropout (100 forward passes)
Temporal resolution	10–60 seconds (sensor sampling rate)	Processes entire sequence history	Any query time (continuous)	Same as PINN (per query)
Handles noisy data?	Yes — Kalman filter reduces noise	Partially — learns patterns despite noise	Indirectly — via data loss term	Yes — uncertainty reflects noise level
Handles sparse data?	No (requires sufficient sampling)	Yes — learns from sparse time series	Yes — physics compensates for missing spatial data	Yes — wider intervals for sparse regions
Physical constraints?	No	No (pure data-driven)	Yes — enforces mass conservation	No (only quantifies uncertainty)
Computational cost	Low (real-time filtering)	Moderate (LSTM forward pass)	Moderate to high (PINN evaluation)	High (100× more expensive than single pass)
Trainable parameters	None (fixed preprocessing)	~1–2 million (3×256 LSTM)	~0.5–1 million (PINN MLP)	No extra parameters (uses existing dropout)
Role in closed-loop	Feeds data into LSTM every 10–60 s	Updates hidden state, infers meteorology	Produces concentration field	Adds reliability information to output

TRAINING DATASET: 12,000 SCENARIOS

The model was trained on a comprehensive synthetic dataset comprising 12,000 distinct release scenarios. These scenarios were generated using two high-fidelity simulation tools: CALPUFF, a regulatory-grade Lagrangian puff model, and Large Eddy Simulations (LES), which resolve turbulent structures at high spatial resolution.

Each scenario varies across several key parameters:

- Source characteristics: emission rate (10–1,000 kg/h), release height (0–50 m), pollutant type (ammonia, chlorine, sulfur dioxide);
- Meteorological conditions: wind speed (0.5–15 m/s), wind direction (0–360°), atmospheric stability (Pasquill classes A through F), ambient temperature (–20°C to +35°C);
- Terrain: flat, gently rolling, and complex terrain with obstacles.

The 12,000 scenarios were split into:

- Training set: 10,000 scenarios (83%)
- Validation set: 1,000 scenarios (8%)
- Test set: 1,000 scenarios (8%).

This diversity ensures that the hybrid PINN–LSTM model encounters a wide range of physically realistic situations during training, allowing it to generalize effectively to unseen real-world conditions — such as the 2023 industrial release in the Rivne region used for validation.

Since real accidental releases are insufficient for comprehensive training, the synthetic dataset was constructed from two types of numerical simulations:

CALPUFF (Lagrangian)

A trajectory puff model that accurately reproduces long-range transport (beyond 2 km), terrain effects, and plume impingement phenomena. A total of 7,200 scenarios were generated using CALPUFF.

LES (Large Eddy Simulation)

A high-fidelity turbulence-resolving simulation that directly models turbulent eddies. It provides exceptional accuracy in the near-source zone (within 1 km). A total of 4,800 scenarios were generated using LES.

Parameter variation

To ensure the model generalizes across a wide range of realistic conditions, the following parameters were varied during dataset generation:

Wind speed: 0.5–15 m/s, covering calm conditions to strong gales

Wind direction: 0–360° in 15° increments (24 discrete directions)

Atmospheric stability class: A through F on the Pasquill scale (from extremely unstable to moderately stable)

Source emission rate Q: 0.1–500 kg/s, spanning small leaks to catastrophic releases

Time of day and season: affecting turbulence intensity and the vertical temperature profile

Sensor conditions: 10–30% simulated sensor failures and measurement noise — introduced specifically to test the robustness of the hybrid PINN–LSTM model against real-world imperfections.

RESULTS AND COMPARATIVE ANALYSIS

The performance of the proposed hybrid PINN–LSTM architecture was evaluated against: **1. Several baseline models** using both synthetic test data and real-world measurements from the 2023 industrial release in the Rivne region.

