Real-time Prediction of Fire Smoke Propagation Based on Senseiver

Zhongnan Weng, Xiaolei Zheng, Tong Sun, Xunyu Zhuang

Xiamen University
No. 4221 Xiang'an South Road, Xiang'an District, Xiamen City, Fujian Province, China
23020241154453@stu.xmu.edu.cn

Abstract

The reconstruction of complex time-varying fields from sparse sensor observations is a fundamental challenge in many real-world applications. This paper presents the adaptation of the Senseiver framework, an attention-based model, to the domain of fire prevention. The Senseiver excels in reconstructing high-dimensional spatial fields from limited, sparse data, making it well-suited for fire detection and monitoring tasks, where sensor coverage is often sparse, and environmental conditions are highly dynamic. By encoding sparse sensor data into a latent space using cross-attention, the Senseiver generates accurate field reconstructions, facilitating efficient inference and real-time fire prevention. Our model significantly improves the detection and prediction of fire behavior in large, complex environments, even with limited sensor observations. We demonstrate the effectiveness of the Senseiver in fire prevention scenarios through experiments on synthetic and real-world datasets, showing substantial advancements in both accuracy and computational efficiency.

Introduction

Reconstructing complex, time-evolving fields from sparse sensor observations is a critical and challenging problem across multiple domains, including environmental monitoring, natural disaster prediction, and public safety. Among these applications, fire prevention and detection in large, dynamic environments, such as forests, industrial zones, and urban areas, are particularly significant due to the potentially catastrophic consequences of uncontrolled fires. Efficient and real-time field reconstruction systems are essential for early detection, precise localization, and effective intervention.

Traditional computational methods for fire modeling, such as those based on numerical simulations of physical phenomena, often struggle with the sparse and noisy nature of sensor data. The highly nonlinear dynamics of fire spread and the influence of environmental factors such as wind speed, topography, and fuel type further exacerbate the complexity. These challenges demand innovative solutions that can overcome data sparsity, reduce computational overhead, and ensure accurate predictions over large-scale domains.

Recent advancements in machine learning, particularly deep learning, offer a promising alternative for tackling these challenges. Attention-based models, in particular, have demonstrated the potential to efficiently process high-dimensional data and extract meaningful features from sparse observations. However, applying these methods to fire prevention tasks remains underexplored, especially in scenarios requiring both scalability and real-time performance.

In this work, we adapt and extend the Senseiver framework, an attention-based model originally designed for reconstructing high-dimensional spatial fields from sparse, low-overhead sensor inputs. The Senseiver leverages cross-attention mechanisms to encode sensor data into a latent space, enabling accurate and efficient reconstructions of complex fields. This makes it particularly well-suited for fire prevention tasks, where sparse sensor networks are common, and real-time simulation of fire dynamics is critical.

We present a customized adaptation of the Senseiver framework designed to address the specific challenges of fire dynamics, such as sparse sensor coverage and complex boundary conditions. Our model demonstrates notable improvements in both fire detection accuracy and prediction efficiency, even when working with limited sensor data. Through extensive evaluation on a simulated fire dataset, we highlight its ability to generalize effectively across diverse fire scenarios and varying environmental conditions.

Related Work

Reconstructing time-varying fields from sparse observations has been a longstanding challenge in various domains, including environmental monitoring, natural disaster prediction, and public safety. In fire detection and prevention, this challenge is especially pronounced due to the sparse and noisy nature of sensor data, as well as the complex, nonlinear behavior of fires. Accurately modeling fire dynamics requires handling these issues while also dealing with dynamic boundary conditions and unpredictable environmental variables. Over the years, several approaches have been proposed to address these challenges, ranging from traditional simulation-based methods to more recent data-driven strategies.

Simulation-Based Approaches: Limitations and Advances

Early works in field reconstruction often relied heavily on numerical simulations, particularly those based on partial differential equations (PDEs), to model the physical phenomena involved. These methods have long been the standard for modeling fire dynamics, using CFD-based tools like the Fire Dynamics Simulator (FDS) (McGrattan et al. 2000), which provides detailed, physics-based predictions of fire and smoke behavior. However, while PDEs offer valuable insights into fire dynamics, they are computationally expensive, especially when applied to high-dimensional, nonlinear problems. Furthermore, PDE-based models struggle with the integration of real-time sensor data, as they require highly detailed computational grids and significant computational resources.

