

KAN-Mamba with Physics-Informed Learning for Landfill Methane Leak Localization

Junsuke Moriya^{1,*}, Kyohei Seki¹ and Mikiro Tokuda¹

¹ Department of Computer Science, Keio University, Yokohama, 223-8522, Japan

*Corresponding author: junsuke.mo@keio.jp

Abstract. Landfill methane leakage is spatially sporadic, meteorologically sensitive, and difficult to pinpoint from sparse concentration measurements. This paper introduces a physics-informed KAN-Mamba model for high-resolution leak localization by combining Kolmogorov-Arnold Networks (KANs) to handle nonlinear source-to-receptor mapping, a selective state-space Mamba backbone for long-range temporal dependency, and differentiable plume constraints that penalize mass-balance and advection-diffusion inconsistencies. A simulated field campaign with 64 fixed sensors, two mobile transects, 46 leak events and 18 meteorological regimes was conducted to test the model against LSTM, Transformer, PINN and wind-conditioned graph neural network baselines. The median localization error of the proposed model was 6.8 m, the source-rate mean absolute error decreased to 0.031 kg h⁻¹, and the event-level F1-score reached 0.941 under mixed wind stability; it exceeded the best baseline by 21.6% in distance error and 8.3% in recall. Based on the ablation experiment, the physics loss reduced false alarms near the downwind boundary by 34.7%, and the KAN layer improved the expected calibration error under weak wind from 0.113 to 0.061. Robustness tests further showed a recall of over 0.89 at 35% receptor dropout and only 2.6 m additional error under +/-10-degree wind perturbation. These results indicate that physically constrained sequence learning can convert sparse methane data into high-quality source maps for engineering inspection and landfill emission control.

Keywords: *Physics-Informed Learning, Landfill Methane, Leak Localization, KAN-Mamba, Plume Constraint*

Received on 12 September 2023, Accepted on 23 January 2024, Published on 02 February 2024

Copyright © 2024 Author(s), licensed to DEA. This is an open access article distributed under the terms of the CC BY-NC-SA 4.0, which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

Introduction

Methane emissions from municipal solid-waste landfills are still a serious engineering problem because they have high climate relevance and unevenly distributed spatial releases [1]. Recent emission inventories show that landfill gas can account for a large proportion of anthropogenic methane in urban areas when cover integrity, leachate routing and gas-collection pressure are not jointly controlled [2]. Leakage is rarely a single, stable-point source; instead, it often takes the form of short-lived hot spots along cover seams, penetrations, daily-cover transitions, wellheads and drainage interfaces [3]. Fixed monitoring stations can detect site-level abnormalities, but the recorded data are highly susceptible to wind shear, atmospheric stability, local roughness and sensor position [4]. Bel Hadj Ali et al. [5] compared surface emission monitoring, drone emission monitoring and downwind plume emission monitoring, showing that mobile and drone-based measurements can add spatial detail after transport has reshaped the source signature. Engineering decisions need to determine whether methane has increased, and if so, where the repair crew should dig, patch, rebalance extraction or inspect buried facilities. Therefore, a localization model needs to recover the source field from incomplete downwind evidence while adhering to the physical behavior of a buoyant, advected and diffused gas plume [6].

Data-driven models have improved methane anomaly recognition, but many of them treat concentration time series as general signals and do not consider the transport mechanism. Recurrent networks have been employed to capture lagged concentration responses at static receptors, but their memory may be lost under abrupt wind

rotation [7]. Attention-based models can relate distant observations in space and time, but their computational cost is relatively high when processing dense sensor histories and mobile transects together [8]. Graph neural networks can naturally encode the monitoring topology, and the graph edges are generally only a proxy for physical connection rather than a direct plume constraint [9]. Physics-informed neural networks have incorporated governing-equation residuals into environmental modelling, but many implementations have been optimized for smooth fields and are thus sensitive to sparse, noisy landfill data [10]. Recently, state-space sequence models have offered a good compromise by handling long-term temporal context with linear scaling and learning selective memory transitions that can adapt to nonstationary forcing [11].

The current work addresses the deficiency of flexible learning systems for physically grounded leak localization. The proposed KAN-Mamba model employs a Mamba backbone for long-receptor-history representation, a Kolmogorov-Arnold Network module to learn compact nonlinear mappings of meteorology, concentration residuals and source coordinates, and a physics-informed objective function to steer predictions towards transport-consistent methane fields. Rather than functioning as a purely empirical predictor, the model learns a leak probability map and source-rate estimate under relaxed advection-diffusion and mass-conservation constraints over the landfill surface. The design of this study is an engineering evaluation rather than a theoretical architecture note; therefore, the evaluation emphasizes localization accuracy, source-rate estimation, robustness to sensor dropout, and behavior under wind conditions that occur frequently in landfill inspections.

To make the argument traceable from monitoring constraints to model design and engineering evidence, the rest of this manuscript will be organized as follows. Section 2 introduces methane-monitoring localization methods and recent learning architectures for sparse environmental sensing. Section 3 introduces the proposed KAN-Mamba framework and defines the data representation, model architecture and physics-informed training objective. Section 4 shows the experimental data set, benchmark design, quantitative comparison, ablation analysis, uncertainty behavior and engineering interpretation of the data figures and tables. Section 5 summarizes the technical contributions, limitations and future research directions of the paper.

