

# BCWildfire: A Long-term Multi-factor Dataset and Deep Learning Benchmark for Boreal Wildfire Risk Prediction

Zhengsen Xu<sup>1</sup>, Sib0 Cheng<sup>2</sup>, Lanying Wang<sup>3</sup>, Hongjie He<sup>3,4</sup>, Wentao Sun<sup>5</sup>, Jonathan Li<sup>3,4,5</sup>,  
Lincoln Linlin Xu<sup>1\*</sup>

<sup>1</sup>Department of Geomatics Engineering, University of Calgary, Canada

<sup>2</sup>CEREA, ENPC, EDF R&D, Institut Polytechnique de Paris, France

<sup>3</sup>Department of Environmental Management, University of Waterloo, Canada

<sup>4</sup>Hinton STAI Institute, East China Normal University, China

<sup>5</sup>Department of Systems Design Engineering, University of Waterloo, Canada

{zhengsen.xu, lincoln.xu}@ucalgary.ca, sib0.cheng@enpc.fr, {hongjie.he, lanying.wang, wentao.sun, junli}@uwaterloo.ca

## Abstract

Wildfire risk prediction remains a critical yet challenging task due to the complex interactions among fuel conditions, meteorology, topography, and human activity. Despite growing interest in data-driven approaches, publicly available benchmark datasets that support long-term temporal modeling, large-scale spatial coverage, and multimodal drivers remain scarce. To address this gap, we present a 25-year, daily-resolution wildfire dataset covering 240 million hectares across British Columbia and surrounding regions. The dataset includes 38 covariates, encompassing active fire detections, weather variables, fuel conditions, terrain features, and anthropogenic factors. Using this benchmark, we evaluate a diverse set of time-series forecasting models, including CNN-based, linear-based, Transformer-based, and Mamba-based architectures. We also investigate effectiveness of position embedding and the relative importance of different fire-driving factors.

## Code and Datasets —

<https://github.com/SynUW/BCWildfire>

## Introduction

Wildfires have experienced unprecedented escalation in frequency, scale, and intensity globally, with climate change exacerbating fire-prone conditions and creating an urgent need for accurate wildfire risk prediction models (Ager et al. 2019). Boreal regions, characterized by substantial carbon storage and high wildfire susceptibility, face particularly acute risks (Stephens et al. 2014; Jain et al. 2024). For instance, British Columbia’s wildfire patterns threaten both regional ecosystems and global climate stability, necessitating robust predictive frameworks for effective fire management strategies (Parisien et al. 2023; Daniels et al. 2024).

Data-driven approaches, especially deep learning architectures, have shown great promise in wildfire prediction due to their capacity to model complex nonlinear interactions among heterogeneous environmental variables (Xu et al.

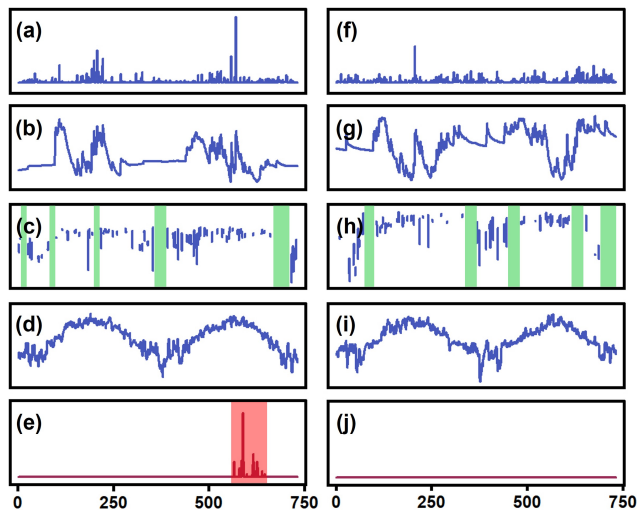


Figure 1: Time series of key variables (2023-2024) showing wildfire prediction challenges. Figure (a)-(e): precipitation, soil moisture L1, MODIS Band 20, 2m temperature (exogenous), and burned area (endogenous) at an ignition point. Figure (f)-(j): same variables in non-ignition areas. Red: wildfire periods; Green: missing data. The figure demonstrates spatiotemporal decoupling between exogenous/endogenous variables and remote sensing data gaps.

