A retrieval-based approach for real-time wildfire analysis using LSTM-driven similarity matching

Iliana Malliou[1]; Giorgos Filandrianos*[1]; Angeliki Dimitriou[1]; David Caballero[2]; Luis Torres[2]; Giorgos Stamou[1]

[1]*National Technical University of Athens, 9, Heroon Polytechniou St, 15780 Zografou, Greece*
[2] *Meteogrid. Modesto Lafuente 45, 28010, Madrid, Spain*

**Corresponding author. Email: geofila@ails.ece.ntua.gr*

Keywords

Similarity-Based Retrieval, Spatiotemporal Embeddings, Early-Stage Situational Awareness, Short-Term Fire Projection

Abstract

In the critical early stages of a wildfire, the ability to project potential fire growth is often hindered by the time required to configure and execute high-fidelity physical simulations. This paper presents a data-driven retrieval engine designed to provide immediate decision support by identifying historical or pre-simulated fire events that share characteristics with an active incident. By shifting the focus from real-time computation to intelligent search and retrieval, this framework allows fire managers to access relevant past scenarios and their outcomes within seconds.

The core of this system is a similarity-matching architecture that utilizes Long Short-Term Memory (LSTM) networks to process the temporal and spatial evolution of fire perimeters. Unlike traditional static overlays, the LSTM components encode the trajectory of a fire's growth, capturing the influence of topography and weather patterns on fire behavior over time. This enables the representation of complex fire spread dynamics as high-dimensional vectors, which can be queried against an extensive database of previously recorded or simulated wildfire events.

The retrieval process identifies "neighboring" historical cases where environmental conditions - such as fuel moisture, wind vectors, and slope - align with current observations. By surfacing these analogous events, the system provides a probabilistic indication of how the current fire may evolve, based on historical outcomes. This approach bypasses the "real-time barrier" associated with complex simulations by leveraging an existing knowledge base, ensuring timely insights during the most critical phases of an incident.

This research demonstrates that integrating recurrent neural networks for feature extraction significantly improves the precision of similarity-based retrieval compared to standard geometric matching. The results suggest that such a framework can effectively complement traditional forecasting methods, offering a rapid, data-driven tool for wildfire situational awareness and resource planning.

1. Introduction

Wildfires are among the most destructive natural hazards, with severe consequences for ecosystems, infrastructure, and human life. Their impact is particularly critical in wildland–urban interface environments, where rapid fire growth can directly threaten settlements and emergency response operations. In such settings, timely situational awareness is essential, especially during the early stages of an incident, when decisions on suppression planning, evacuation, and resource allocation must be made under significant time pressure.

Although physics-based and semi-empirical wildfire simulators remain fundamental tools for fire behaviour analysis, their operational use in early response is often constrained by the time and expertise required for setup and execution. This limitation motivates the development of complementary approaches that can provide immediate insight without relying on real-time simulation. One promising alternative is to frame wildfire support not only as a forecasting problem, but also as a retrieval problem: given an active incident, the system searches for historically observed or pre-simulated fires with similar evolution patterns and uses these analogous cases to inform likely future behaviour. This is especially applicable when using physic-based models with high computational

demand that, in essence, do not allow for operational pipelines. Such is the case with slope-scale or urbanization-scale simulations, covering several hectares, using models such as FDS/WFDS.

A central challenge in such a framework is the representation of wildfire similarity. Static geometric comparisons are insufficient, since wildfire behaviour is inherently dynamic and depends on the temporal evolution of shape, growth, and directional spread. To address this, the present work employs Long Short-Term Memory (LSTM) networks (Vennerød et al., 2021) to model wildfire evolution as multivariate sequential data and to learn latent embeddings that capture the spatiotemporal structure of fire trajectories. These learned representations are then used for similarity-based retrieval, enabling rapid identification of relevant prior cases. In this way, the proposed framework aims to complement traditional forecasting pipelines with a fast, data-driven mechanism for wildfire situational awareness and early decision support.