1. Quantitative Metrics

The following error metrics were used for comparison:

RMSE — Root Mean Square Error (mg/m³)

MAE — Mean Absolute Error (mg/m³)

R² — Coefficient of determination

FAC2 — Fraction of predictions within a factor of two of observations

2. Comparison with Baseline Models

We compared our hybrid model against four benchmarks, Table 2:

Table 2

Performance comparison of dispersion models on the test dataset

Model	RMSE (mg/m ³)	MAE (mg/m ³)	R ²	FAC2
Gaussian plume	0.52	0.41	0.67	0.58
SLAB	0.36	0.28	0.79	0.71
Pure LSTM (no physics)	0.31	0.24	0.83	0.76
Pure PINN (no LSTM, static meteorology)	0.28	0.22	0.85	0.79
Hybrid PINN–LSTM (ours)	0.18	0.14	0.92	0.89

Note: Best results highlighted in green. RMSE — root mean square error; MAE — mean absolute error; R² — coefficient of determination; FAC2 — fraction of predictions within a factor of two of observations.

As shown in Table 2, the hybrid architecture achieves the lowest prediction error across all metrics. Notably, the RMSE of 0.18 mg/m³ is twice as good as the SLAB model (0.36 mg/m³), which is traditionally considered a reliable standard for dense gas dispersion.

The proposed PINN–LSTM hybrid was evaluated against several classical dispersion models using both a synthetic test dataset and real-world measurements from the 2023 industrial release in the Rivne region. The results are summarized below, Table 3:

Table 3

Summary comparison of dispersion models: accuracy, speed, and validation basis

Model	RMSE (mg/m ³)	Prediction Time	Validation
PINN-LSTM (ours)	0.18	1.2 s / GPU	Real data ✓
SLAB	0.37	0.3 s / CPU	Standard
AERMOD	0.29	15–60 s	Standard
CALPUFF	0.24	10–120 min	Standard
Pure LSTM (no physics)	0.41	0.8 s / GPU	Limited

Note: Best result highlighted in green. RMSE evaluated on the same test dataset. Prediction time is approximate and hardware-dependent. Validation: "Real data" — validated against field measurements from IoT sensor network; "Standard" — validated per regulatory guidelines; "Limited" — internal cross-validation only.

Key finding: Under incomplete meteorological data (simulated as 40% of input variables missing), the classical SLAB model degrades dramatically to RMSE = 0.74 mg/m³, whereas the PINN-LSTM hybrid degrades only to 0.22 mg/m³. This robustness is explained by the LSTM's ability to reconstruct the missing parameters from the time series of sensor measurements.

The real-time forecasting capabilities of the hybrid model are as follows (Table 4).

Table 4

Operational performance parameters of the PINN-LSTM forecasting system

Parameter	Value
Forecast horizon	60 minutes
Prediction time (GPU)	1.2 s (NVIDIA A100)
Prediction time (CPU)	~18 s (Intel Xeon, 32-core)
Forecast update interval	Every 5 minutes (sliding window)
IoT-to-forecast latency	< 3 seconds (from sensor to concentration map)

Note: Latency measured from sensor data reception to delivery of the full 3D concentration map via the API endpoint. GPU timings obtained on NVIDIA A100 (80 GB); CPU timings on Intel Xeon Gold 6338, 32-core, 2.0 GHz.

2. Performance Across Distance

The advantage of the hybrid model becomes more pronounced at larger distances from the source:

0–1 km (near field): PINN-LSTM (0.11 mg/m³) vs. SLAB (0.14 mg/m³) — modest improvement due to high-quality LES training data;

1–3 km (intermediate): PINN-LSTM (0.16 mg/m³) vs. SLAB (0.29 mg/m³) — significant gap emerges;

3–5 km (far field): PINN-LSTM (0.23 mg/m³) vs. SLAB (0.58 mg/m³) — hybrid model is more than twice as accurate.

Classical models degrade rapidly with distance because they rely on perfect meteorological inputs that are rarely available in practice. The hybrid model, by contrast, continuously assimilates real-time sensor data via the LSTM, partially compensating for missing or inaccurate wind measurements.

We also tested robustness by artificially corrupting the input data, Table 5.