2025). However, the deployment of such models is often hindered by the lack of comprehensive and high-quality training datasets. Despite extensive research on wildfires, publicly available standardized datasets remain limited in terms of scope and accessibility (Gerard, Zhao, and Sullivan 2023).

Current wildfire prediction systems face fundamental challenges across multiple dimensions. A key limitation lies in the lack of standardized, publicly accessible datasets with broad geographic coverage, particularly in carbon-rich boreal ecosystems. Most existing research focuses on localized fire spread modeling (see Table 1), typically using input sequences constrained to short temporal windows such as 1 or 7 days (Eddin, Roscher, and Gall 2023; Deng et al. 2025;

\*Corresponding author.

Dataset	Temporal Span (Lookback/Forecast)	Resolution
FireCube	Daily/Daily	Daily
WildfireDB	Daily/Daily	Daily
WildfireSpreadTS	14–94 d/4–94 d	Daily
Next Day Wildfire Spread	Daily/Daily	Daily
SeasFire Cube	N/A/N/A	8 d
BCWildfire (ours)	$\leq 25$ y/ $\leq 25$ y	Daily

Table 1: Comparison between current next day or event-based wildfire prediction datasets and BCWildfire.

Zhou et al. 2025b). This narrow temporal scope fails to account for large-scale, long-term wildfire risk driven by fuel accumulation, prolonged drought, and anthropogenic activity over seasonal to annual timescales. In addition, the integration of diverse and heterogeneous fire-driving factors remains limited, which restricts models’ ability to capture the complex and multi-scale mechanisms that govern wildfire ignition and spread.

By contrast, methodological challenges arise from the inherent complexity of wildfire systems, where risk is driven by long-term cumulative interactions among diverse geospatial covariates, including meteorological conditions (e.g., temperature, humidity, precipitation), vegetation characteristics (e.g., fuel load, moisture content), human activity, and topography (Xu et al. 2025). The interdisciplinary nature of this domain has led to a knowledge gap: machine learning practitioners often lack the geospatial expertise required to effectively apply advanced algorithms to wildfire prediction, while wildfire experts continue to rely on traditional models such as Random Forest and XGBoost (Di Giuseppe et al. 2025). As a result, data-driven wildfire research has yet to fully leverage the recent advances in deep learning.

To address these fundamental limitations, we present BCWildfire, a comprehensive spatiotemporal dataset that covers British Columbia, Canada, and adjacent regions, spanning approximately 240 million hectares over the period 2000 to 2024 and incorporating 38 multimodal wildfire driving factors. We further conduct a systematic benchmark of deep learning architectures, including CNN-based, linear-based, Transformer-based, and Mamba-based models, for the next day wildfire risk prediction to evaluate their effectiveness on this challenging task.

The primary contributions of this work are:

- We introduce BCWildfire, a large-scale spatiotemporal dataset for wildfire risk prediction, consisting of 2.4 million samples across 240 million hectares of boreal landscapes over a 25-year period. The dataset integrates 38 multimodal wildfire-driving factors, addressing major limitations in existing resources related to geographic coverage, feature completeness, and long-term temporal representation.
- In contrast to existing datasets that primarily focus on wildfire spread and behavior modeling based on spatial information, BCWildfire is specifically designed to sup-

port time series forecasting. This design enables the modeling of cumulative interactions among wildfire drivers over time, thereby facilitating the development of models with enhanced temporal reasoning capabilities and the potential to establish new benchmarks for wildfire prediction.

- We conduct a systematic benchmarking of state-of-the-art deep learning architectures across 4 major paradigms including CNN-based, linear-based, Transformer-based, and Mamba-based models for spatiotemporal wildfire prediction. Our evaluation highlights the strength and limitations of each paradigm in handling complex variable dependencies, non-periodic patterns, and high stochastic inherent in wildfire dynamics.

