Scalable Air-Quality Sensor Placement via Gradient-Based Mutual Information Maximization

Zeel B Patel, Vinayak Rana, Nipun Batra

Indian Institute of Technology Gandhinagar, India
{patel_zeel, vinayak.rana, nipun.batra}@iitgn.ac.in

Abstract

Air pollution is a leading global health threat, yet many developing countries lack the dense monitoring infrastructure needed for accurate exposure assessment and informed policy. Optimal Sensor Placement (OSP) is a foundational challenge in expanding monitoring capacity. While mutual information (MI) offers a principled criterion for selecting informative sensor locations, its computational cost grows with both the number of placements and the density of the candidate grid. We present a scalable, continuous optimization framework that treats sensor coordinates as differentiable parameters and directly maximizes MI. Unlike standard approaches, our method is computationally efficient—its runtime is independent of both the number of placements and the size of the search grid—making MI-based acquisition feasible over large spatial domains. On a continental-scale PM2.5 dataset, our method outperforms random placement and the widely-used Maximum Predictive Variance heuristic. In a focused regional study, it approaches the performance of greedy MI while being orders of magnitude faster. Our framework enables practical, information-theoretic sensor placement for real-world environmental monitoring.

Code — https://github.com/sustainability-lab/gdmi-aqs

Introduction

Air pollution is one of the most urgent public health challenges of our time. According to the World Health Organization (WHO), over 99% of the global population breathes air that exceeds recommended pollutant thresholds (World Health Organization 2024). A central tenet of environmental policy is that we cannot manage what we do not measure. Yet in many developing countries, the lack of ground-based monitoring infrastructure poses a fundamental barrier to both exposure assessment and effective intervention.

India exemplifies this challenge. Despite a population exceeding 1.4 billion, the country operates only around 500 official monitoring stations (Central Pollution Control Board 2025). Without high-resolution local data, policymakers lack the evidence needed to design targeted public health interventions. To address this measurement gap, recent work has explored using machine learning to estimate pollution levels in unmonitored regions (van Donkelaar et al. 2016). However, the accuracy of these models depends critically on the spatial distribution of the sensors. This raises a foundational question: where should new sensors be placed to maximize predictive performance and minimize error? Given the high cost of regulatory-grade stations, universal deployment is economically unfeasible. As a result, Optimal Sensor Placement (OSP) has emerged as a core problem in data-driven environmental science.

Principled methods grounded in information theory, such as maximizing mutual information (MI) between selected sensors and the unobserved field (Krause, Singh, and Guestrin 2008), can yield high-quality placements. However, MI-based acquisition requires a look-ahead evaluation at each candidate site to estimate potential information gain, making its runtime scale linearly with the number of candidates—a bottleneck in large domains. On the other hand, lightweight heuristics like greedy selection based on predictive variance offer scalability but sacrifice accuracy. These often result in spatial clustering or redundant placements, as they fail to account for interactions among sensors. Moreover, such methods frequently place sensors along domain boundaries where model uncertainty is artificially inflated due to extrapolation, further degrading the quality and coverage of the monitoring network.

In this work, we propose a method that directly addresses the trade-off between accuracy and scalability in sensor placement. We recast the discrete selection problem as a continuous optimization task, enabling gradient-based search over sensor coordinates. This formulation resolves two core limitations of prior work. First, by treating sensor locations as differentiable parameters, our approach avoids evaluating every candidate site and scales independently of grid resolution—making mutual information (MI) acquisition tractable in large spatial domains (Krause, Singh, and Guestrin 2008). Second, unlike greedy or distance-based heuristics that select sensors sequentially (Sun et al. 2025), we jointly optimize the full batch of new placements. This enables the model to account for spatial correlations and redundancy across sensors, resulting in more informative and geographically balanced monitoring networks.

We validate our approach on a continental-scale reanalysis dataset of surface-level $PM_{2.5}$. Our method consistently outperforms the widely-used Maximum Predictive Variance

heuristic, achieving a **4%** reduction in predictive RMSE. Compared to classical mutual information (MI)–based selection, it is orders of magnitude faster while delivering competitive placement quality.

This combination of accuracy and scalability yields a practical tool for designing high-impact air quality monitoring networks that better support public health decision-making.