The main contributions of this work can be summarized as follows:

1. We formulate early wildfire support as a similarity-based retrieval task, enabling rapid access to historical or pre-simulated fire trajectories relevant to an active incident.
2. We propose an LSTM-based representation learning framework that encodes wildfire evolution as compact sequence embeddings, capturing temporal and geometric characteristics beyond static perimeter matching.
3. We show that the learned representations are effective for structure-preserving retrieval and sequential modelling, indicating their usefulness as a fast complementary tool for wildfire situational awareness.

2. Related Work

Wildfire spread has traditionally been studied through empirical, semi-empirical, and physics-based models. These approaches remain fundamental because they provide physically grounded descriptions of fire behavior, but they often involve a trade-off between realism, generalizability, and computational efficiency. In particular, empirical and quasi-empirical approaches are attractive for operational use due to their simplicity, whereas physical and quasi-physical models can capture fire–atmosphere interactions in greater detail at substantially higher computational cost (Sullivan, 2007a, 2007b).

More recently, machine learning and deep learning methods have emerged as data-driven alternatives for wildfire prediction. Traditional machine learning approaches, such as those explored by Khanmohammadi et al. (2022) and Singh et al. (2023), have been used to model wildfire spread from environmental and geospatial variables. At the deep learning level, convolutional architectures have been used to capture spatial patterns from remote sensing and terrain data, as in FireCast (Radke et al., 2019) and the U-Net-based approach of Khennou and Akhloufi (2023). Recurrent and hybrid spatiotemporal models have also been investigated, including LSTM-based wildfire spread prediction (Li et al., 2021), ConvLSTM models for time-resolved fire dynamics (Burge et al., 2020, 2022), and attention-enhanced ConvLSTM architectures that improve the modeling of wildfire spread dynamics and transferability across regions (Masrur et al., 2024).

Recent work has further expanded toward Transformer-based and hybrid decision-support systems. Dubey and Dubey (2025) combine Transformer-based spatiotemporal prediction with fuzzy inference, while Chen et al. (2024) propose AutoST-Net, a framework integrating convolutional and Transformer-based components for fire spread prediction from remote sensing data. Overall, these studies show a clear shift toward richer spatiotemporal modeling. However, most prior work focuses primarily on direct prediction. By contrast, the present study emphasizes **representation learning for similarity-based retrieval**, investigating whether LSTM-based embeddings can organize wildfire trajectories into

a latent space that supports retrieval of analogous fire evolutions while retaining predictive information about spread.

3. Methodology

3.1 Problem formulation

Let a wildfire event be represented as a multivariate temporal sequence $X^{(i)} = \{x_1^{(i)}, x_2^{(i)}, \ldots, x_{T_i}^{(i)}\}$, $x_t^{(i)} \in \mathbb{R}^F$, where T_i denotes the length of the i-th fire trajectory and F the number of features describing the fire state at each timestep. In the present setting, each timestep summarizes the geometric state of the evolving fire perimeter through variables such as centroid coordinates, area, bounding-box dimensions, eccentricity, and moment-based shape descriptors, together with engineered angular features. Since wildfire trajectories have variable duration, T_i is not constant across events.

The objective is not only to model wildfire evolution as a sequential process, but to support similarity-based retrieval. Given a query wildfire trajectory X^q, observed either fully or only up to an early prefix, the task is to retrieve from a database $\mathcal{D}$=$\{X^{(1)}, \ldots, X^{(N)}\}$ those historical or pre-simulated fire events whose temporal evolution is most similar to the query. Formally, we seek a representation function $f_\theta : \mathcal{X} \to \mathbb{R}^d$ that maps each variable-length wildfire sequence to a fixed-dimensional embedding vector $z^{(i)} = f_\theta(X^{(i)}), \qquad z^{(i)} \in \mathbb{R}^d$ such that similarity in the latent space reflects similarity in wildfire evolution dynamics.

Under this formulation, retrieval reduces to a nearest-neighbor search problem in the embedding space. For a query event X^q, its latent representation is $z^q = f_\theta(X^q)$ and the retrieved set consists of the k closest database embeddings according to a distance function $\delta(\cdot, \cdot)$, here taken to be Euclidean distance:

$$\mathcal{N}_k(X^q) = \mathrm{TopK}_{X^{(i)} \in \mathcal{D}} \ \delta(z^q, z^{(i)}).$$

Thus, instead of performing real-time fire-spread simulation, the system identifies analogous historical or simulated trajectories whose subsequent development may provide immediate operational insight.