## Related Work

### Datasets

Several open-source wildfire datasets have been released in recently, including WildfireDB (Singla et al. 2020), Next Day Wildfire Spread (Huot et al. 2022), FireCube (Prapas, Kondylatos, and Papoutsis 2022; Kondylatos et al. 2023), WildfireSpreadTS (Gerard, Zhao, and Sullivan 2023), and SeasFire Cube (Alonso et al. 2023) (see Table 1). These datasets compile environmental covariates related to wildfire events across the United States and the Mediterranean, incorporating meteorological conditions, topography, vegetation (fuel) status, human activity, and fire behavior. Most of these datasets are designed to support models that use recent covariate states and fire activity as inputs to predict wildfire spread at the next time step. Consequently, as summarized in Table 1, most studies generally adopt a daily temporal resolution with short look back periods and limited prediction horizons. Moreover, they mainly emphasize small scale wildfire spread and behavior prediction after ignition, rather than large scale wildfire risk prediction.

Overall, the number and spatiotemporal coverage of publicly available wildfire datasets remain limited. Existing datasets tend to focus on fire spread modeling for individual events within short time windows, primarily in the U.S. and Mediterranean regions. This limited scope constrains the development of wildfire ignition risk prediction models, which rely on long-term accumulative drivers and are crucial for proactive wildfire management (Xu et al. 2025). Moreover, due to the high sensitivity of wildfire behavior to local environmental and climatic conditions, models trained on region-specific or structurally constrained datasets often fail to generalize to boreal forests (Gerard, Zhao, and Sullivan 2023).

From a time series forecasting perspective, although there exists a wide range of datasets in domains such as electricity trading, energy consumption, traffic flow, and weather forecasting, there remains a notable scarcity of datasets focused on natural disaster prediction, particularly those involving complex and stochastic hazards like wildfires (Chen et al. 2001; Lai et al. 2018; Wu et al. 2021; Zhou et al. 2021). This scarcity limits the ability of machine learning researchers to apply and evaluate advanced time series forecasting algorithms in real-world disaster risk scenarios.