Related Work

Classical Sensor Placement: Information Theory and Submodularity

Classical approaches to Optimal Sensor Placement (OSP) are rooted in information theory and have traditionally relied on probabilistic surrogate models like Gaussian Processes (GPs). GPs are prized for their ability to provide well-calibrated uncertainty estimates, making them a natural fit for model-based placement strategies (Williams and Rasmussen 2006). Within this framework, a dominant paradigm is to select locations that maximize an information-theoretic criterion, such as the mutual information (MI) between the selected sensors and the unobserved field. A key theoretical breakthrough was the proof that the MI objective is submodular. This property allows a simple greedy algorithm to find a near-optimal set of sensor locations with strong performance guarantees (Krause, Singh, and Guestrin 2008).

However, these methods have significant computational limitations. Standard GPs require inversion of large covariance matrices, resulting in cubic time complexity ($O(N^3)$). Moreover, evaluating the MI acquisition function requires look-ahead computations at each candidate location to assess the marginal gain in information. As a result, the total runtime scales linearly with both the number of candidate sites and the number of placements. This makes classical MI-based approaches impractical for high-resolution or large-scale domains such as continental air quality monitoring, motivating the development of more scalable alternatives.

Modern Learning-Based Approaches

The scalability issues of classical methods have motivated a shift towards modern, data-driven strategies for OSP. Many approaches frame the problem as active learning, where a model iteratively queries for sensor data at the most informative locations to improve its performance (Settles 2009). More recently, meta-learning has emerged as a powerful paradigm (Finn, Abbeel, and Levine 2017). Models like Neural Processes (NPs) are trained across a wide distribution of tasks (e.g., different pollution fields) to learn an efficient placement strategy that can rapidly adapt to new environments from only a few observations (Garnelo et al. 2018b). This NP-based approach has been successfully applied to environmental sensor placement using sophisticated variants like Convolutional Gaussian Neural Process (ConvGNP) (Andersson et al. 2024).

However, a common limitation of these strategies is their reliance on discrete, sequential selection policies. The learned policies can be sensitive to the training distribution and are typically non-differentiable, preventing the direct, joint optimization of an entire batch of sensor locations at once.

Domain-Specific Work: Air Quality Monitoring

Finally, our work is situated within the specific domain of air quality monitoring, where sensor placement has significant societal implications. A large body of research has focused on forecasting air quality using deep learning, leveraging spatiotemporal models to predict pollutant concentrations (Xing et al. 2020). In parallel, a number of studies have addressed OSP for air quality, often focusing on specific urban areas and using methods like greedy selection or land-use regression to inform placements (Chang-Silva et al. 2024).

Recognizing the social dimensions of monitoring, a growing and vital area of research focuses on equitable placement strategies. These works aim to design networks that explicitly address environmental justice concerns by prioritizing coverage in vulnerable or historically underserved communities (Holstein et al. 2019).

However, these domain-specific approaches are typically limited by framing sensor placement as a discrete selection problem solved with suboptimal, greedy methods. Many existing methods are designed for smaller, city-scale problems and do not scale to continental-level domains, nor do they offer a mechanism to find the jointly optimal set of locations that maximally improves the overall predictive accuracy.

Problem Formulation

Let X_c and Y_c denote the locations and corresponding observations of the existing sensors. Given a deployment budget of k additional sensors to be selected from a candidate pool X_{pool} with cardinality $|X_{\text{pool}}| = n$, the objective is to identify a subset of locations $X_{\text{new}} \subseteq X_{\text{pool}}$, with $|X_{\text{new}}| = k$ minimizing a prescribed acquisition function defined over a target set of locations X_t.

Methodology

The optimization objective defined in the previous section involves selecting an optimal set of k new sensor locations from the pool X_{pool}, resulting in a combinatorial search space of size $\binom{|X_{\text{pool}}|}{k}$. Exploring this space exhaustively is computationally infeasible for any real-world deployment (Guestrin, Krause, and Singh 2005). To overcome this, we recast sensor placement as an active learning problem and adopt uncertainty sampling as our acquisition strategy, in line with prior work that has leveraged uncertainty-driven heuristics for sensor selection (Andersson et al. 2024). This approach prioritizes adding sensors in regions where the model is least confident, under the assumption that predictive uncertainty correlates with the potential generalization error.

A greedy naive placement of sensors based on maximum uncertainty fails to account for interaction effects between sensor locations and can lead to suboptimal joint configurations (Krause, Singh, and Guestrin 2008). Instead, we

jointly optimize all k sensor positions. This requires a surrogate model that produces principled uncertainty estimates over the entire spatial field.

To meet these requirements, we employ *Neural Processes* (NP) (Garnelo et al. 2018b) as our predictive model. NPs combine the flexibility of neural networks with the ability to produce coherent predictive distributions over functions, making them particularly well suited for data-efficient learning in spatiotemporal domains. This section details our modeling choices, acquisition strategies, and training procedures.