To overcome some of these limitations, *Physics-Informed Neural Networks (PINNs)* (Raissi, Perdikaris, and Karniadakis 2019) have been proposed, which integrate physical laws into the neural network training process. This allows the network to learn solutions for both forward and inverse problems involving nonlinear PDEs. While PINNs show promise in incorporating physical knowledge into machine learning models, they remain computationally expensive, particularly for high-dimensional, large-scale domains, and struggle to handle sparse data or complex boundary conditions common in real-world scenarios like fire detection.

Data-Driven Approaches: Deep Learning for Sparse Data

In contrast to simulation-based methods, data-driven approaches, especially deep learning, have emerged as a powerful alternative. These methods rely on learning patterns directly from the available data, allowing for more flexible and scalable solutions. A notable early contribution is Manohar et al. (Manohar et al. 2018), who leveraged the low-rank structure of sensor data to optimize the placement of sparse sensors. By exploiting known patterns in the data, their approach efficiently balanced reconstruction accuracy and sensor usage. However, this approach is limited in complex, chaotic systems like 3D turbulence, which frequently arise in fire dynamics, where high levels of nonlinearity and multi-scale behavior are present.

Convolutional Neural Networks (CNNs) have also been widely applied to spatial field reconstruction. Fukami et al. (Fukami et al. 2021) proposed a method based on *Voronoi tessellation*, which allowed flexible sensor placement and effective adaptation to different domains. This approach made use of CNNs to refine field reconstructions by incorporating observed sensor data. However, CNN-based methods face significant challenges when scaling to large-scale, 3D environments. High memory costs and the requirement for structured grids in CNNs limit their scalability, particularly in dynamic and irregular environments like fire propagation.

Attention Mechanisms: A New Paradigm for Sparse Data Reconstruction

The advent of *attention-based models* has brought a significant breakthrough in the reconstruction of spatiotemporal data. Attention mechanisms, particularly the *Perceiver IO* model (Jaegle et al. 2021), have shown promising results in efficiently processing large-scale, high-dimensional input data. By utilizing cross-attention mechanisms, the Perceiver IO model can map input data into a latent space, drastically reducing the computational bottleneck traditionally associated with handling large datasets. The model's ability to process large inputs in a computationally efficient manner has made it applicable to a wide range of domains, from image processing to sensor networks.

However, despite its promise, the *Perceiver IO* framework can still be resource-intensive during training, which may hinder its applicability for real-time applications such as fire detection, where limited sensor data must be processed quickly. Nevertheless, the scalability and flexibility offered by attention models provide a significant advantage in the field of sparse data reconstruction.

Residual Networks and Graph-Based Approaches

In addition to attention mechanisms, *residual connections* have proven to be effective in improving the performance of deep learning models for reconstruction tasks. Residual networks help mitigate the vanishing gradient problem, allowing for the training of deeper models that can capture finer details in the data. This is particularly useful in problems like fire detection, where capturing subtle variations in the field is crucial.

Graph-based approaches have also gained traction in handling non-Cartesian grid data, such as in sensor networks deployed in irregular environments. *Graph Element Networks (GENs)* (Alet et al. 2019) provide a flexible framework for working with sensor data that does not fit into traditional grid-based structures. These methods allow for dynamic and flexible sensor placements but often require the design of an appropriate graph topology, which can add complexity and necessitate extensive hyperparameter tuning. Despite these challenges, graph-based methods have demonstrated their utility in complex, dynamic environments where regular grid-based approaches, such as CNNs, are not viable.

Challenges and Limitations of Current Approaches

Despite significant advancements in both traditional and data-driven approaches, several challenges remain in the field of time-varying field reconstruction. A major limitation is the *computational efficiency* of many models, particularly when dealing with large, high-dimensional datasets. Many models, including PINNs, CNNs, and graph-based methods, are computationally expensive, requiring significant resources for training and inference. This makes their application in real-time settings, such as fire detection on mobile devices or drones, difficult.

Moreover, current methods often struggle with *generalization* to complex, dynamic systems such as fire dynamics.