Related Work

Landfill Methane Monitoring and Source Localization

In the past, the locations of landfill methane have been determined by emission screening, plume analysis and specific-area investigations. Static sensors can continuously monitor an abnormal concentration, but when a small number of receptors need to cover a large, heterogeneous landfill surface, the inverse problem is ill-posed [12]. Mobile inspection systems expand spatial coverage by sampling along roads or covering grids, and their value depends on when the transects are taken in relation to the wind direction and source activity [13]. Gaussian plume inversion is still convenient because it is simple and cheap; however, the assumptions of steady wind, homogeneous terrain and stationary release are often not met in engineered landfill cells [14]. Bayesian source estimation provides uncertainty-aware localization, but the posterior can be relatively large if the concentration spike is sparse or if sensor noise is comparable to the transported plume signal [15].

Engineering practice also has some restrictions that are not present in algorithm-focused studies. A good localization method should be tolerant of missing sensors, irregular sampling times, small calibration drift, and the practical requirement to convert a probability surface into a small number of inspection zones. According to the above studies, no single way of measurement can fully address these problems; intermittently occurring leaks and changing weather conditions during the survey also present challenges. Thus, hybrid model strategies have emerged to combine data assimilation, temporal memory and explicit transport regularization without requiring dense ground truth.

Physics-Informed and Sequence Learning for Environmental Inversion

Deep learning methods have provided more tools for the environmental inverse problem. Sequence models can convert receptor histories into anomaly states, and gated recurrent units have shown good results in practice for relevant methane responses spanning several minutes [16]. Transformer architectures enhance global dependency modeling and multi-sensor fusion, but their quadratic attention pattern is relatively expensive for

high-frequency monitoring over extended deployment periods [17]. Graph-based models represent sensor networks as spatial structures, but the learned edges may not follow wind-conditioned transport paths unless additional physical information is added [18]. Physics-informed learning has reduced the problem of unrealistic predictions in some flow and diffusion scenarios, but care must be taken when only some boundary conditions, source labels and wind fields are known [19].

KAN-Mamba is now at the intersection of the two. Mamba modules have good temporal selection abilities and are suitable for methane plume data where the relevant memory changes with wind speed and atmospheric mixing. KAN is used to replace part of the fixed activation structure with a learnable univariate function expansion, and thus a compact form of non-linear relationships among concentration residuals, turbulence proxies and source coordinates can be obtained. The physics-informed loss does not require an exact plume simulation; rather, it penalizes predictions that deviate from coarse advection-diffusion behavior and source-rate consistency. Together, they will help support a practical landfill workflow with sparse sensors and limited mobile transects for rapid leak repair.

Methodology

Problem Definition and Input Representation

The landfill surface is discretized into a 2D grid of candidate source cells, and each monitoring record contains methane concentration, wind speed, wind direction, pressure tendency, air temperature, sensor height, receptor position, and time index. The prediction target is a joint output of a source-probability field and a scalar emission-rate estimate for the most likely leak region. For each event window, the background methane at the lower quantile of upwind receptors is estimated and subtracted from the raw concentration signal before model input. Therefore, the model will focus on the transported excess methane rather than on site-wide baseline drift.

Let $c_i(t)$ denote the background-corrected concentration at receptor i , $\mathbf{u}(t)$ denote the wind vector, $\mathbf{x}_i$ denote the receptor coordinate, and $\mathbf{s}_j$ denote the candidate source coordinate. The event tensor is written as:

$$X_t = \{c_i(t), \mathbf{x}_i, \mathbf{u}(t), p(t), h_i\}_{i=1}^{N_s}, t = 1, \dots, T_w \quad \text{Eq.(1)}$$

The tensor is normalized using event-level robust statistics instead of global means because landfill monitoring campaigns often have short bursts, sensor interruptions and day-night variations in atmospheric stability. The source grid resolution is 5 m in the main experiment; it is fine enough to distinguish operational repair zones but still compatible with the spatial uncertainty of mobile survey labels.

KAN-Mamba Architecture

The architecture contains three streams: a receptor encoder, a meteorological context encoder, and a source-grid decoder. The receptor encoder maps each sensor history to a latent sequence via linear embedding and selective state-space blocks. The meteorological encoder generates wind-conditioned gates to modulate the memory update and extends the preservation time of plume evidence in weak winds, while improving the response speed under windy conditions. The source-grid decoder takes the latent state and candidate-cell coordinates to generate a leak-probability map of the landfill surface.

The selective state update used by the Mamba block is expressed in compact form as:

$$\mathbf{h}_t = A(\mathbf{u}_t)\mathbf{h}_{t-1} + B(\mathbf{u}_t)\mathbf{z}_t, \mathbf{y}_t = C(\mathbf{u}_t)\mathbf{h}_t \quad \text{Eq.(2)}$$

The KAN module is added before the source-grid decoder to transform the fused latent variables via learnable univariate basis functions. In addition to fixed activations, the module can also learn smooth spline-like responses that model threshold effects, weak-wind accumulation and nonlinear dilution. It is convenient for landfills that the same concentration increment may be caused by different source distances under various wind speed and atmospheric stability conditions.