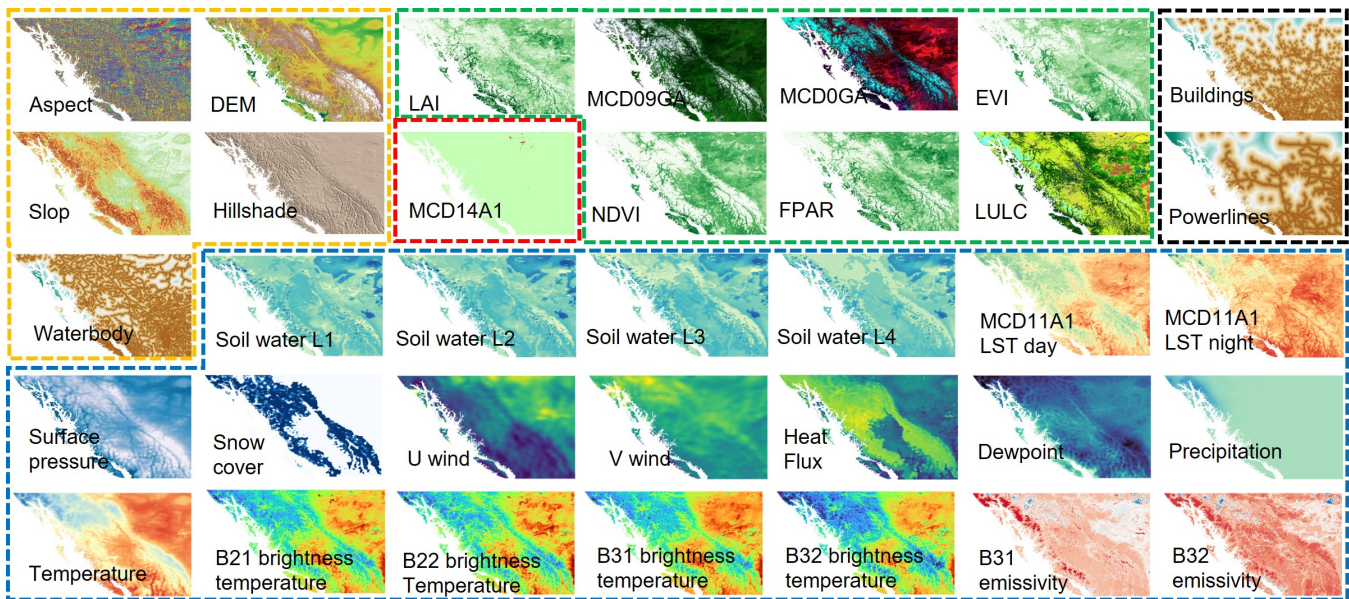


Figure 2: Partial overview of the dataset, illustrating fuel (green), fire detection (red), topography (yellow), human activities (black), and meteorological factors (blue).

## Methods

Current wildfire risk prediction research predominantly treats the problem as a spatiotemporal forecasting task, utilizing wildfire burning states and covariate conditions from previous days as inputs to predict wildfire coverage extent and behavior at next day (Shams Eddin, Roscher, and Gall 2023; Bhowmik et al. 2023; Ali et al. 2024). Consequently, algorithmic approaches frequently employ various CNN and Vision Transformer-based variants, such as UNet, UNETR, and T4-Fire, with primary emphasis on spatial information consideration, while neglecting the cumulative effects of long-term driving factors (Ronneberger, Fischer, and Brox 2015; Hatamizadeh et al. 2022; Zhao, Gerard, and Ban 2024). On the other hand, wildfire prediction algorithms based on long-term time series information remain predominantly grounded in traditional machine learning methodologies, including XGBoost, Random Forest, simple feed-forward neural networks, and Long Short-Term Memory (LSTM) (Kondylatos et al. 2022; Di Giuseppe et al. 2025).

Deep learning has greatly advanced time series forecasting, enabling models to infer future dynamics from historical observations. Multiple model architectures have emerged based on CNN, Linear, Transformer, and Mamba frameworks (Wang et al. 2025), including CNN-based TimesNet (Wu et al. 2022) and SCINet (Liu et al. 2022), Linear-based TiDE (Das et al. 2024), and CrossLinear (Zhou et al. 2025a), as well as Transformer-based models such as TimesNet (Wu et al. 2022), Autoformer (Wu et al. 2021), FEDformer (Zhou et al. 2022), Crossformer (Zhang and Yan 2023), PatchTST (Nie et al. 2023), and iTransformer (Liu et al. 2024). These models have demonstrated significant advantages across diverse prediction tasks in transportation, meteorology, and trading domains (Wang et al. 2025; Zhou et al. 2025a). How-

ever, exploration in natural disaster applications, particularly in wildfire risk prediction, remains notably limited.

## BCWildfire Benchmark Dataset

### Overview

Overall, we construct a long-term time series forecasting dataset with consistent spatiotemporal resolution, covering British Columbia, Canada, and surrounding regions from 2000 to 2024. Each sample includes 38 daily covariates grouped into five categories: vegetation (fuel) conditions, active fire detections, meteorological factors, human activity, and topographic features, as shown in Figure 2. The dataset primarily sources data from MODIS satellite remote sensing products, ERA5-Land reanalysis data (Muñoz-Sabater et al. 2021), OpenStreetMap (OpenStreetMap contributors 2021), and the ASTER digital elevation model (DEM) (NASA et al. 2019). As illustrated in Figure 2, each day’s observation is organized into a structured data cube with all features aligned. Given the varying native resolutions of the data sources, we resample all inputs to a uniform spatial resolution of 1 km and a daily temporal resolution to ensure consistency across both dimensions.