3.2 Retrieval-oriented representation learning

The central methodology learns a latent space optimized for wildfire trajectory retrieval. The embedding must be both informative (preserving temporal structure via an autoencoding objective) and structure-preserving, modeling fire evolution as a continuous manifold rather than discrete clusters. This design choice is consistent with the empirical observation that wildfire evolution varies along a continuum and does not naturally partition into sharply separated regimes.

Accordingly, the methodology adopts a two-stage perspective:

1. Representation learning through an LSTM autoencoder, which compresses each wildfire trajectory into a fixed-dimensional latent vector.
2. Similarity-based retrieval in the learned latent space, where embeddings are compared using nearest-neighbor search.

3.3 Sequence representation and preprocessing

Each fire simulation is transformed into a multivariate time series. For the reconstruction setting, the primitive features include centroid coordinates, area, width, height, eccentricity, moment-related information, and rectangle-center coordinates. To better represent angular behavior, the original

rotation angle is transformed into sine and cosine components: $x_t^{\text{angle-sin}} = \sin(\phi_t), \quad x_t^{\text{angle-cos}} = \cos(\phi_t)$, where ϕ_t is the rotation angle at timestep t. This transformation avoids discontinuities at the $0°/360°$ boundary and yields a numerically smoother representation for learning.

For the forecasting setting, an additional engineered variable is introduced: $\Delta a_t = a_t - a_{t-1}$, where a_t denotes the fire area at timestep t. This feature explicitly captures local fire-growth dynamics and was introduced because preliminary experiments showed that the model had difficulty learning expansion and shrinkage patterns from raw area values alone. The dataset is split at the sequence level into training and test subsets using an 80/20 partition, ensuring that all timesteps of a given simulation remain in the same subset and preventing data leakage. Feature-wise Min–Max normalization is applied with the scaling parameters estimated only on the training set. Variable-length sequences are padded for batch training and masked to exclude padded steps from computations.

3.4 Autoencoder for latent embedding learning

To learn wildfire embeddings, the study employs a sequence-to-sequence LSTM autoencoder. Let $X = \{x_1, \ldots, x_T\}, \quad x_t \in \mathbb{R}^F$ be an input wildfire sequence. The encoder processes the sequence through one or more stacked LSTM layers and produces hidden states $h_t = \text{LSTM}(x_t, h_{t-1}), \quad h_t \in \mathbb{R}^H$.

The encoder follows a funnel-shaped architecture, where hidden dimensionality is progressively reduced across layers, thereby compressing temporal information into a compact latent vector. In the implementation, one to three stacked LSTM layers are used, with latent dimensionality selected from $\{8, 16, 32, 64, 128, 256\}$.

Two alternative strategies are considered for obtaining the latent representation.

- **Final-state encoding:** Only the final hidden state is retained: $z = h_T$, providing a compact summary of the full input trajectory.
- **Attention-like temporal pooling:** The encoder outputs all hidden states $\{h_1, \ldots, h_T\}$, and a masked attention-like pooling mechanism is applied. For each timestep, a scalar relevance score is computed: $e_t = w^\top h_t + b$, where $w \in \mathbb{R}^H$ and $b \in \mathbb{R}$ are trainable parameters. After masking padded positions, normalized importance weights are obtained through softmax:

$$\alpha_t = \frac{\exp(e_t)}{\sum_{\tau=1}^{T} \exp(e_\tau)}, \qquad \sum_{t=1}^{T} \alpha_t = 1$$

The latent representation is then computed as the weighted sum of encoder states: $z = \sum_{t=1}^{T} \alpha_t h_t$.

This mechanism allows the model to emphasize the most informative timesteps in the wildfire trajectory. The decoder reconstructs the original sequence from z, and the model is trained by minimizing reconstruction loss over valid, non-padded timesteps. If $\hat{x}_t$ denotes the reconstructed feature vector, the training objective is

$$\mathcal{L}_{\text{rec}} = \frac{1}{\sum_t m_t} \sum_{t=1}^{T} m_t \, \|x_t - \hat{x}_t\|_2^2,$$

where $m_t \in \{0, 1\}$ is a mask indicating whether timestep t is valid or padded. This masked mean squared error ensures that training focuses only on actual observations.