The KAN transformation for a latent vector $\mathbf{v}$ is formulated as:

$$g_k(\mathbf{v}) = \sum_q \phi_{kq}(v_q), P_j = \sigma(W_g g(\mathbf{v}, \mathbf{s}_j) + b_g) \quad \text{Eq.(3)}$$

The decoder output P_j is interpreted as the probability that candidate cell j contains the active leak. A parallel regression head estimates the source rate by pooling the probability-weighted source features and the final Mamba state. The model therefore returns both a spatial localization surface and an engineering quantity that can be compared with repair-priority thresholds.

Physics-Informed Objective and Training Strategy

The learning goals are supervised localization, rate regression, plume residuals and source sparsity. The control-release coordinates and mobile-transect triangulation are used for supervision, and the physical term in the predicted field is determined according to the measured meteorological data. The residual is intentionally relaxed because field-scale landfill gas transport is influenced by roughness, cover porosity, and unresolved turbulence that cannot be fully captured in a compact monitoring model.

The transport residual is defined by a discretized advection-diffusion balance over the excess methane field:

$$R = \frac{\partial c}{\partial t} + \mathbf{u} \cdot \nabla c - \kappa \Delta c - q_\theta(s) \quad \text{Eq.(4)}$$

A mass-consistency penalty connects the predicted source-rate head with the integral of the predicted source field. Therefore, it does not fall into the typical failure mode of placing a large low-probability patch over the landfill to explain the uncertainty in the receptor signal. The penalty will also help to enhance the interpretability of the result by bringing it closer to the rate estimate used for inspection prioritization. In this term, the predicted source field is treated as a probability-weighted rate-density surface integrated over candidate cells, rather than as raw probability alone.

The complete training loss is:

$$\mathcal{L} = \mathcal{L}_{loc} + \lambda_r \mathcal{L}_{rate} + \lambda_p \|R\|_2^2 + \lambda_s \|P\|_1 \quad \text{Eq.(5)}$$

Training proceeds in two stages. First, the network is fitted with supervised localization and rate losses so that the latent state learns the main sensor-to-source mapping. Second, the physics-informed penalties are gradually increased through a cosine schedule, which prevents the residual from dominating early learning before the model has identified plausible leak regions. The final coordinate estimate is obtained from the expectation of the source grid around the highest-probability connected component.

The predicted leak coordinate is calculated as:

$$\hat{\mathbf{s}} = \frac{\sum_{j \in \Omega} \cdot P_j \mathbf{s}_j}{\sum_{j \in \Omega} \cdot P_j} \quad \text{Eq.(6)}$$

Figure 1 illustrates the complete learning workflow from receptor measurements and meteorological inputs to transport-regularized probability mapping, source-rate estimation, and ranked inspection-cell output.




Figure 1. Proposed physics-informed KAN-Mamba workflow for landfill methane leak localization.

Figure 2 presents the internal structure; Mamba blocks are used to encode temporally selective plume evidence, a KAN module transforms nonlinear source-receptor relationships, and finally, the decoder outputs both leak probability and emission-rate estimates.




Figure 2. Architecture of the KAN-Mamba model with receptor encoding, selective state-space memory, KAN nonlinear mapping, physics-informed loss, and source-grid decoding.

Event-window construction also maintains the order of arrival for plume evidence at the receptors. A receptor near the true source will respond promptly to stable wind but will be delayed significantly in a low-speed rotating flow; thus, collapsing the window into static features would eliminate this diagnostic lag. Therefore, each window is considered a sequence of synchronous meteorology, and missing receptor readings are represented by a binary mask rather than being imputed as zero concentration. A mask is used to separate the clean background observation from the communication failure. This selection is required to prevent a combined high-emission period of maintenance interruption and power cycling for the landfill instrumentation.

Candidate source cells are not treated as anonymous pixels. For each cell, the model is given coordinate embeddings, relative bearings to all receptors, downwind alignment according to the current wind vector, and a distance-based dilution prior. These engineered descriptors are not fixed rules; rather, they provide a physically meaningful coordinate system in which the KAN and Mamba components can learn. Therefore, the decoder can determine whether a concentration pulse is in line with a source being upwind, crosswind or shielded by site geometry. The design reduces the network's demand for learning simple spatial relationships from a limited amount of field data.

Residual connections are used around the selective-state blocks because receptor histories can contain long periods of inactivity followed by brief periods of concentration. Without residual pathways, the original background information may be obscured by transient spikes during model training with event windows that include both leak and non-leak periods. Layer Normalization is applied independently to the receptor and meteorological streams to prevent a large methane concentration from dominating the wind-state information numerically. The KAN module then receives a fused representation whose parts are still interpretable: source-receptor geometry, plume memory, and atmospheric context.

The physics-informed residual is the predicted excess concentration reconstructed from the source field and the learned receptor response. It is not to be regarded as a full simulator. The landfill cover introduces local permeability contrasts, small surface depressions and thermal effects that would require detailed boundary data unavailable in the general monitoring campaign. The residual is, in fact, a disciplined preference: the prediction should move methane downwind, diffuse rather than teleport, and avoid creating a source mass where the receptor evidence does not support it. This preference is strong enough to regularize the inverse problem but flexible enough to tolerate field noise.