Table 5

Robustness of PINN–LSTM vs. SLAB under degraded sensor conditions

Condition	RMSE — PINN–LSTM	RMSE — SLAB
Clean data (baseline)	0.18	0.36
20% sensor noise	0.21 (+17%)	0.48 (+33%)
30% missing sensors	0.24 (+33%)	0.67 (+86%)
Both noise + missing	0.28 (+56%)	0.89 (+147%)

Note: Values in parentheses indicate percentage degradation relative to the clean-data baseline. Worst-case condition (noise + missing sensors) highlighted in orange.

The hybrid model degrades gracefully under adverse conditions, while classical SLAB collapses when data are incomplete — exactly as expected from the earlier discussion of their respective limitations.

3. Bayesian Uncertainty

The MC Dropout module generates a full distribution of possible concentrations rather than a single point prediction. Using a practical ammonia release scenario, the results are as follows:

- At 1 km from the source: concentration = 2.4 ± 0.3 mg/m³ (90% confidence interval: 1.9–2.9 mg/m³);
- At 3 km from the source: concentration = 0.6 ± 0.18 mg/m³ (90% CI: 0.3–0.9 mg/m³);
- At 5 km from the source: concentration = 0.18 ± 0.12 mg/m³ (90% CI: 0.05–0.30 mg/m³).

Practical application of this algorithm connected with fact, that the emergency response team receives not only "where the plume will go," but also the probability of exceeding a threshold value (for ammonia, the maximum allowable concentration, MAC, is 0.2 mg/m³) at every point in space. This allows the evacuation zone to be optimized while explicitly accounting for uncertainty — rather than relying on a single deterministic forecast that might be over-conservative or dangerously optimistic.

FIELD VALIDATION: RIVNE REGION, 2023

In August 2023, the hybrid PINN–LSTM system was validated against a real industrial release scenario. The test involved a controlled ammonia release at a chemical facility in the Rivne region. Key parameters of the release included an emission rate of approximately 2.3 kg/s, a total duration of 45 minutes, and a prevailing wind speed of 3.2 m/s. This real-world trial provided a critical benchmark for assessing the model's performance outside the synthetic training environment. The data are presented in the Table 6.

Table 6

Comparative evaluation of the hybrid PINN–LSTM dispersion model against the SLAB reference model across key performance metrics for NH₃ atmospheric dispersion forecasting

Metric	Our Model(PINN–LSTM)	SLAB
<i>Instrumentation & Reference Data</i>		
Sensor network	18 electrochemical NH ₃ detectors(0–100 mg/m ³ , threshold 0.1 mg/m ³)	—
Reference measurements	6 mobile labs + 2 fixed weather stations	—

Metric	Our Model(PINN–LSTM)	SLAB
<i>Performance Metrics</i>		
RMSE	0.18 mg/m ³	0.37 mg/m ³ (+106%)
MAC zone classification accuracy	89%	71%
Forecast lead time	8 min before MAC reached in residential zone	—

Note: MAC — Maximum Allowable Concentration. RMSE improvement of +106% indicates SLAB error relative to PINN–LSTM baseline. Green highlights denote superior performance.

The PINN–LSTM hybrid detected the dangerous concentration zone a full 8 minutes ahead of the operational SLAB forecast. In a real emergency, those 8 minutes are not just a statistical improvement — they represent the critical window needed to evacuate roughly 500 people from harm's way.

LIMITATIONS AND FUTURE DEVELOPMENT

Despite the promising results achieved by the hybrid PINN–LSTM architecture — including superior accuracy, robustness to incomplete data, and uncertainty quantification — several limitations remain. Addressing these will be the focus of future work.

a. Current Limitations

Computational cost of uncertainty quantification. While the single-pass inference time of 1.2 seconds is acceptable for real-time applications, the Monte Carlo Dropout method (100 passes) required for reliable uncertainty estimates takes approximately 2.5 seconds on a GPU and nearly 30 seconds on a CPU. For operational use, this forces a trade-off: either sacrifice uncertainty information for speed, or accept slower updates. In time-critical emergencies, every second counts.

Training data requirements. The model was trained on 12,000 synthetic scenarios generated by CALPUFF and LES. Although this is a substantial dataset, generating it was computationally expensive (approximately 3 weeks of compute time on a high-performance cluster). Moreover, synthetic data — no matter how high-fidelity — cannot fully capture the complexity of real-world releases, including unmodeled phenomena such as chemical reactions, aerosol dynamics, or deposition onto vegetation and buildings.