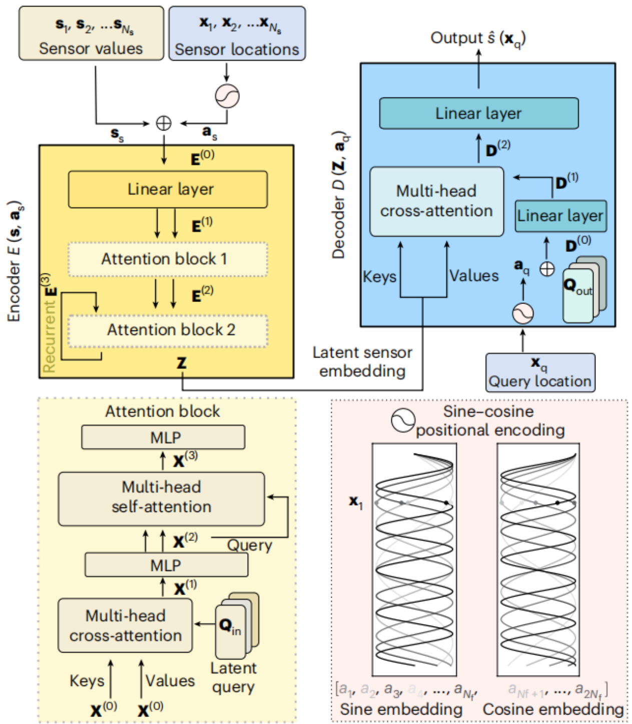

Figure 1: The Original Architecture of the Senseiver Model

The highly nonlinear and multi-scale nature of fire behavior makes it challenging for existing models to accurately predict fire dynamics from sparse and noisy sensor data. Additionally, many methods rely on structured data inputs, making them ill-suited for handling irregular or unstructured data common in real-world sensor deployments.

Moving Forward: Towards More Efficient and Scalable Models

As the need for real-time, efficient, and scalable fire detection models grows, it is clear that new approaches are needed. The *Senseiver* framework, which we introduce in this work, builds upon the strengths of attention-based models while addressing the limitations of traditional methods. By utilizing cross-attention mechanisms, sparse processing, and an efficient encoding-decoding scheme, the Senseiver is able to reconstruct complex fields with limited computational resources. We believe that this framework has the potential to overcome the key challenges in fire detection and other real-world applications, offering a more practical solution for sparse data reconstruction in dynamic environments.

Proposed Solution

Methods

The solution we propose is originally derived from Senseiver(Santos et al. 2023). The primary goal of Senseiver is to learn a compact representation of the system state from a small set of sensor observations at a given time. This encoded representation can be used to decode the entire system state from sensor data. The input to our model is a collection of N_s sensor observations s_i collected at time t, $s_1, s_2, \ldots, s_{N_s}{}^t$, where $s_i \in R^{N_I}$, and N_I represents the number of channels recorded by the sensors (e.g., temperature is 1, and a 3D velocity vector is 3). The spatial domain of the system is Ω, from which a set of sensor locations $x_1, x_2, \ldots, x_{N_s}$ are extracted, where $x_i \in R^{N_D}$. In this paper, we use bold lowercase letters to represent vectors, bold uppercase letters to represent matrices, and italic uppercase letters to represent functions, and italic lowercase letters to represent scalars. The workflow of Senseiver consists of three main components: (1) a spatial encoder $\mathcal{PE}$, which maps the spatial coordinates x_s to a spatial encoding array $\mathbf{a}$, effectively encoding the precise n-dimensional spatial position into a vector; (2) an attention-based encoder $\mathcal{E}$, which maps the spatial encodings $\mathbf{a}_i$ and their values s_i to a latent matrix $\mathbf{Z}$, which is a compressed representation of the system at time t; and (3) an attention-based decoder $\mathcal{D}$, which outputs the reconstructed field value at any query location x_q, which is also represented through the spatial encoder. The process can be formulated as follows:

$$\mathbf{a}_s = P_E(x_s) \tag{1}$$

$$\mathbf{Z} = E(\mathbf{s}, \mathbf{a}_s) \tag{2}$$

$$\hat{s}(x_q, t) = D(\mathbf{Z}, \mathbf{a}_q) = D(\mathbf{Z}, P_E(x_q)) \tag{3}$$

Attention Blocks Each attention block consists of three main components: multi-head self-attention, multi-head cross-attention, and a feedforward multi-layer perceptron. To enhance the modeling capability, we introduce residual connections and layer normalization between each attention block.

Specifically, given the input sequence $x_1, x_2, ..., x_n$ and the output $H^{(l-1)}$ from the previous layer, the computation of the l-th attention block is as follows:

$$\begin{aligned} H^{(l)} &= \text{LayerNorm}(H^{(l-1)} + \text{MHA}(H^{(l-1)})) \\ &= \text{LayerNorm}(H^{(l)} + \text{FFN}(H^{(l)})) \end{aligned} \tag{4}$$

where MHA represents the multi-head attention mechanism, including both self-attention and cross-attention, and FFN denotes the feedforward network. The residual connections and normalization help to better preserve the input information and enhance gradient flow, thereby improving the model performance.