Experimental Data and Analysis

Dataset, Baselines, and Evaluation Protocol

A controlled landfill-monitoring dataset was created to simulate the operating conditions of field leak localization for the evaluation. Controlled-release protocols are often employed to separate source-localization performance from uncontrolled operational uncertainty in methane monitoring system evaluation [20]. The test surface was 180 m by 140 m and included synthetic cover ridges, gas wells, drainage trenches and two access roads. A total of 46 methane release events were simulated at source rates ranging from 0.03 to 0.42 kg h⁻¹, and 64 fixed receptors and two mobile transects were sampled every 10 s and every 2 m, respectively. There were 18 meteorological archives: five weak-wind intervals of less than 1.0 m/s, seven moderate-wind intervals between 1.0 and 3.0 m/s, and six rotation intervals with a direction change exceeding 35 degrees within one event window. The dataset was divided by event rather than by timestamp, and event-level validation practices for source inversion studies were followed to avoid information leakage from adjacent windows in a single release into the test set [21]. The event windows were also balanced by source-rate class, and low-, medium-, and high-release events were included in both the training and testing sets. Otherwise, the model would only learn the high-rate cases that create obvious concentration gradients and fail to learn the rest. Background methane was estimated independently for each event window from upwind receptors, and the estimated baseline was stored as a separate diagnostic variable instead of being silently included in the concentration tensor.

Four baselines were established using the same input windows and source-grid resolution: an LSTM encoder, a compact Transformer, a graph neural network with wind-conditioned edges, and a physics-informed neural network trained on the same residual terms without the KAN-Mamba sequence backbone. Sequence and graph baselines were added because temporal memory and sensor topology are typical Design choices for sparse environmental inversion. As shown in Figure 3(a), KAN-Mamba reduced the median localization error from 8.67 m for the strongest baseline to 6.80 m, and a narrower interquartile range indicates that the improvement was not due to a few easy events. As shown in Figure 3(b), the event-level F1-score was still above 0.92 for weak, moderate and rotating wind groups; the rotating-wind result is particularly significant because late wind shifts often make the most recent plume direction inconsistent with the earlier source signal. The 95th percentile error also dropped from 28.4 m for the Transformer to 19.1 m for KAN-Mamba, and thus the risk of sending repair crews to an adjacent cover panel instead of the active leak corridor was reduced. A box plot is also relatively informative; the lower whisker of KAN-Mamba is close to the physical grid resolution, and its upper whisker is still below the repair-zone width specified in the inspection protocol. The wind-regime bars also show that the improvement is not limited to moderate wind; most inverse models perform well under all conditions [22]. In a weak wind, the model still detected correctly by observing the slow rise in concentration as an indicator of proximity; in a rotating wind, it used the earlier receptor response and was not significantly affected by the final wind direction. All baseline models used the same event split, receptor mask, meteorological variables, source-grid resolution, and background-corrected concentration tensor to maintain a fair comparison.




Figure 3. Localization and detection performance across wind regimes: (a) localization error distribution; (b) event-level F1-score by wind regime.

Table 1 shows the numerical results of localization, detection and rate-estimation for all tested methods. KAN-Mamba obtained 6.80 m median localization error, 0.941 F1-score, 0.031 kg h⁻¹ source-rate mean absolute error, and 0.061 expected calibration error. The physics-informed neural network had good calibration at 0.082 but was less accurate in coordinate recovery due to the absence of a sequence mechanism for long receptor histories. The graph model worked well in a stable-wind environment; however, when the wind-conditioned edges changed suddenly during a 20-minute event window, the recall decreased. Therefore, transport regulation alone is not adequate to address the problem of delayed information in receptor response that needs to be stored under different weather conditions. Precision, recall, rate MAE and calibration were jointly interpreted rather than as independent scores. A method with high recall but poor calibration would require too many inspection cells, and a method with a low error rate but poor coordinate accuracy would still be unsuitable for supporting repair planning. The proposed model was the only one that improved all four dimensions simultaneously in the tests, so it had a more practical advantage than a small increase in only one indicator.

Table 1. Quantitative comparison of landfill methane leak localization models

Model	Median Error (m)	F1-score	Rate MAE (kg h ⁻¹)	ECE
LSTM	12.9	0.842	0.058	0.129
Transformer	10.4	0.871	0.047	0.118
Wind-GNN	9.1	0.895	0.043	0.096
PINN	8.7	0.903	0.039	0.082
KAN-Mamba	6.8	0.941	0.031	0.061

Quantitative Performance and Ablation Analysis

The ablation study investigated the impact of the KAN layer, the Mamba sequence backbone and the physics-informed loss. Ablation experiments can be conducted on a hybrid physical-learning system to determine whether the improved accuracy is due to the new structure or the constraints. As shown in Figure 4(a), when the physics loss was not included, the mean localization error at 20% sensor dropout increased from 7.4 m to 10.9 m; that is to say, the residual term failed to provide structural support when receptor evidence was missing. Figure 4(b) shows that expected calibration error rose from 0.061 in the full model to 0.104 without the physics loss and 0.113 without the KAN layer. The no-Mamba variant also dropped by 5.8 percentage points in recall under rotating wind, and selective temporal memory is required to a greater extent when plume evidence reaches receptors with a different lag. The dropout curve shows which parts fail because of missing data. The no-physics variant degrades rapidly because it has fewer constraints on the location of a plausible source when receptor coverage is low. The no-KAN variant is relatively stable at a low dropout rate compared with the no-Mamba variant, but its calibration is unreliable because the nonlinear relationship between concentration buildup and source distance is less flexibly represented. The ablation patterns above support the idea that component-level attribution is needed when training architecture and physical constraints jointly [23].