Neural Processes

Neural processes (NP) are a family of meta-learning models that combine the flexibility and scalability of deep neural networks with the principled uncertainty quantification traditionally provided by Gaussian processes. At their core, NPs learn to map a set of observed inputs and outputs (context) (X_c, Y_c) to a full posterior predictive distribution over a set of unlabelled points (target) (X_t). By training across many functions (in our case, monthly $PM_{2.5}$ fields), NPs learn an amortized inference procedure that can rapidly adapt to new spatial patterns from only a sparse set of sensor readings, making them especially suited for data-efficient, active learning tasks. Crucially, unlike deterministic models, NPs provide coherent predictive variances that reflect both model and data uncertainty, enabling principled sensor selection via uncertainty-driven acquisition functions. This rich combination of accuracy and uncertainty quantification establishes NP as an ideal surrogate for our placement task. Our methodology therefore begins with a rigorous empirical benchmark, comparing several NP variants against classical and deterministic baselines to identify the most robust model for our active learning pipeline.

Surrogate Model Benchmarking

To this end, we conducted a comprehensive benchmark of diverse modeling approaches—including classical baselines like Gaussian Process, deterministic regressors, and multiple variants of Neural Processes (NPs) —on the WUSTL $PM_{2.5}$ dataset (Shen et al. 2024). All models were evaluated on their ability to predict held-out $PM_{2.5}$ concentrations over a strictly chronological validation set, a standard practice for mitigating information leakage in spatiotemporal modeling (Andersson et al. 2024). We used Negative Log-Likelihood (NLL) to assess uncertainty calibration and Root Mean Square Error (RMSE) for predictive accuracy. Among all contenders, the Diagonal Transformer Neural Process (TNP-D) (Nguyen and Grover 2023) delivered the strongest performance on the validation set. Its ability to combine strong predictive accuracy with well-calibrated uncertainty estimates makes it a highly robust candidate. Consequently, on the basis of its superior validation metrics, we adopt TNP-D as the surrogate model for all subsequent experiments. The complete benchmark results, model details, and implementation specifics are provided in .

Sensor Placement Acquisitions

In this section, we formally define the employed acquisitions for sensor placements. We begin by introducing standard baseline strategies used for comparison, followed by our proposed gradient-based method for jointly optimizing sensor placements.

Maximum Variance (MaxVar) This is a greedy, uncertainty-driven heuristic. In each step, it selects the single location from the pool of candidates (X_{pool}) with the highest predictive variance, conditioned on the set of already placed sensors (X_c). The selected location is:

$$x_{new}^{(i)} = \arg\max_{x \in X_{pool}} \text{Var}(y|x, X_c, Y_c) \tag{1}$$

This location is then added to X_c, and the process is repeated.

Random Placement For this baseline, new sensor locations are selected by sampling uniformly at random and without replacement from the pool of valid candidate locations, X_{pool}. This strategy serves as a lower-bound on performance and helps quantify the gains achieved by more intelligent placement methods.

Mutual Information (MI). A principled strategy for active sensor placement is to select the location that maximizes the expected reduction in predictive uncertainty over the target region. This corresponds to maximizing the Mutual Information (MI) between the new observation and the unobserved target outputs. Formally, the new sensor locations are selected by solving:

$$\begin{aligned} X^*_{\text{new}} = \arg\max_{X_{\text{new}} \subset X_{\text{pool}}} \; & \mathbb{H}[Y_t \mid X_c, Y_c] \\ & - \mathbb{H}[Y_t \mid X_c, Y_c, X_{\text{new}}, \hat{Y}_{\text{new}}] \end{aligned} \tag{2}$$

where $\mathbb{H}(\cdot)$ denotes entropy, and $\hat{Y}_{\text{new}}$ refers to the predicted measurements at candidate points (used for a look-ahead estimate).

While theoretically appealing, computing this objective for every candidate subset X_{new} is combinatorial and therefore computationally infeasible. In practice, this objective is approximated with "greedy MI" which greedily selects the next candidate until k candidates are selected. Owing to the submodularity of mutual information, this greedy strategy is guaranteed to achieve at least a $(1 - 1/e)$-approximation (approx 63%) of the optimal solution (Krause, Singh, and Guestrin 2008). Nevertheless, even this greedy approach remains computationally intensive, especially in large-scale domains such as all of India. Hence, we evaluate the greedy MI only in a focused regional case study and treat it as a strong but computationally expensive baseline.