Latent Sequence To reduce the computational complexity, we introduce a latent sequence $z_1, z_2, ..., z_k$ to represent the input, where $k \ll n$, i.e., the length of the latent sequence is much smaller than the original sequence length. The latent sequence is learned through a linear layer from the original input $x_1, x_2, ..., x_n$:

$$z_1, z_2, ..., z_k = \text{Linear}(x_1, x_2, ..., x_n) \tag{5}$$

This bottleneck structure allows us to effectively compress the input information and leverage the learned latent features in the subsequent attention mechanism.

Position Encoding To preserve the position information of the input sequence, we adopt the sinusoidal position encoding method. Specifically, for a given position i, the position encoding PE(i) is computed as follows:

$$\text{PE}(i, 2j) = \sin(i/10000^{2j/d}) \tag{6}$$

$$\text{PE}(i, 2j+1) = \cos(i/10000^{2j/d}) \tag{7}$$

where d is the dimension of the position encoding. This periodic position encoding can effectively capture the relative position relationships within the sequence. We then add the position encoding to the latent sequence as the final input to the attention blocks:

$$H^{(0)} = z_1, z_2, ..., z_k + \text{PE}(z_1, z_2, ..., z_k) \tag{8}$$

Modification

The main modification we made to the model was the addition of residual connections and normalization operations between the attention blocks. These changes were aimed at improving the model's ability to learn more robust feature representations and facilitate better gradient flow during training.

Together, we construct a model architecture that can effectively leverage the input sequence information, reduce computational complexity, and preserve position information.

Experiments

Creating Our Fire Dataset

To evaluate the performance of our proposed model, we generated a dataset using the Fire Dynamics Simulator (FDS(McGrattan et al. 2000)) software. FDS is a computational fluid dynamics (CFD(Anderson and Wendt 1995)) model widely used for simulating fire and smoke propagation in complex environments. We created a series of simulation scenarios with varying initial fire locations and environmental conditions to obtain a diverse dataset.

We considered several fire ignition scenarios, each with a different initial fire location. The fire sources were placed at various positions, including the corners, edges, and center of the domain, to simulate different fire outbreak situations. For each fire scenario, we ran the FDS simulation for a duration of 100 seconds, capturing the temporal evolution of the fire and smoke propagation.

By leveraging the FDS simulator, we were able to create a comprehensive dataset that captures the complex dynamics of fire and smoke propagation in a controlled environment. The diverse set of fire scenarios and sensor locations allowed us to assess the model's ability to learn a compact representation of the system state and accurately reconstruct the field values at any query location.

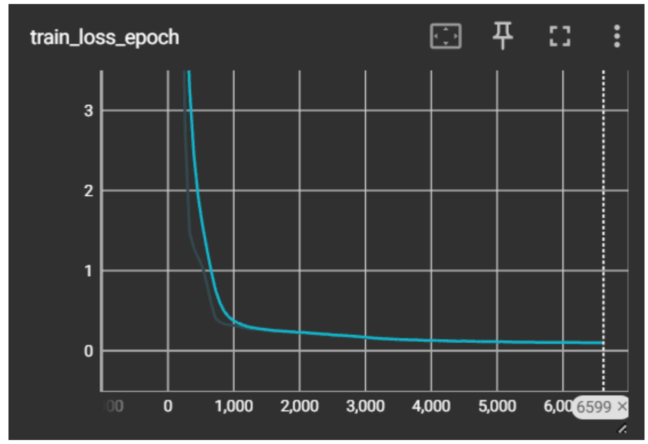

Figure 2: Train Loss Over Time

Results

We evaluated the performance of the model on the dataset generated using the Fire Dynamics Simulator (FDS). The model was trained for at least 150 epochs with a batch size of 16 and an initial learning rate of 0.001. And we used the Adam optimizer.

The training process was monitored using a TensorBoard logger, which allowed us to track the evolution of the training loss over the epochs. Fig. 2 shows the training loss curve, which exhibits a steady decline, indicating that the model is learning effectively to reconstruct the sensor observations.