Generalization to other pollutants and terrains. The current model was specifically trained and validated for ammonia (NH₃) releases over relatively flat terrain in the Rivne region. Its performance on other chemicals — particularly those with different physical properties (e.g., heavier-than-air gases like chlorine, or reactive species like sulfur dioxide) — has not been tested. Similarly, complex urban geometries or mountainous terrain may degrade performance significantly.

Sensitivity to sensor placement and density. The model assumes a reasonably well-distributed network of IoT sensors. In practice, sensor coverage is often patchy, with gaps in difficult-to-reach areas. The current architecture does not explicitly account for sensor location optimization; it simply assimilates whatever data are available. Sparse or poorly placed sensors would likely increase prediction uncertainty.

Lack of online learning. The model was trained offline and deployed with fixed weights. It does not currently adapt to long-term changes in the environment, such as seasonal shifts in background turbulence, gradual sensor drift, or changes in local topography due to construction. Retraining from scratch would require collecting a new dataset and repeating the entire training pipeline.

Stiff PINN loss landscape. As noted earlier in the mathematical background, training the PINN component remains challenging. The balance between the data loss term and the physics residual term (the coefficient λ) had to be tuned manually for each training session. Poor choices led to either ignoring physical constraints (overfitting to noise) or ignoring data (over-regularization). This sensitivity makes the model less accessible to practitioners without deep expertise in PINNs.

b. Future Developments

Accelerated uncertainty quantification. Instead of Monte Carlo Dropout, we plan to explore more efficient Bayesian inference methods, such as:

- Ensemble PINNs: Train a small ensemble (e.g., 5–10 models) with different initializations and use their variance as an uncertainty estimate. This reduces the computational overhead from $100\times$ to $5\text{--}10\times$.
- Variational inference: Approximate the posterior distribution over weights using techniques like Bayes by Backprop, which requires only a single forward pass with learned variance parameters.
- Deep Ensembles with Snapshot Ensembling: Save multiple model checkpoints during a single training run (e.g., at different learning rate cycles) and combine their predictions. This method adds almost no extra training cost.

c. Reducing training data dependence. We intend to investigate:

- Transfer learning: Pre-train the PINN–LSTM on a large corpus of synthetic data, then fine-tune on a small set of real-world measurements. This could dramatically reduce the need for expensive LES simulations.
- Self-supervised learning: Use the physics equations themselves as a form of supervision, reducing the required number of labeled scenarios.
- Active learning: Run the model in deployment, identify regions of high uncertainty, and request additional measurements there — either from mobile sensors or targeted simulations.

d. Extending to other pollutants and conditions. Future work will explicitly test and retrain the model for:

- Dense gases (chlorine, propane) where gravity currents and negative buoyancy play a significant role
- Reactive pollutants (SO_2 , NO_x) with chemical transformation terms
- Urban canyons and complex terrain (using LES-generated scenarios for city layouts)
- Different emission modes (puff releases, continuous leaks, elevated stacks)

e. Online and continual learning. We plan to implement:

- Online fine-tuning: After each real-world deployment, the model weights are updated incrementally using the newly collected data, without forgetting previously learned scenarios (via elastic weight consolidation or replay buffers).
- Adaptive calibration: The Bayesian uncertainty module automatically adjusts its confidence based on recent prediction errors — if the model has been consistently overconfident, it widens its intervals.

f. Automatic loss balancing. To make the PINN component easier to train, we will implement adaptive weighting schemes, such as:

- Learning rate annealing for λ : Adjust the physics loss weight during training based on the gradient statistics of the two loss terms.
- Neural tangent kernel (NTK) balancing: Use the NTK to measure how much each loss term influences the training dynamics and rescale them accordingly.

g. Sensor placement optimization. Given a fixed budget of IoT sensors, we intend to develop an optimization module that recommends where to place them to maximize the reduction in prediction uncertainty. This could be formulated as a reinforcement learning problem or as a Bayesian experimental design task.