Gradient Descent Mutual Information (GD-MI)

Our proposed placement strategy directly optimizes the coordinates of a batch of k new sensors, $X_{\text{new}} = \{x_1, \ldots, x_k\}$, with the objective of approximating the mutual information criterion introduced in Eq. 2. This approach treats the sensor locations themselves as trainable parameters within an optimization loop. The acquisition function is based on a

look-ahead procedure that estimates the expected reduction in future uncertainty. Since the true ground-truth observations at the proposed new locations are unknown, we must approximate the expected information gain. We do this by first imputing the likely sensor readings at those locations. These imputed values, $\hat{Y}_{\text{new}}$, are the model's mean predictions conditioned on the currently existing sensors (X_c, Y_c):

$$\hat{Y}_{\text{new}} = \mathbb{E}[Y|X_{\text{new}}, X_c, Y_c] \tag{3}$$

This use of the predictive mean is a standard technique in active learning, serving as a computationally tractable approximation for the full integral over the predictive distribution. It allows us to simulate the state of our knowledge after a new sensor is placed by using the model's own predictive mean as an estimate of what that sensor will observe. We then form a hypothetical, augmented context set,

$$C_{\text{aug}} = (X_c \cup X_{\text{new}}, Y_c \cup \hat{Y}_{\text{new}})$$

The primary objective, $\mathcal{L}_{\text{var}}$, is to minimize the expected predictive variance across the entire target set, X_t, conditioned on this augmented context:

$$\mathcal{L}_{\text{var}}(X_{\text{new}}) = \mathbb{E}_{x_t \in X_t}[\text{Var}(y_t|C_{\text{aug}})] \tag{4}$$

A key challenge in unconstrained optimization is that the sensor coordinates may drift into invalid locations (e.g., oceans). To address this, we introduce a differentiable out-of-region penalty, $\mathcal{L}_{\text{OOR}}$, which grows exponentially as a point moves away from the valid land grid.

$$\mathcal{L}_{\text{OOR}}(X_{\text{new}}) = \sum_{i=1}^{K} \Big[\exp\Big(\text{ReLU}\Big(\min_{\mathbf{x} \in X_t} \|\mathbf{x}_{\text{new}}^{(i)} - \mathbf{x}\|_2 - \delta \Big)\Big) - 1 \Big] \tag{5}$$

Here δ is the difference between any neighboring grid points. The final objective is the following:

$$\mathcal{L}(X_{\text{new}}) = \mathcal{L}_{\text{var}}(X_{\text{new}}) + \mathcal{L}_{\text{OOR}}(X_{\text{new}}) \tag{6}$$

The locations in X_{new} are initialized from the pool of locations X_{pool}, and their positions are updated iteratively using the Adam optimizer (Kingma and Ba 2017). This end-to-end, gradient-based approach allows the model to account for complex interaction effects between all k new sensor locations simultaneously. Since the other baseline methods select from a discrete set of grid points, for a fair comparison, the final continuous coordinates found by our method are snapped to the nearest grid point in X_{pool} for evaluation.

Experiments and Results

Dataset and Preprocessing

Following prior work that leverages global reanalysis data for air quality modeling and sensor placement research (Andersson et al. 2024), we use the monthly surface-level $PM_{2.5}$ dataset from Washington University in St. Louis (WUSTL) at a 0.1° spatial resolution. This dataset integrates satellite aerosol optical depth (AOD) retrievals, chemical transport model outputs, and ground-based monitoring observations using Gaussian process regression to produce globally consistent $PM_{2.5}$ estimates. We focus on a regional subset covering the Indian subcontinent, spanning latitudes $[5°, 39°]$ and longitudes $[67°, 99°]$. To eliminate oceanic regions, we apply a land mask and retain only grid points over land. We use data from January 1998 to December 2018 to construct a temporally rich dataset, capturing seasonal and interannual variability while avoiding distortions from the COVID-19 pandemic period.

To ensure a realistic and temporally coherent evaluation setting, we partition the dataset chronologically: the training set spans January 1998 to December 2008, the validation set includes January 2009 to December 2010, and the test set covers January 2011 to December 2018. This forward-in-time split mitigates information leakage and better reflects the deployment scenario of learning from historical data to inform future predictions. For preprocessing, raw $PM_{2.5}$ concentrations are log-transformed to reduce skewness and then standardized to have zero mean and unit variance. This transformation stabilizes model training and ensures consistency across different spatial and temporal contexts. The resulting processed dataset serves as input to all neural process models evaluated in this study. Further implementation details, including masking and interpolation routines, are provided in supplementary material.