3.5 Retrieval stage

Once the autoencoder has been trained, only the encoder f_θ is retained for downstream retrieval. Each wildfire sequence in the database is embedded once: $z^{(i)} = f_\theta(X^{(i)}), \qquad i = 1, \ldots, N$. Given a new query event X^q, the same encoder produces z^q, and retrieval is performed by ranking database elements according to latent distance: $d_i = \|z^q - z^{(i)}\|_2$.

Small distances in the latent space indicate highly similar trajectories, enabling rapid retrieval of historical fires. This is especially valuable for early-stage decision support based on incomplete observations. Crucially, our results show that the learned embeddings remain stable across shorter sequence prefixes, confirming their reliability for early retrieval scenarios.

3.6 Forecasting architecture

In addition to retrieval-oriented representation learning, the methodology includes a forecasting component based on an LSTM encoder–decoder for short-term wildfire prediction. The purpose of this module is not only predictive evaluation, but also to assess whether the learned sequential modelling framework captures meaningful wildfire dynamics.

Training samples are generated using a growing sliding-window strategy. For a sequence $X = \{x_1, \ldots, x_T\}$, an input window contains the observed history up to timestep t, $X_{1:t} = \{x_1, \ldots, x_t\}$, while the target corresponds to the future change after a forecast horizon h: $y_t^{(h)} = x_{t+h} - x_t$. Thus, instead of predicting absolute future states directly, the model predicts temporal increments: $\hat{y}_t^{(h)} \approx x_{t+h} - x_t$. This reformulation was adopted because direct prediction of absolute values led to regression toward the mean and poor capture of rapid growth phases. Predicting differences instead stabilizes optimization and encourages the network to model local temporal dynamics.

Like the autoencoder, the forecasting model uses pre-padded and masked variable-length inputs. Training is augmented by generating multiple input-target pairs per sequence, enforcing a minimum three-timestep history. Finally, a stride of two is used during sample generation to reduce redundancy.

3.7 Evaluation of retrieval-oriented embeddings

Since, the downstream objective is retrieval, embedding quality is not evaluated only through reconstruction error, but also through the extent to which the latent space preserves similarity relations among wildfire trajectories.

Let $N_k^{\text{orig}}(i)$ and $N_k^{\text{emb}}(i)$ denote the sets of the k-nearest neighbors of sample i in the original and embedding spaces, respectively. The kNN overlap is defined as:

$$\text{kNNOverlap}(i) = \frac{|N_k^{\text{orig}}(i) \cap N_k^{\text{emb}}(i)|}{k},$$

and measures the fraction of shared neighbors.

To further assess local structure preservation, we use trustworthiness and continuity, defined as:

$$\text{Trustworthiness}(k) = 1 - \frac{2}{nk(2n-3k-1)} \sum_{i=1}^{n} \sum_{j \in U_k(i)} \left(r_{ij}^{\text{orig}} - k\right),$$

$$\text{Continuity}(k) = 1 - \frac{2}{nk(2n-3k-1)} \sum_{i=1}^{n} \sum_{j \in V_k(i)} \left(r_{ij}^{\text{emb}} - k\right),$$

where $U_k(i)$ contains neighbors that appear in the embedding space but not in the original space, and $V_k(i)$ contains true neighbors from the original space that are not preserved after embedding.

Finally, global structure is assessed using Spearman distance-rank correlation to compare original and embedded pairwise distance rankings. Together, these metrics evaluate if the latent space supports reliable nearest-neighbor retrieval.

4. Experimental Evaluation

4.1 Experimental setup

The experimental dataset consists of wildfire simulation trajectories generated for two representative case-study areas in Spain, namely Rectoret (Catalonia) and La Pedriza (Madrid). Wildfire events were represented as variable-length multivariate time series describing the geometric evolution of the fire perimeter over time. Beyond primary retrieval, the framework was evaluated via two experiments: **representation learning**, to test for similarity-preserving embeddings, and **forecasting**, to determine if the model captures predictive wildfire dynamics. The data were split at the sequence level into training and test subsets. Feature normalization was fitted on the training set only and then applied to the test set. Variable-length sequences were handled through padding and masking, while angular variables were encoded through sine and cosine transformations.