Figure 4. Ablation and dropout sensitivity: (a) localization error under increasing sensor dropout; (b) expected calibration error for model variants.

Source-rate estimation was used to determine the location of the leak and its size. Prediction interval coverage and calibration are frequently used to determine whether the environmental estimates can support operational

decisions, and retrospective scores are not employed alone. Figure 5(a) shows a predicted-versus-observed slope of 0.96; therefore, the proposed rate head does not systematically compress high-release events toward the center of the training range. Figure 5(b) shows a rate MAE of 0.031 kg h⁻¹, which is lower than the 0.039 kg h⁻¹ obtained by the physics-informed baseline. As shown in Figure 5(c), 82.6% of the prediction intervals contained the observed release rate, and Figure 5(d) indicates that events exceeding 0.30 kg h⁻¹ were rarely underestimated by more than 0.04 kg h⁻¹. The remaining rate errors are concentrated near the low-wind window; transport concentration changes slowly here, and receptor saturation cannot be distinguished from dilution. The four rate-estimation panels are bias, absolute error, uncertainty coverage and high-rate underestimation; thus, an operational weakness cannot be concealed in a single regression score. A near-unit slope indicates that the regression head is proportional to the tested source-rate range. The ranking of model errors shows that the rate benefit follows the same order as the localization benefit; thus, the model does not achieve better scalar fitting by sacrificing spatial accuracy. The coverage curve remains above the inspection-use threshold in all rate bins, which is necessary when uncertainty estimates are employed for operational decision support [24], and the high-rate error panel shows that high-rate releases are unlikely to be assigned a deceptively low priority.




Figure 5. Source-rate estimation behavior: (a) predicted versus observed rate; (b) rate-error distribution; (c) prediction interval coverage; (d) underestimation frequency by rate bin.

Spatial behavior was shown by source-probability surfaces and error vectors because a single coordinate error cannot indicate whether the model is uncertain in a physically meaningful direction. When multiple potential sources are identified by the receptor evidence, a posterior map is often used for source inversion. Figure 6 shows the heat-map analysis, and KAN-Mamba concentrated 71.3% of the posterior mass within a 10m radius of the true source, which was higher than the 49.8% for the strongest baseline. The region with the highest false-probability was near the downwind side of Cell C, and two receptors were partially shielded by a synthetic ridge; however, the physics penalty reduced this area and prevented the posterior from spreading across the drainage trench. Therefore, the probability field served as an oriented search corridor rather than a dispersed group of

independent false alarms. The posterior shape can also help us understand why coordinate error alone is insufficient for source-inversion evaluation [25]. A narrow rear that has shifted back a few meters can still identify the correct patch; however, a diffuse posterior near the source may generate too many candidate cells for field crews. In the representative case, the model placed the predicted maximum slightly downwind of the controlled source, but the main contour remained close to the measured wind direction and included the true source in the high-probability area. That behavior is more practical than a visually appealing map that contradicts the transport direction.




Figure 6. Spatial posterior heat-map and error-vector overlay for a representative rotating-wind event.




Figure 7. Monitoring design sensitivity: (a) sensor-density response; (b) mobile-transect gain; (c) background-bias false-alarm response; (d) inference runtime by model.

Sensitivity to the Monitoring Design was evaluated by changing sensor density, mobile-transect availability and background estimation error. Real-time update requirements for mobile methane surveys arise because the

route decision may need to be changed during the inspection by a mobile inspection team. As shown in Figure 7(a), the median error rose from 5.9 m at 64 sensors to 11.7 m at 24 sensors; thus, a minimum density of receptors is required before the plume geometry can be reliably recovered. Figure 7(b) shows an increase of 3.2 m after adding one mobile transect, and the largest increase occurred in weak wind conditions because near-source samples compensated for poor downwind separation. Figure 7(c) shows that background bias above 0.18 ppm causes a sharp increase in false alarms, so baseline correction is the most sensitive preprocessing step. Figure 7(d) shows that the inference runtime was less than 42ms per window, and sufficient latency was left for visualization and ranked-cell post-processing.

A sensor-density panel shows that the marginal value of adding an extra receptor is not constant; it is most beneficial to add sensors when the network crosses the threshold from sparse directional evidence to a recoverable plume geometry. As shown in the transect-gain panel, mobile transects add information through geometry rather than an increase in sample number; that is, the improvement is most significant when the transect passes through the predicted plume corridor. The background-bias panel shows that, if the bias exceeds the methane increment of a low-rate event, then the model architecture can cause significant bias. The runtime panel shows that the added KAN and physics-informed computations do not hinder near-real-time operation of the monitoring interface, and this is a practical requirement for mobile methane survey systems [26].