Surrogate model Evaluation and Selection

As established in our methodology, the foundational step of our approach is the selection of a robust surrogate model. This section provides the detailed experimental setup and comparative results from the comprehensive benchmark used to make this selection.

Training and Evaluation. We evaluated the models on two primary metrics: Root Mean Square Error (RMSE) for predictive accuracy and Negative Log-Likelihood (NLL) for uncertainty calibration. All neural models were implemented in PyTorch and trained with early stopping based on validation NLL. Following the meta-learning paradigm (Finn, Abbeel, and Levine 2017), each training iteration involved sampling a random monthly snapshot from the training set and constructing a new task by randomly selecting 700 context points and 5,000 target points. The final performance was evaluated on the held-out validation and test sets. Further implementation details such as optimizer and hardware are provided in the supplementary material.

Results. The complete performance metrics for all benchmarked models are presented in Table 1 (Validation) and Table 2 (Test). The validation results reveal a clear performance hierarchy. Purely deterministic methods like Random Forest serve as strong RMSE baselines. The classical Gaussian Process (GP) performs well but suffers from cubic scaling complexity that limits its practicality for large-scale applications (Titsias 2009). Among the Neural Process family, performance scales with architectural complexity: the basic Conditional NP (CNP) (Garnelo et al. 2018a) is significantly outperformed by attentive (CANP) (Kim et al. 2019) and convolutional variants. Notably, both the ConvCNP (Gordon et al. 2020) and the ConvGNP (Andersson et al. 2024)—which models a full covariance—show strong performance,

Model	NLL (↓)	RMSE (↓)
CNP	0.48 ± 0.00	11.46 ± 0.09
Random Forest	-0.11 ± 0.00	6.55 ± 0.08
Gaussian Process	-0.19 ± 0.00	5.16 ± 0.08
CANP	-0.19 ± 0.00	6.15 ± 0.05
ConvCNP	-0.27 ± 0.00	5.28 ± 0.07
ConvGNP	-0.30 ± 0.00	5.31 ± 0.06
TabPFN	-0.37 ± 0.00	5.09 ± 0.05
TNP-D	**-0.44 ± 0.00**	**4.90 ± 0.04**

Table 1: Performance of various models on the validation dataset spanning over 2009 and 2010. NLL = Negative Log Likelihood (lower is better) and RMSE = Root Mean Squared Error (lower is better)

Model	NLL (↓)	RMSE (↓)
CNP	0.98	24.99
Random Forest	2.43	15.37
CANP	0.82	14.48
ConvGNP	1.04	12.25
TabPFN	0.39	11.23
ConvCNP	0.59	10.78
Gaussian Process	0.42	10.31
TNP-D	0.51	9.11

Table 2: Performance of various models on the test dataset. NLL = Negative Log Likelihood (lower is better) and RMSE = Root Mean Squared Error (lower is better)

demonstrating the power of convolutional architectures for spatial data. However, the Transformer-based models, including the pre-trained TabPFN (Hollmann et al. 2023) and our task-trained TNP-D (Nguyen and Grover 2023), are the decisive top performers. The TNP-D achieves the best overall result on the validation set, securing a 4% lower RMSE (4.90 vs. 5.16 μg/m^3) and a 57% lower NLL (-0.44 vs. -0.19) than its closest non-pre-trained competitor, the GP. TabPFN, used here in a zero-shot capacity without any fine-tuning, also performs exceptionally well. The TNP-D's architecture, which assumes conditional independence between target points for faster, parallelized inference, is a highly effective modeling choice for this spatiotemporal task. Based on its strong validation performance and scalability, we adopt TNP-D as the surrogate model for all subsequent active learning experiments. Since sensor selection relies heavily on the quality of uncertainty estimates from the surrogate, TNP-D's superior calibration, as demonstrated by its low NLL, is particularly crucial for our active learning pipeline.