### Driving Factors

**Fuel Conditions** Vegetation serves as the primary fuel for wildfires; therefore, capturing its dynamic changes is essential for predicting wildfire occurrence and spread (Smith et al. 2023). Given the rapid dynamics of wildfire behavior, stable products that provide long-term time series with high temporal resolution are extremely scarce. Accordingly, we selected MODIS products to characterize fuel conditions. First, we used the stable 500 m, 4-day Leaf Area Index (LAI) composite product (MCD15A3H), which provides LAI and

the Fraction of Photosynthetically Active Radiation (FPAR) (Myneni, Knyazikhin, and Park 2015). These variables serve as proxies for fuel load, structure, moisture, and vegetation activity. In addition, to better capture the rapid temporal dynamics of vegetation, we employed the daily 500 m MOD/MYD09GA reflectance data, specifically Bands 1, 2, 3, and 7, together with the derived Normalized Difference Vegetation Index (NDVI) and Enhanced Vegetation Index (EVI), to jointly represent fuel conditions (Vermote and Wolfe 2015). The combination of these bands and products not only provides robust proxies for fuel status but also effectively distinguishes healthy vegetation from burned areas while maintaining sensitivity to clouds and smoke (Gerard, Zhao, and Sullivan 2023).

**Meteorological Factors** Weather and climate factors are primary drivers of wildfire activity, influencing fire occurrence, burned area, and fire behavior, thus occupying a central position in various wildfire prediction models (Xu et al. 2025). High temperature, low relative humidity, insufficient precipitation, and strong winds are key short-term meteorological conditions affecting fire occurrence and spread (Kondylatos et al. 2022). This study employs daily composite ERA5-Land products (with approximately 11 km spatial resolution) to extract meteorological driving factors (Muñoz-Sabater et al. 2021). In variable selection, beyond key factors that directly influence fuel moisture, wind transport, and fire propagation, such as 2 m air temperature, 10 m eastward and northward wind components, atmospheric pressure, total precipitation, surface latent heat flux, and 2 m dewpoint temperature, we additionally include snow cover as a seasonal variable with a significant impact in northern regions. Snow cover not only modulates surface temperature and humidity but also significantly delays the exposure and drying process of combustible materials during winter and spring seasons, thereby affecting wildfire occurrence probability and initial spread rate. Furthermore, to more comprehensively reflect the cumulative effects of environmental humidity over longer time scales, this study also incorporates soil moisture data from the 0–289 cm depth range.

It is important to note that the relatively coarse spatial resolution of ERA5-Land leads to spatially smoothed temperature fields, which are insufficient to capture fine-scale thermal anomalies associated with wildfire ignition and spread. To overcome this limitation, we additionally incorporate MODIS thermal products, including MOD/MYD11A1 and MOD/MYD09CMG, with spatial resolutions of approximately 1 km and 5.6 km, respectively (Wan, Hook, and Hulley 2021; Vermote 2015). The MOD/MYD11A1 product provides day- and night-time land surface temperature and emissivity derived from Bands 31 and 32, whereas the MOD/MYD09CMG product contains brightness temperature from Bands 20, 21, 31, and 32. By combining these MODIS products with ERA5-Land temperature data, we aim to enhance the spatial detail and responsiveness to local thermal anomalies.

**Topographical Factors** Topographical factors also play an important role in the ignition and spread of wildfires (Coen et al. 2013). Steep slopes help fire to rise quickly up-

hill, increasing the rate of spread; aspect affects sunlight exposure and dryness, thereby influencing the flammability of fuels; and hillshade reflects, to some extent, the shading of terrain on solar radiation, indirectly affecting local microclimate and burning conditions (Deng et al. 2025). Therefore, we selected 30-meter resolution ASTER DEM data and calculated slope, aspect, and hillshade based on it (NASA et al. 2019). Additionally, water bodies often serve as natural barriers that can significantly reduce the probability of fire spread, so we also calculated the distance to the nearest water body using OpenStreetMap's water distribution map, further enhancing the model's ability to represent fire behavior (OpenStreetMap contributors 2021).

**Human Activity Factors** Human activities represent a major source of wildfire ignition (Bowring et al. 2024). To capture their potential influence, we incorporated the yearly 500 m resolution MODIS Land Use product (MCD12Q1) to account for the impact of different land use types on wildfire occurrence, as substantial contrasts exist between agricultural and forested areas in terms of fuel characteristics and human intervention (Friedl and Sulla-Menashe 2019). In addition, using settlement and power grid distribution data from OpenStreetMap, we computed the distance to the nearest infrastructure for each pixel. Since human activities tend to cluster around infrastructure, such proximity may increase the likelihood of ignition events (OpenStreetMap contributors 2021).