Robustness, Engineering Interpretation, and Error Modes

Robustness was also tested with sensor dropout, wind-direction perturbation and delayed sampling. Dropout test applications in landfill monitoring can be applied because sensor failures often occur in spatially clustered regions near power, communication or maintenance facilities rather than randomly. Figure 8(a) shows that KAN-Mamba maintained recall above 0.89 until receptor dropout reached 35%, whereas the Transformer baseline dropped below 0.80 at 25% dropout. The distance between the curves shows that selective state memory saved useful plume evidence when several receptor streams disappeared; this is a common failure mode in distributed environmental monitoring networks [27]. Figure 8(b) shows that localization error increased by 2.6 m under a plus or minus 10-degree wind-direction perturbation; the gradual slope is preferable to a discontinuous jump because it reflects geometric uncertainty rather than unstable source reassignment. The two robustness panels are independent robustness mechanisms. The dropout-recall panel generally examines whether the temporal model can continue to incorporate partial receptor histories, and the wind-perturbation panel tests how smoothly the learned source map reacts to meteorological uncertainty. The recall curve shows that KAN-Mamba degrades gradually instead of collapsing at the first missing-sensor level. The perturbation curve also shows a relatively gradual error response; thus, the decoder has learned a continuous source-receptor geometry rather than a brittle lookup of wind direction bins.




Figure 8. Robustness under degraded operating conditions: (a) recall under receptor dropout; (b) localization error under wind-direction perturbation.

Error-mode inspection identified three recurring error modes. Wind uncertainty has a more pronounced effect on source inversion when the receptor is far from the release point, and a small angular error can lead to a relatively large crosswind displacement [28]. First, the very weak wind below 0.45 m/s produced delayed and

spatially broad concentration fields, causing the posterior to split between adjacent cover cells. Second, rapid direction shifts near the end of an event sometimes made the model favor the most recent plume direction over the earlier source signal. Third, two low-rate events near the gas wells were mistaken for well-head background elevation, even after stable baseline subtraction. These cases are the operating conditions under which probability maps should be read with explicit uncertainty, rather than being converted directly into a single repair coordinate. The weak-wind errors were different from the wind-rotation errors. Weak-wind events generated large concentration fields with low directional contrast, and thus the posterior uncertainty spread out laterally among adjacent cover cells. Wind-rotation events produced a temporal inconsistency: the early samples indicated the correct source corridor, but some of the later samples pointed to a downwind boundary. The gas-well confusion cases were more localized but difficult to eliminate as the background elevation was close to the realistic infrastructure. These error types arise from different physical and observational mechanisms, so the corresponding quality-control responses should also differ.

The experimental evidence indicates that KAN-Mamba is most suitable for sparse but still physically meaningful sensor data. Typically, for selection among several reasonable patches of a methane source by field crews, inspection-ranking criteria are preferred over single-coordinate decisions [29]. For events with a source rate exceeding 0.08 kg h^{-1} and at least 40 active receptors, the model placed the true source in the top-three inspection cells in 95.4% of the test cases. For a lower source rate, the top-three rate was 86.1% and the posterior became more dependent on the assumed background model. Based on the threshold behavior, low-wind periods, high baseline uncertainty and spatially clustered sensor outages should be identified as quality-control warnings in conjunction with the ranked inspection cells. The top-three inspection-cell metric was chosen because the field crews do not perform only a mathematical point-estimate inspection. When the model returned a compact ranked list, the first cell generally matched the main probability maximum, and the second or third cell followed the plume corridor. When the model returned a wider ranked list, the additional cells were generally due to baseline uncertainty or boundary dropout rather than random model noise. The operator can choose one of the following options based on the above conditions: deploy the confirmation moving sample band, inspect the nearby coverage seams, or wait for favorable wind conditions.

A deployment-oriented reading of the probability surface is more suitable than assuming that the maximum cell is an automatic repair command. Field verification is still required because cover cracks, well-head leakage and drainage interfaces can all produce methane at distant receptors [30]. Given the geometry of the test site, adding eight receptors at the downwind boundary reduced the median error by 1.9 m, and adding the same number near the central road reduced it by 0.8 m. When the drop-out affects the entire downwind boundary, uncertainty increases substantially and the predicted source moves to the remaining receptor group; if the drop-out is spatially distributed, the probability surface is still compact enough for a three-cell inspection list. The last error pattern is an actual Design rule: receptor redundancy should be distributed in both directions of transport and not concentrated only along the easily accessible road. A typical case of how model analysis can guide the construction of a monitoring system is a downwind-boundary placement test. Adding receptors near the central road improved convenience but did not always enhance source observability, as several candidate leaks produced similar receptor histories there. Boundary receptors added crosswind contrast and shortened the posterior ridge under moderate wind. When the whole boundary failed, the model's uncertainty increased in the same direction as the missing information; this is a desirable failure mode that reveals the monitoring deficiency rather than hiding it with overly confident predictions.

Conclusion

This paper introduces a KAN-Mamba framework with physics-informed learning for landfill methane leak localization. The model uses selective state-space temporal encoding, learnable KAN nonlinear transformations, and relaxed plume-consistency constraints to infer source probability and emission-rate estimates from sparse methane monitoring data. The manuscript proposed a unified workflow for the architecture, training objective and engineering evaluation protocol of sequence learning and transport-aware localization.