Case Study: Benchmarking Against a Strong Heuristic

To benchmark our method's accuracy against the powerful but computationally expensive Mutual Information (MI) heuristic, we conducted a focused case study. This experiment was performed on the Madhya Pradesh region of India, a geographical scope where the MI approach was computationally feasible. For this comparison, we used the existing Central Pollution Control Board (CPCB) sensor locations as the initial context set (X_c) (more details in supplementary material). We compare our GD-MI approach against MI, along with the MaxVar and Random baselines. The results, shown in Figure 1, demonstrate the effectiveness of our joint optimization approach. As expected given its theoretical grounding, the greedy MI heuristic achieves the lowest predictive error (RMSE of 5.6 at k=9), establishing a clear performance benchmark. Crucially, our GD-MI approach is highly competitive. For a budget of k=9 sensors, it achieves an RMSE of 5.8, closing most of the performance gap with the MI heuristic while consistently outperforming the simpler MaxVar method (RMSE of 6.3). These findings suggest that our approach achieves accuracy approaching that of the powerful greedy method. This result is significant because it shows that the substantial gains in computational efficiency, which we demonstrate in the following section, do not come at a major cost to deployment accuracy, making our method a practical choice for real-world applications.

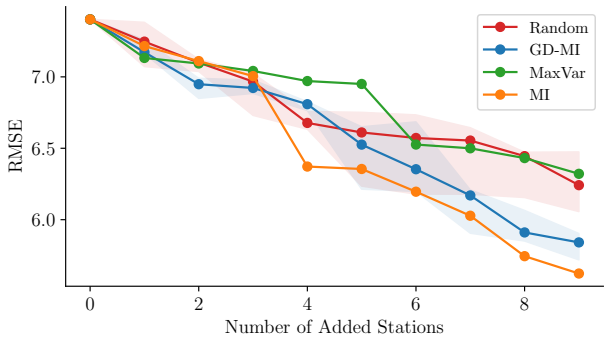

Figure 1: Case study comparing placement effectiveness on a regional subset. The plot shows the RMSE on the validation set as a function of the number of new stations deployed. Our proposed GD-MI method is compared against the computationally expensive Mutual Information (MI), the greedy MaxVar method, and a Random baseline. While MI achieves the lowest error, our GD-MI method is highly competitive and consistently outperforms the other baselines. The shaded regions for the stochastic methods (GD-MI and Random) represent the interquartile range (25th to 75th percentile) over 5 independent runs.

Sensor Placement Results

We now present the results of our primary experiment: a full active learning pipeline deployed across the entire Indian subcontinent, starting from the existing Central Pollution Control Board (CPCB) sensor locations as the initial context set (X_c). This large-scale evaluation assesses performance in a realistic, high-dimensional setting. We compare our GD-MI method against the widely-used MaxVar heuristic and a Random baseline over 5 independent runs (with fixed random seeds to ensure reproducibility). The test set RMSE, as a function of newly deployed sensors, is shown in

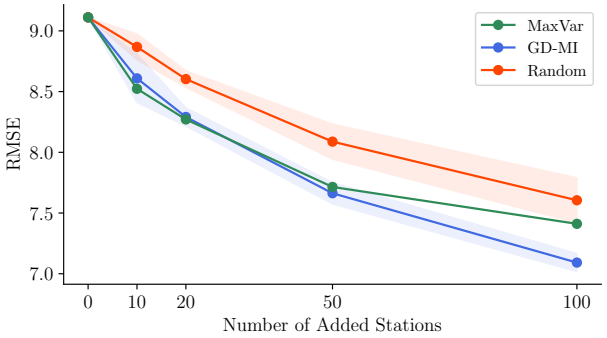

Figure 2: Main active learning results on the test set. The plot shows the RMSE as a function of the number of new stations added to the existing CPCB network across India. Our GD-MI method is compared against the greedy MaxVar and a Random baseline. GD-MI consistently outperforms the baselines, with the performance gap widening as the sensor budget increases. Shaded regions represent the full min-max range of results over 5 random initializations.

Figure 2. The results demonstrate a clear performance hierarchy. Our GD-MI method yields a lower mean RMSE than both the greedy MaxVar heuristic and the Random baseline at nearly all budget sizes. The performance gain of GD-MI grows with the budget size, suggesting its joint optimization benefits become more pronounced in richer deployment regimes where avoiding spatial redundancy is critical. At a final deployment of 100 new stations, GD-MI achieves an RMSE of approximately 7.1, a 4% relative reduction over MaxVar's 7.4. The non-overlapping error bars at this point suggest this improvement is consistent. These results highlight the practical utility of our approach for large-scale environmental monitoring tasks, demonstrating that joint optimization leads to a more efficient network design and lower overall predictive error than the widely-used MaxVar heuristic.