**Wildfire Data** To construct temporally continuous and spatially consistent records of wildfire activity over the extended study period, we utilize daily MODIS Active Fire products MOD/MYD14A1, both with 1 km spatial resolution (Giglio and Justice 2015). These products are based on thermal anomaly detection in mid-infrared bands (primarily Bands 20 and 21), providing global coverage with up to four observations per day. MODIS Active Fire products offer several advantages, including extensive historical coverage, high temporal frequency, and consistent detection algorithms, making them particularly suitable for large-scale and long-term fire trend analysis. We merge MOD and MYD detection data daily to generate composite active fire masks, retaining only pixels with high confidence levels to reduce false positives. When a pixel is detected as actively burning, its value reflects the most recent observation time for that day. This approach reliably identifies fire location and intensity, enabling precise delineation of wildfire ignition timing and spatial extent. The resulting time series provides a robust foundation for downstream tasks such as wildfire risk modeling and fire spread analysis.

## Dataset Preparation

As our dataset integrates multimodal drivers from diverse sources, several challenges arise during its construction, including differences in geographic reference systems, spatial and temporal resolutions, and cloud contamination in optical remote sensing imagery. Moreover, since wildfires are extreme and relatively rare events, their spatiotemporal distribution is highly sparse compared to non-fire occurrences. Consequently, when wildfire risk prediction is formulated