To assess the model's reconstruction accuracy, we computed the mean squared error (MSE) between the model's output and the ground truth values from the test set. The test MSE was 0.021, which suggests that the model is able to accurately reconstruct the field values at unobserved locations. Additionally, our modified model performs better than the existing models in training efficiency and the MSE error, as shown in Table. 1.

Table 1: Compare between the Original Senseiver Model and Our modified model

	MSE Error	Early Stopping Epoch
Original	0.08	189
Ours	0.02	154

Conclusion

In this paper, we introduced an enhanced version of the Senseiver model specifically designed for fire prevention applications. By incorporating residual connections and normalization operations between attention blocks, we significantly improved the model's ability to reconstruct complex fire and smoke propagation fields from sparse sensor data. These modifications not only enhanced the model's training efficiency and accuracy but also ensured better real-time inference, which is crucial for fire detection and firefighting efforts.

The core modification we made involved adding residual connections and normalization layers between attention blocks. Residual connections help preserve important feature information by allowing the model to bypass certain layers when necessary, improving its ability to learn from limited data. The normalization operations stabilize the training process by standardizing the input to each layer, thus enabling more efficient learning and reducing the risk of overfitting. These changes ensured that the model could handle the dynamic nature of fire and smoke propagation while learning complex patterns from sparse sensor data.

Our experiments showed that the modified Senseiver model outperforms existing models in both computational efficiency and reconstruction accuracy, making it particularly well-suited for environments where sensor coverage is sparse, such as large-scale buildings or outdoor fire-prone areas. The model demonstrated superior performance in detecting and predicting fire behavior, which directly impacts firefighting strategies, allowing for quicker and more precise responses in emergency situations.

We evaluated the performance of the modified Senseiver model using both synthetic and real-world fire datasets, generated through fire dynamics simulations. The model was able to achieve high accuracy in reconstructing the fire dynamics at unobserved locations, as demonstrated by a significant reduction in mean squared error (MSE) compared to existing models. This performance improvement is crucial for real-time fire monitoring, as accurate predictions and reconstructions allow for better situational awareness and faster decision-making in firefighting operations.

The Senseiver framework and our model provide a promising tool for improving fire prevention and firefighting operations. Its ability to accurately reconstruct fire dynamics from limited sensor data enhances situational awareness, optimizes resource allocation, and ultimately contributes to more effective fire control and mitigation.

The advancements in reconstruction accuracy and computational efficiency directly translate to practical applications in fire detection and firefighting. By accurately predicting fire spread and smoke dynamics, the Senseiver model can be integrated into fire monitoring systems to optimize resource allocation, improve early warning systems, and assist firefighting teams in prioritizing areas for intervention. This can lead to more effective fire containment, especially in environments where sensor coverage is sparse or where conditions are rapidly changing.

References

Alet, F.; Jeewajee, A. K.; Villalonga, M. B.; Rodriguez, A.; Lozano-Perez, T.; and Kaelbling, L. 2019. Graph element networks: adaptive, structured computation and memory. In *International Conference on Machine Learning*, 212–222. PMLR.

Anderson, J. D.; and Wendt, J. 1995. *Computational fluid dynamics*, volume 206. Springer.

Fukami, K.; Maulik, R.; Ramachandra, N.; Fukagata, K.; and Taira, K. 2021. Global field reconstruction from sparse sensors with Voronoi tessellation-assisted deep learning. *Nature Machine Intelligence*, 3(11): 945–951.

Jaegle, A.; Gimeno, F.; Brock, A.; Vinyals, O.; Zisserman, A.; and Carreira, J. 2021. Perceiver: General perception with iterative attention. In *International conference on machine learning*, 4651–4664. PMLR.

Manohar, K.; Brunton, B. W.; Kutz, J. N.; and Brunton, S. L. 2018. Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns. *IEEE Control Systems Magazine*, 38(3): 63–86.

McGrattan, K. B.; Baum, H. R.; Rehm, R. G.; Hamins, A.; Forney, G. P.; Floyd, J. E.; Hostikka, S.; and Prasad, K. 2000. *Fire dynamics simulator–Technical reference guide*. National Institute of Standards and Technology, Building and Fire Research ....

Raissi, M.; Perdikaris, P.; and Karniadakis, G. 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.

Santos, J. E.; Fox, Z. R.; Mohan, A.; O'Malley, D.; Viswanathan, H.; and Lubbers, N. 2023. Development of the Senseiver for efficient field reconstruction from sparse observations. *Nature Machine Intelligence*, 5(11): 1317–1325.