Qualitative Analysis

To provide an intuitive understanding of placement behavior, we visualize the locations selected by the MaxVar heuristic and our GD-MI method for a representative time slice from February 1, 2012 (Figure 3). The analysis reveals clear differences in the resulting network designs. The greedy MaxVar strategy exhibits two suboptimal patterns. First, it creates dense spatial clusters by aggressively targeting high-error regions. Second, it frequently places sensors along the geographical boundaries, reducing effective inland coverage. This combined behavior results in a redundant and poorly distributed network. In contrast, our GD-MI method produces a more spatially balanced and efficient network. The joint optimization process naturally avoids both clustering and boundary-placement issues. It also learns to place sensors away from existing CPCB stations, targeting unexplored regions to maximize new information. The resulting placements provide superior geographical coverage, con-

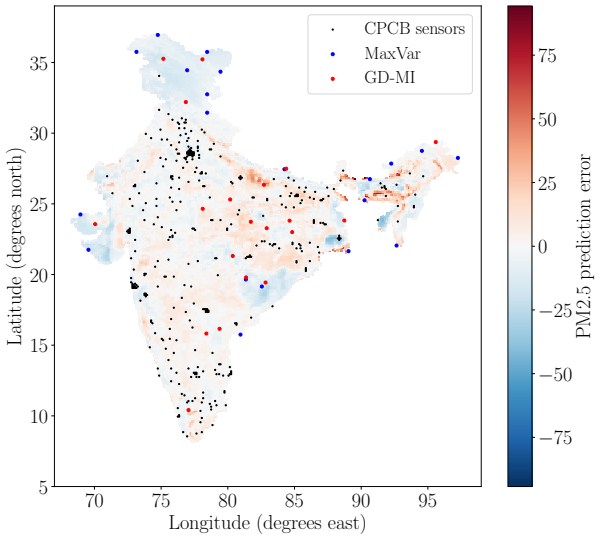

Figure 3: Qualitative comparison of sensor placements on 2012-02-01. Red markers show sensors selected by GD-MI; blue markers show those selected by MaxVar. Black dots indicate existing CPCB stations. GD-MI yields better spatial coverage, avoiding clustering and boundary bias common in MaxVar placements.

firming that joint optimization leads to a more effective and informative network design than a simple greedy approach. Placements for other budget sizes exhibit similar patterns and are included in supplementary material.

Discussion

Our results demonstrate that reframing sensor placement as a continuous, gradient-based optimization problem can effectively overcome the longstanding trade-off between accuracy and scalability. By jointly optimizing sensor coordinates and decoupling runtime from the number of candidate sites, our approach provides a practical solution for large-scale air quality monitoring. The qualitative analysis further shows that our method produces more spatially balanced sensor configurations than greedy heuristics, mitigating issues like clustering and boundary redundancy.

However, several limitations open avenues for future research. A primary challenge is the sensitivity of gradient descent to initialization, which can lead to convergence at suboptimal local minima. This could be mitigated by running multiple optimizations with different seeds or by initializing from a fast heuristic placement. Additionally, our current objective function minimizes predictive variance—a useful statistical proxy for information gain—but does not explicitly consider equity or societal impact. Incorporating fairness-aware criteria, such as prioritizing underserved or high-risk populations, is an important extension. Finally, our method currently optimizes for a fixed snapshot in time. Extending it to a spatio-temporal setting, where sensors are dynamically scheduled across space and time, is a promising direction for

increasing long-term robustness.

Despite these limitations, our work provides a strong foundation for a new class of fast, differentiable sensor placement algorithms. By making high-quality monitoring network design computationally feasible at national and continental scales, this approach opens up new possibilities for scalable, data-driven environmental policy and public health planning.

Conclusion

We address the trade-off between accuracy and scalability in optimal sensor placement for air quality monitoring. Our method reframes the discrete placement task as a continuous optimization problem. By directly optimizing sensor coordinates, we overcome the limitations of greedy heuristics. Experiments on a continental-scale $PM_{2.5}$ reanalysis dataset show that our approach achieves high placement quality while being orders of magnitude faster than traditional information-theoretic baselines. Unlike methods that scale with the number of candidate sites, our runtime remains constant, enabling efficient large-scale deployment. This approach offers a practical tool for AI-driven social impact, helping policymakers design more informative and cost-effective monitoring networks. Future extensions could incorporate fairness-aware or multi-objective criteria and explore joint spatiotemporal optimization.

Acknowledgments

This work was supported by funding from GUJCOST under grant RES/GCOST/CS/P0275/2324/0009, which enabled the core research activities reported here. Additional support came from the Microsoft Research India PhD Fellowship under grant RES/MRLI/CS/P0275/2425/0015, which sustained the broader doctoral research effort connected to this study.