For representation learning, the LSTM autoencoder was evaluated using latent dimensions $d \in \{8, 16, 32, 64, 128, 256\}$, 1-3 LSTM layers, dropout rates $\{0.1, 0.2, 0.3, 0.4, 0.5\}$, batch sizes $\{16, 32, 64\}$, and learning rates $h \in \{10^{-2}, 10^{-3}, 10^{-4}\}$. Two sequence representation strategies were examined: the final hidden state and attention-based pooling over encoder states. For forecasting, prediction horizons of 1, 3, and 5 timesteps ahead were considered, while input history was analyzed over the ranges $[5, 10), [10, 15), [15, 20), [20, 25)$, and $[25, 99)$ timesteps. Hyperparameter optimization was performed using the Optuna framework (Akiba et al., 2019).

4.2 Representation Learning Experiment

To determine if wildfire trajectories can be encoded for similarity-based retrieval, an LSTM autoencoder was trained to compress sequences while retaining their temporal and geometric structure. Evaluation focused on two areas: reconstruction quality (masked MSE), which drove the learning process, and embedding quality (k-NN overlap, trustworthiness, continuity, and Spearman correlation), which determined the space's usefulness for retrieval. To verify robustness under partial observation, this analysis was repeated across sequences of 10, 20, and 37 timesteps.

5. Results

5.1 Embedding quality for similarity-based retrieval

The first set of results examines whether the learned latent representations are suitable for similarity-based retrieval of wildfire trajectories. As shown in **Table 1**, the learned embeddings preserve important structural relations in the original sequence space. The strongest results reach **0.4507** in kNN overlap, 0.9902 in continuity, **0.9865** in trustworthiness, and **0.8994** in Spearman distance-rank correlation, indicating that the latent space retains both local neighborhood structure and broader relational geometry. These findings support the use of the learned representations as a similarity space for retrieval. Evaluating initial fire 'prefix' durations (10, 20, and 37 timesteps) confirms that while longer observations maximize precision, the model maintains robust trustworthiness and continuity even at 10 timesteps, demonstrating its effectiveness for rapid, early-stage decision support.

Table 1 - Embedding quality summary for sequence length = 20 and k=10 neighbors.

Metric Best value	Metric
kNN overlap	0.4507
Continuity	0.9902
Trustworthiness	0.9865
Spearman distance-rank correlation	0.8994

Figure 1 provides a qualitative view of the learned embedding space using PCA (Shlens, 2014) , t-SNE (van der Maaten and Hinton, 2008), and UMAP (McInnes et al., 2020) projections for a single experiment. The visualizations suggest that the latent space preserves meaningful neighborhood structure, as several wildfire trajectories appear organized into local regions of proximity, while others remain distant. This indicates that the model captures graded similarity relations among fire evolutions, allowing some events to be represented as clearly more similar to each other than to the rest. The overall organization remains relatively continuous, which is consistent with the gradual and heterogeneous nature of wildfire evolution.