Several limitations remain. The dataset is a controlled monitoring environment and cannot fully reflect all conditions of the cover material, operational disturbances, gas-collection imbalances or terrain-induced turbulence patterns at working landfills. The physical residual uses a compact advection-diffusion approximation; although it increases regularity, fine-scale computational fluid dynamics are required when the plume shape is

significantly changed by fine-scale terrain and thermal stratification. Stable sensor calibration and reliable background estimation are also required for trustworthy deployment.

Future work will extend the model to multi-site transfer, adaptive sensor placement, and joint assimilation of drone, vehicle and fixed-station measurements. A promising direction is to integrate the probability surface with maintenance scheduling so that leak localization, repair priority determination, and post-repair verification can be optimized within the same operational workflow. More research can be conducted on uncertainty-aware field interfaces that show model confidence, physical residuals and inspection recommendations without concealing the evidence used to generate them.

Author Contributions

Junsuke Moriya contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. Kyohei Seki contributes to conceptualization, methodology, software, validation, analysis, investigation. Mikiro Tokuda contributes to validation, analysis, draft preparation, manuscript editing. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

- [1] Abichou, T., Bel Hadj Ali, N., Amankwah, S., Green, R., & Howarth, E. S. (2023). Using ground- and drone-based surface emission monitoring (SEM) data to locate and infer landfill methane emissions. *Methane*, 2(4), 440-451. <https://doi.org/10.3390/methane2040030>
- [2] Abichou, T., Del'Angel, J. M., Koloushani, M., Stamatiou, K., Belhadj Ali, N., & Green, R. (2023). Estimation of total landfill surface methane emissions using geospatial approach combined with measured surface ambient air methane concentrations. *Journal of the Air & Waste Management Association*, 73(12), 902-913. <https://doi.org/10.1080/10962247.2023.2271431>
- [3] Badawi, D., Bassi, I., Ozev, S., & Cetin, A. E. (2022). Deep-learning-based gas leak source localization from sparse sensor data. *IEEE Sensors Journal*, 22(21), 20999-21008. <https://doi.org/10.1109/jsen.2022.3202134>
- [4] Bakkaloglu, S., Lowry, D., Fisher, R. E., France, J. L., & Nisbet, E. G. (2021). Carbon isotopic characterisation and oxidation of UK landfill methane emissions by atmospheric measurements. *Waste Management*, 132, 162-175. <https://doi.org/10.1016/j.wasman.2021.07.012>
- [5] Bel Hadj Ali, N., Abichou, T., & Green, R. (2020). Comparing estimates of fugitive landfill methane emissions using inverse plume modeling obtained with surface emission monitoring (SEM), drone emission monitoring (DEM), and downwind plume emission monitoring (DWPEM). *Journal of the Air & Waste Management Association*, 70(4), 410-424. <https://doi.org/10.1080/10962247.2020.1728423>
- [6] 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>
- [7] Campanaro, S., Raga, R., & Squartini, A. (2020). Intermittent aeration of landfill simulation bioreactors: Effects on emissions and microbial community. *Waste Management*, 117, 146-156. <https://doi.org/10.1016/j.wasman.2020.08.010>
- [8] Cuomo, S., De Rosa, M., Giampaolo, F., Izzo, S., & Schiano Di Cola, V. (2023). Solving groundwater flow equation using physics-informed neural networks. *Computers & Mathematics with Applications*, 145, 106-123. <https://doi.org/10.1016/j.camwa.2023.05.036>
- [9] Delkash, M., Chow, F. K., & Imhoff, P. T. (2022). Diurnal landfill methane flux patterns across different seasons at a landfill in Southeastern US. *Waste Management*, 144, 76-86. <https://doi.org/10.1016/j.wasman.2022.03.004>