References

Andersson, T. R.; Bruinsma, W. P.; Markou, S.; Requeima, J.; Coca-Castro, A.; Vaughan, A.; Ellis, A.-L.; Lazzara, M. A.; Jones, D.; Hosking, J. S.; and Turner, R. E. 2024. Environmental Sensor Placement with Convolutional Gaussian Neural Processes. *arXiv preprint arXiv:2402.11824*.

Central Pollution Control Board. 2025. National Air Quality Data and Reports. The official count of monitoring stations is available through portals on this main government website.

Chang-Silva, R.; Tariq, S.; Kim, S.; Moosazadeh, M.; Park, S.; and Yoo, C. 2024. Satellite-informed smart sensor placement framework for near-optimal PM2.5 monitoring in urban areas. *Environmental Science and Pollution Research International*. Epub ahead of print.

Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. arXiv:1703.03400.

Garnelo, M.; Rosenbaum, D.; Maddison, C. J.; Ramalho, T.; Saxton, D.; Shanahan, M.; Teh, Y. W.; Rezende, D. J.; and Eslami, S. M. A. 2018a. Conditional Neural Processes. arXiv:1807.01613.

Garnelo, M.; Schwarz, J.; Rosenbaum, D.; Viola, F.; Rezende, D. J.; Eslami, S. M. A.; and Teh, Y. W. 2018b. Neural Processes. In *ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*.

Gordon, J.; Bruinsma, W. P.; Foong, A. Y. K.; Requeima, J.; Dubois, Y.; and Turner, R. E. 2020. Convolutional Conditional Neural Processes. arXiv:1910.13556.

Guestrin, C.; Krause, A.; and Singh, A. P. 2005. Near-optimal sensor placements in Gaussian processes. In *ICML*, 265–272.

Hollmann, N.; Müller, S.; Eggensperger, K.; and Hutter, F. 2023. TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. arXiv:2207.01848.

Holstein, K.; Wortman Vaughan, J.; Daumé, H.; Dudik, M.; and Wallach, H. 2019. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need? In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, 1–16. ACM.

Kim, H.; Mnih, A.; Schwarz, J.; Garnelo, M.; Eslami, A.; Rosenbaum, D.; Vinyals, O.; and Teh, Y. W. 2019. Attentive Neural Processes. arXiv:1901.05761.

Kingma, D. P.; and Ba, J. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980.

Krause, A.; Singh, A.; and Guestrin, C. 2008. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *J. Mach. Learn. Res.*, 9: 235–284.

Nguyen, T.; and Grover, A. 2023. Transformer Neural Processes: Uncertainty-Aware Meta Learning Via Sequence Modeling. arXiv:2207.04179.

Settles, B. 2009. Active learning literature survey.

Shen, S.; Li, C.; van Donkelaar, A.; Jacobs, N.; Wang, C.; and Martin, R. V. 2024. Enhancing Global Estimation of Fine Particulate Matter Concentrations by Including Geophysical a Priori Information in Deep Learning. *ACS ES&T Air*, 1(5): 332–345.

Sun, Z.; Mahmoodian, M.; Sidiq, A.; Jayasinghe, S.; Shahrivar, F.; and Setunge, S. 2025. Optimal Sensor Placement for Structural Health Monitoring: A Comprehensive Review. *Journal of Sensor and Actuator Networks*, 14(2).

Titsias, M. 2009. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial intelligence and statistics*, 567–574. PMLR.

van Donkelaar, A.; Martin, R. V.; Brauer, M.; Hsu, N. C.; Kahn, R. A.; Levy, R. C.; Lyapustin, A.; Sayer, A. M.; and Winker, D. M. 2016. Global Estimates of Fine Particulate Matter using a Combined Geophysical-Statistical Method with Information from Satellites, Models, and Monitors. *Environmental Science & Technology*, 50(7): 3762–3772. PMID: 26953851.

Williams, C. K.; and Rasmussen, C. E. 2006. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.

World Health Organization. 2024. Ambient (outdoor) air quality and health. Accessed: 2025-07-31.

Xing, J.; Zheng, S.; Ding, D.; Kelly, J. T.; Wang, S.; Li, S.; Qin, T.; Ma, M.; Dong, Z.; Jang, C.; Zhu, Y.; Zheng, H.; Ren, L.; Liu, T.; and Hao, J. 2020. Deep Learning for Prediction of the Air Quality Response to Emission Changes. *Environmental Science & Technology*, 54(14): 8589–8600.