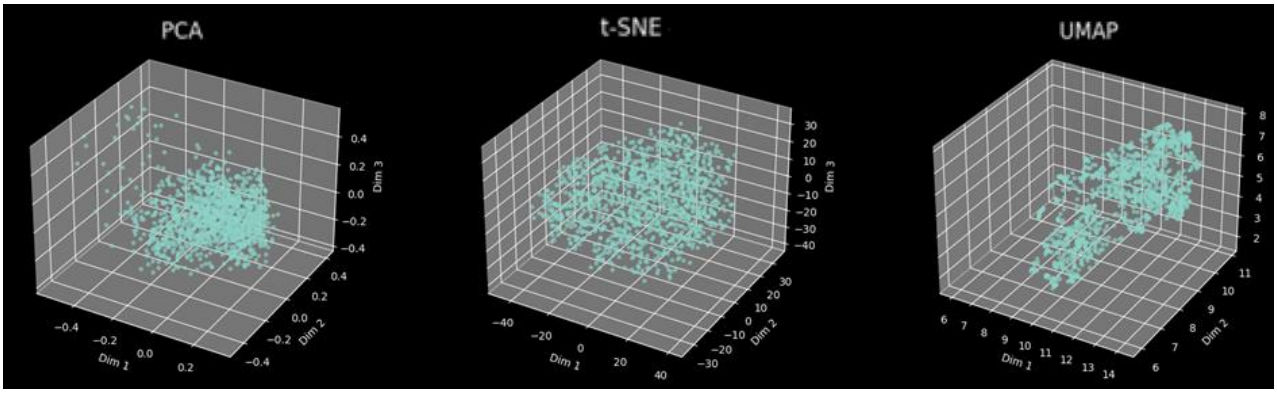


Figure 1 - PCA, t-SNE, and UMAP projections of wildfire embeddings learned by a representative LSTM-based autoencoder configuration.

5.2 One-step forecasting performance

The one-step forecasting results show that the proposed LSTM framework improves overall predictive accuracy relative to the baselines. As shown in **Table 2**, the persistence baseline remains strong, which is expected given the short five-minute interval between consecutive observations. Nevertheless, the LSTM model achieves lower **overall RMSE** (**0.0870** vs. **0.1006**) and lower **area-relative error** (**0.0541** vs. **0.0951**), while maintaining comparable geometric consistency in terms of bounding-box IoU (**0.8143** vs. **0.8249**). These results indicate that the recurrent model captures wildfire growth dynamics more accurately while preserving the predicted fire footprint geometry at a similar level.

5.3 Effect of input-history length

A particularly important result concerns the effect of available temporal context. As shown in **Table 3**, one-step forecasting performance improves consistently as longer input histories become available. Specifically, the overall RMSE decreases from **0.1050** in the **[5,10)** bucket to **0.0713** in the **[25,99)**

bucket, while the bounding-box IoU increases from **0.7056** to **0.8857**. The same trend is observed for centroid, area, aspect-ratio, and rotation-related errors, showing that richer temporal context substantially improves both numerical accuracy and geometric consistency.

This pattern also persists for longer forecasting horizons. For **three-step-ahead prediction**, the bounding-box IoU increases from **0.5711** to **0.7824** as the available input history grows, while for **five-step-ahead prediction** it increases from **0.5035** to **0.7165**. These findings indicate that the proposed sequence model remains effective even under more challenging multi-step forecasting settings, although performance naturally degrades as the horizon increases.

Table 2 - One-step forecasting performance.

Model	*Overall RMSE*	*Centroid error*	*Relative centroid error*	*Area relative error median*	*Aspect ratio relative error median*	*Rotation error (rad)*	*Bounding box IoU*
Last-value baseline	*0.1006*	*2.9532*	*0.024*	*0.0951*	*0.0371*	*0.1644*	*0.8249*
Linear baseline	*0.1344*	*8.5328*	*0.0475*	*0.1848*	*0.1197*	*0.3003*	*0.3232*
Best LSTM model	*0.087*	*3.2286*	*0.0309*	*0.0541*	*0.0547*	*0.1794*	*0.8143*

Table 3 - Effect of input-history length on one-step forecasting performance.

Input length bucket	Overall RMSE	Centroid error	Relative centroid error	Area relative error median	Aspect ratio relative error median	Rotation error (rad)	Bounding box IoU
[5,10)	0.1050	3.5719	0.0612	0.2208	0.1297	0.2622	0.7056
[10,15)	0.0940	3.3735	0.0323	0.0947	0.0773	0.2100	0.7913
[15,20)	0.0854	3.1644	0.0213	0.0565	0.0545	0.1729	0.8350

[20,25)	0.0791	3.1159	0.0160	0.0407	0.0437	0.1477	0.8599
[25,99)	0.0713	2.9488	0.0111	0.0284	0.0324	0.1201	0.8857

6. Conclusion

This paper presented an LSTM-based framework for similarity-driven wildfire analysis, aimed at supporting decision-making during the critical early stages of an incident. Instead of relying exclusively on time-consuming real-time simulation, the proposed approach represents wildfire evolution as sequential data and learns latent embeddings that enable retrieval of similar historical or pre-simulated fire trajectories.