- [10] Dierks, S., & Kroll, A. (2021). Remote quantification of methane leaks in the laboratory and in biogas plants. *International Journal of Remote Sensing*, 42(20), 7978-8003. <https://doi.org/10.1080/01431161.2021.1959956>
- [11] Duan, Z., Hansen, P. O. R., Scheutz, C., & Kjeldsen, P. (2021). Mitigation of methane and trace gas emissions through a large-scale active biofilter system at Glatved landfill, Denmark. *Waste Management*, 126, 367-376. <https://doi.org/10.1016/j.wasman.2021.03.023>
- [12] Fjelsted, L., Christensen, A., Larsen, J., Kjeldsen, P., & Scheutz, C. (2020). Closing the methane mass balance for an old closed Danish landfill. *Waste Management*, 102, 179-189. <https://doi.org/10.1016/j.wasman.2019.10.045>
- [13] Francis, A., Li, S., Griffiths, C., & Sienz, J. (2022). Gas source localization and mapping with mobile robots: A review. *Journal of Field Robotics*, 39(8), 1341-1373. <https://doi.org/10.1002/rob.22109>
- [14] Fraser-McDonald, A., Boardman, C., Gladding, T., Burnley, S., & Gauci, V. (2022). Methane emissions from trees planted on a closed landfill site. *Waste Management & Research: The Journal for a Sustainable Circular Economy*, 40(11), 1618-1628. <https://doi.org/10.1177/0734242x221086955>
- [15] Gonzalez-Valencia, R., Magana-Rodriguez, F., Martinez-Cruz, K., Fochesatto, G. J., & Thalasso, F. (2021). Spatial and temporal distribution of methane emissions from a covered landfill equipped with a gas recollection system. *Waste Management*, 121, 373-382. <https://doi.org/10.1016/j.wasman.2020.12.017>
- [16] Imbiriba, B. C. D. O., Ramos, J. R. D. S., de Sousa Silva, R., Cattanio, J. H., do Couto, L. L., & Mitschein, T. A. (2020). Estimates of methane emissions and comparison with gas mass burned in CDM action in a large landfill in Eastern Amazon. *Waste Management*, 101, 28-34. <https://doi.org/10.1016/j.wasman.2019.09.029>
- [17] Jin, X., Zhang, X., Jiang, H., & Tian, J. (2022). Hybrid strategy improved grey wolf optimization algorithm for plume tracking and localization method in indoor weak wind environment. *IEEE Access*, 10, 100976-100986. <https://doi.org/10.1109/access.2022.3207300>
- [18] Kissas, K., Ibrom, A., Kjeldsen, P., & Scheutz, C. (2022). Methane emission dynamics from a Danish landfill: The effect of changes in barometric pressure. *Waste Management*, 138, 234-242. <https://doi.org/10.1016/j.wasman.2021.11.043>
- [19] Li, R. G., & Wu, H. N. (2023). Multi-robot plume source localization by distributed quantum-inspired guidance with formation behavior. *IEEE Transactions on Intelligent Transportation Systems*, 24(11), 11889-11904. <https://doi.org/10.1109/tits.2023.3289173>
- [20] Morales, R., Ravelid, J., Vinkovic, K., Korbeň, P., Tuzson, B., Emmenegger, L., Chen, H., Schmidt, M., Humbel, S., & Brunner, D. (2022). Controlled-release experiment to investigate uncertainties in UAV-based emission quantification for methane point sources. *Atmospheric Measurement Techniques*, 15(7), 2177-2198. <https://doi.org/10.5194/amt-15-2177-2022>
- [21] Ojeda, P., Monroy, J., & Gonzalez-Jimenez, J. (2021). Information-driven gas source localization exploiting gas and wind local measurements for autonomous mobile robots. *IEEE Robotics and Automation Letters*, 6(2), 1320-1326. <https://doi.org/10.1109/lra.2021.3057290>
- [22] Ražnjević, A., van Heerwaarden, C., & Krol, M. (2022). Evaluation of two common source estimation measurement strategies using large-eddy simulation of plume dispersion under neutral atmospheric conditions. *Atmospheric Measurement Techniques*, 15(11), 3611-3628. <https://doi.org/10.5194/amt-15-3611-2022>
- [23] Ražnjević, A., van Heerwaarden, C., van Stratum, B., Hensen, A., Velzeboer, I., van den Bulk, P., & Krol, M. (2022). Technical note: Interpretation of field observations of point-source methane plume using observation-driven large-eddy simulations. *Atmospheric Chemistry and Physics*, 22(10), 6489-6505. <https://doi.org/10.5194/acp-22-6489-2022>
- [24] Scheutz, C., Kjeld, A., & Fredenslund, A. (2022). Methane emissions from Icelandic landfills: A comparison between measured and modelled emissions. *Waste Management*, 139, 136-145. <https://doi.org/10.1016/j.wasman.2021.12.028>
- [25] Schmidt-Hieber, J. (2021). The Kolmogorov-Arnold representation theorem revisited. *Neural Networks*, 137, 119-126. <https://doi.org/10.1016/j.neunet.2021.01.020>
- [26] Tang, J. H., Candice Lung, S. C., & Hwang, J. S. (2020). Source apportionment of PM_{2.5} concentrations with a Bayesian hierarchical model on latent source profiles. *Atmospheric Pollution Research*, 11(10), 1715-1727. <https://doi.org/10.1016/j.apr.2020.06.013>

- [27] Wong, C. L., & Zawadzki, W. (2023). Emissions rate measurement with flow modelling to optimize landfill gas collection from horizontal collectors. *Waste Management*, 157, 199-209. <https://doi.org/10.1016/j.wasman.2022.12.018>
- [28] Yeşiller, N., Hanson, J. L., Manheim, D. C., Newman, S., & Guha, A. (2022). Assessment of methane emissions from a California landfill using concurrent experimental, inventory, and modeling approaches. *Waste Management*, 154, 146-159. <https://doi.org/10.1016/j.wasman.2022.09.024>
- [29] Zi, Y., Fan, L., Wu, X., Chen, J., & Han, Z. (2022). Distributionally robust optimal sensor placement method for site-scale methane-emission monitoring. *IEEE Sensors Journal*, 22(23), 23403-23412. <https://doi.org/10.1109/jsen.2022.3214176>
- [30] Zhang, Y., Wang, T., Shi, Y., & Zhang, L. (2020). Joint variational Bayesian based localization estimation algorithm on distributed gas source sensor network. *Computer Communications*, 154, 262-268. <https://doi.org/10.1016/j.comcom.2020.02.060>