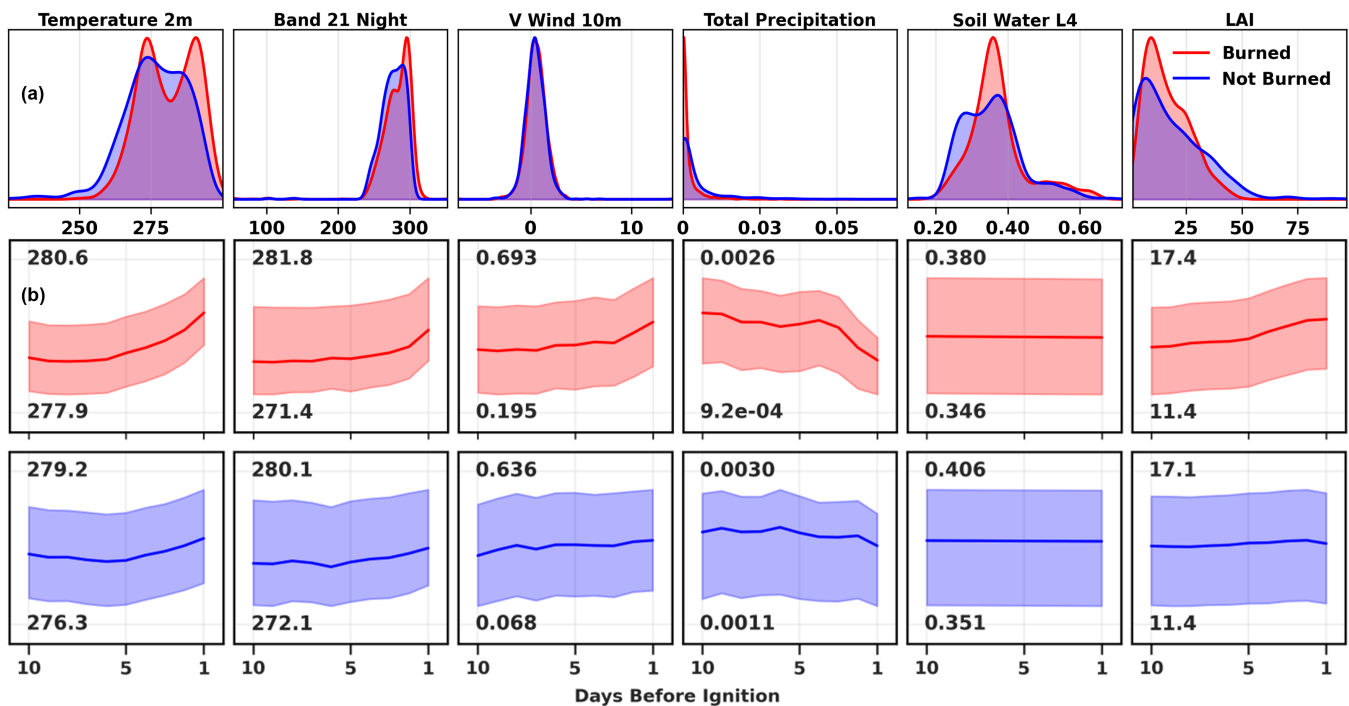


Figure 3: Comparison of key environmental drivers between burned and unburned areas prior to wildfire ignition. (a) Kernel density distributions of six selected variables within 10 days before ignition. (b) Mean temporal trends of the same variables.

as a binary classification problem, it suffers from severe class imbalance between positive (fire) and negative (non-fire) samples.

The study area is frequently influenced by maritime air masses. Moist westerlies from the Pacific Ocean are forced to ascend over the Rocky Mountains, resulting in orographic precipitation and persistent cloud accumulation, which cause substantial cloud contamination and data gaps in MODIS optical imagery (see Figure 1). To mitigate this issue, we applied the quality control (QC) bands provided by the MODIS products to mask low-quality observations, followed by temporal filling using historical records. This approach produces cloud-free observations without introducing future data leakage.

Given the substantial discrepancies in coordinate reference systems and spatiotemporal resolutions among different drivers, all spatial inputs were standardized to the WGS84 coordinate system to ensure geographic consistency. We further resampled all inputs to a uniform 1 km grid, using nearest-neighbor interpolation for categorical data (e.g., land use and fire occurrence) and bilinear interpolation for continuous variables. To address temporal heterogeneity, we forward-filled non-daily data sources under the assumption of short-term temporal stability, thereby aligning all variables into a unified daily sequence.

Ultimately, we constructed a large-scale, preprocessed wildfire risk assessment dataset in GeoTIFF format, covering British Columbia and adjacent regions from 2000 to 2024. The dataset comprises 38 daily variables over a spatial domain of  $2,782 \times 1,302 \text{ km}^2$  at 1 km resolution. All

files share identical spatial extent and resolution across and within variables, providing a flexible structure that facilitates customized preprocessing and dataset assembly.

In our experiments, wildfire risk prediction is formulated as a binary time-series prediction problem, where the goal is to predict the probability of wildfire occurrence at time  $t$  based on historical fire ignition and driver conditions from  $t - 1$  to  $t - n$  at the pixel level. Because wildfires are rare events, negative samples (non-fire pixels) greatly outnumber positive samples (fire pixels). To alleviate this imbalance, we adopted an undersampling strategy (Kondylatos, Camps-Valls, and Papoutsis 2025). In this benchmark, rather than randomly selecting negative samples, we excluded those located within a 60 km and 3-day spatiotemporal buffer around positive samples to avoid sampling from high-risk regions. Following the experimental setup of Kondylatos et al., we ensured that for each land cover type and across all years, the number of negative samples remained a fixed multiple of the positive samples (two times in the training and validation sets, and equal in the test set), preventing the model from learning trivial mappings. We also used the driver conditions from the previous 10 days as input to predict wildfire risk on the following day.

### Dataset Analysis

Following the aforementioned data collection and processing procedures, our final dataset for model training, validation, and testing comprises 1,015,275 samples, including 338,425 positive samples and 676,850 negative samples. To examine the distributional characteristics and temporal

Models	w/o Position				w/ Position			
	P	R	F1	PR_AUC	P	R	F1	PR_AUC
SCINet	84.77	88.05	86.38	94.46	85.18	86.63	85.90	94.33
TSMixer	85.69	90.39	87.97	96.24	86.98	89.29	88.12	96.10
CrossLinear	88.04	87.59	87.81	96.07	86.31	89.56	87.90	95.74
Crossformer	88.74	87.49	88.11	96.28	87.89	89.52	88.70	96.34
FEDformer	82.95	91.18	86.87	94.93	85.17	89.82	87.43	95.23
S_Mamba	84.21	86.44	85.31	94.83	84.99	90.09	87.46	95.74