The results showed that the learned representations preserve meaningful similarity structure, making them suitable for retrieval-oriented analysis, while the forecasting experiments confirmed that the same sequential modelling framework captures important wildfire dynamics. In particular, prediction quality improved consistently as longer temporal context became available, highlighting the importance of sequential information for both retrieval and short-term forecasting. Varying initial fire prefix duration ablations demonstrated that the retrieval mechanism maintains reliable precision even with limited early-stage data, reinforcing its utility for rapid situational awareness.

Overall, the findings suggest that LSTM-based representation learning can provide a practical complement to traditional wildfire forecasting pipelines. By combining similarity-based retrieval with predictive sequence modelling, the proposed framework offers a fast, data-driven tool for wildfire situational awareness, early assessment, and resource planning.

7. Acknowledgements

This work was supported by the HIDALGO2 project, funded by the European High Performance Computing Joint Undertaking (EuroHPC JU) under grant agreement No. 101093590. Computing resources were provided by LUMI (CSC, Finland) and MeluXina (LuxProvide, Luxembourg) through EuroHPC project allocations project_465002443 and p200983.

8. References

Burge, J., Bonanni, M., Ihme, M., & Hu, R. L. (2020). Convolutional LSTM Neural Networks for Modeling Wildland Fire Dynamics. *arXiv preprint https://arxiv.org/abs/2012.06679*.

Burge, J., Bonanni, M. R., Hu, R. L., & Ihme, M. (2022). Recurrent convolutional deep neural networks for modeling time-resolved wildfire spread behavior. Fire Technology, 59(6), 3327–3354.

Chen, X., Tian, Y., Zheng, C., & Liu, X. (2024). AutoST-Net: A spatiotemporal feature-driven approach for accurate forest fire spread prediction from remote sensing data. Forests, 15(4).

Dubey, P., & Dubey, P. (2025). Bridging spatiotemporal wildfire prediction and decision modeling using transformer networks and fuzzy inference systems. MethodsX, 15, 103498.

Khanmohammadi, S., Arashpour, M., Golafshani, E. M., Cruz, M. G., Rajabifard, A., & Bai, Y. (2022). Prediction of wildfire rate of spread in grasslands using machine learning methods. Environmental Modelling & Software, 156, 105507.

Khennou, F., & Akhloufi, M. A. (2023). Improving wildland fire spread prediction using deep U-Nets. Science of Remote Sensing, 8, 100101.

Li, X., et al. (2021). Prediction of forest fire spread rate using UAV images and an LSTM model considering the interaction between fire and wind. Remote Sensing, 13(21).

Masrur, A., Yu, M., & Taylor, A. (2024). Capturing and interpreting wildfire spread dynamics: attention-based spatiotemporal models using ConvLSTM networks. Ecological Informatics, 82, 102760.

Radke, D., Hessler, A., & Ellsworth, D. (2019). FireCast: Leveraging deep learning to predict wildfire spread. Proceedings of the International Joint Conference on Artificial Intelligence, 4575–4581.

Singh, A., Yadav, R., Sudhamshu, G., Basnet, A., & Ali, R. (2023). Wildfire spread prediction using machine learning algorithms. 2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT).

Sullivan, A. L. (2007a). A review of wildland fire spread modelling, 1990–present 1: Physical and quasi-physical models.

Sullivan, A. L. (2007b). A review of wildland fire spread modelling, 1990–present 2: Empirical and quasi-empirical models.

Vennerød, C. B., Kjærran, A., & Bugge, E. S. (2021). *Long short-term memory RNN*. arXiv preprint arXiv:2105.06756.

Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*(pp. 2623–2631).

van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research, 9*(86), 2579–2605.

Shlens, J. (2014). *A tutorial on principal component analysis*. arXiv preprint arXiv:1404.1100.

McInnes, L., Healy, J., & Melville, J. (2020). *UMAP: Uniform manifold approximation and projection for dimension reduction*. arXiv preprint arXiv:1802.03426.