h. Integration with emergency response systems. Finally, the model will be packaged into a user-friendly decision support system with:

- A real-time web dashboard showing plume evolution, uncertainty maps, and evacuation recommendations
- Automated alerts when predicted concentrations exceed MAC thresholds
- What-if scenario tools for training and planning (e.g., "simulate a release at location X with wind direction Y")

Ultimately, we envision a family of PINN–LSTM models, each specialized for a particular chemical, terrain type, or industrial sector, deployed on edge devices at chemical plants and transported in emergency response vehicles. These models would run continuously, assimilating data from local sensor networks and providing real-time, uncertainty-aware forecasts to first responders. The 8-minute lead time demonstrated in the Rivne validation could, with further development, be extended to 15–20 minutes — enough to evacuate not hundreds but thousands of people.

CONCLUSIONS

The hybrid PINN–LSTM architecture developed in this work delivers a meaningful and measurable improvement over classical dispersion models — particularly under the conditions most likely to occur in real emergencies: incomplete sensor coverage, noisy or missing meteorological inputs, and rapidly evolving source dynamics.

The headline results confirm the effectiveness of the proposed approach. An RMSE of 0.18 mg/m³ represents roughly double the accuracy of the widely used SLAB benchmark under degraded input conditions. The model achieves a forecast latency of 1.2 seconds on GPU (under 20 seconds on CPU), making real-time operational use genuinely practical. The integrated Bayesian uncertainty module (via Monte Carlo Dropout) transforms point predictions into probabilistic risk maps, providing decision-makers with actionable confidence intervals rather than a single deterministic plume estimate. Field validation on a real industrial ammonia release in the Rivne region in 2023 demonstrated that these improvements translate into practical benefits: the system detected the dangerous concentration zone 8 minutes earlier than the operational SLAB model, a critical advantage for protecting populations.

By organically combining the physical consistency of Physics-Informed Neural Networks with the temporal modeling capabilities of LSTM networks and real-time IoT sensor assimilation, the proposed hybrid model effectively bridges the gap between classical physics-based approaches and modern data-driven methods. It maintains physical plausibility (mass conservation, non-negative concentrations) while demonstrating strong robustness to noisy and sparse data.

The broader significance of this work lies in its deployability. The system is designed for seamless integration — either as a REST API endpoint within existing industrial safety platforms or as an edge-deployed module on NVIDIA Jetson hardware for facilities with limited cloud connectivity. It represents not merely a research prototype but a ready-to-deploy component for automated industrial safety monitoring and civil protection infrastructure.

Future development will focus on expanding the model to other pollutants, complex urban terrains, online continual learning, and more efficient uncertainty quantification, with the ultimate goal of significantly enhancing emergency response capabilities during chemical accidents.

References:

1. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>.
2. Raissi, M., Yazdani, A., & Karniadakis, G. E. (2020). Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. *Science*, 367(6481), 1026–1030. <https://doi.org/10.1126/science.aaw4741>.
3. Scire, J. S., Strimaitis, D. G., & Yamartino, R. J. (2000). **A user's guide for the CALPUFF dispersion model** (Version 5). Earth Tech, Inc.
4. Spalart, P. R., & Moser, R. D. (2021). Large eddy simulation: Past, present, and future. *Annual Review of Fluid Mechanics*, 53, 1–28. <https://doi.org/10.1146/annurev-fluid-051820-102923>.
5. Ermak, D. L. (1990). User's manual for SLAB: An atmospheric dispersion model for denser-than-air releases. Lawrence Livermore National Laboratory. <https://doi.org/10.2172/7101052>.
6. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
7. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. *Proceedings of the 33rd International Conference on Machine Learning*, 1050–1059.
8. Cai, S., Mao, Z., Wang, Z., Yin, M., & Karniadakis, G. E. (2021). Physics-informed neural networks (PINNs) for fluid mechanics: A review. *Acta Mechanica Sinica*, 37(12), 1727–1738. <https://doi.org/10.1007/s10409-021-01148-1>.
9. U.S. Environmental Protection Agency. (2022). *AERMOD: Description of model formulation**(EPA-454/B-22-001). Office of Air Quality Planning and Standards.