Table 2: Comparison of model performance for next-day wildfire risk prediction using driver conditions from the previous 10 days and evaluating the effectiveness of positional embedding.

dynamics of wildfire-driving factors across multiple time scales, we computed the probability density functions for wildfire and non-wildfire pixels over the preceding 365 days and analyzed their temporal trends during the last 10 days prior to ignition, as shown in Figure 3(a) and Figure 3(b), respectively. The results indicate that temperature increases, fuel accumulation, precipitation decreases, and surface soil moisture reductions are strongly associated with wildfire ignition. However, substantial distributional overlap remains between burned and unburned areas, underscoring the intrinsic stochasticity of wildfire occurrence in space and time, whereby under similar environmental conditions, fires may either ignite or remain dormant even within high-risk zones. Interestingly, regions experiencing wildfire ignition tend to exhibit slightly higher deep-layer (100–289 cm) soil water content, suggesting that wildfires often occur in ecologically moist, fuel-rich environments where deep soil moisture does not directly mitigate surface drying and ignition risk.

## Experiments

### Evaluation Setups

**Evaluated Models.** We evaluated 6 time series forecasting models on the BCWildfire dataset, covering a diverse range of architectures, including CNN-based, linear-based, Transformer-based, and Mamba-based approaches. Specifically, the models include the CNN-based SCINet (Liu et al. 2022), linear-based models such as CrossLinear (Zhou et al. 2025a), and TSMixer (Wang et al. 2024), Transformer-based models including Corssformer (Zhang and Yan 2023) and FEDFormer (Zhou et al. 2022), as well as the Mamba-based S\_Mamba (Wang et al. 2025).

Through extensive comparative experiments across multiple models, we aim to address several key challenges at the intersection of wildfire risk assessment and time-series forecasting. These challenges include identifying the performance limits of existing time-series models for wildfire prediction, understanding the relative importance of wildfire-driving factors within the study area, and investigating how factors such as spatial encoding influence predictive performance and model interpretability. Our findings provide clearer guidance for the effective use of datasets and the design of models in this field.

**Evaluation Protocols.** We evaluate model performance using recall (R) and precision (P), F1-score (F1), and

precision-recall area under curve (PR-AUC). Binary cross-entropy loss is used for model training. All models are trained using the Adam optimizer on two NVIDIA A6000 GPUs. We use a batch size of 128 and train for 50 epochs. The learning rate is scheduled to  $1 \times 10^{-5}$  during training. For data partitioning, we use the years 2000-2020 as the training set, 2021-2022 as the validation set, and 2023-2024 as the test set.

### Main Results

**Mechanistic Drivers of Wildfire Risk.** To identify the short term mechanisms underlying wildfire ignition, we conducted a SHAP analysis on the Crossformer model using a 10 day input sequence to predict next day wildfire risk. The analysis indicates that ignition probability in our study region is primarily governed by recent fire activity, soil moisture dynamics, surface energy flux, and fuel conditions derived from multispectral reflectance.

The most influential feature is fire detection, whose strong negative SHAP contribution suggests that the absence of recent fire activity is a key signal for identifying potential new ignitions, while persistent fire signals correspond to already burned areas. Following this, soil moisture L3 (28–100 cm depth) and surface latent heat flux show strong positive effects, indicating that intermediate soil moisture and high surface energy exchange promote rapid drying of near surface fuels and create favorable ignition conditions. Snow cover exerts a pronounced negative influence, confirming that the presence of snow effectively suppresses ignition risk.

Spectral variables such as Band 3 (459–479 nm reflectance), Band 32 (11.77–12.27  $\mu\text{m}$  emissivity), and vegetation indices including EVI and FPAR contribute positively, revealing that bright visible reflectance, high emissivity, and productive vegetation jointly reflect exposed, desiccated, and fuel rich surfaces prone to burning. Topographic variables such as hillshade, slope, and aspect rank within the top 15 predictors, suggesting that terrain controlled radiation and wind exposure further modulate ignition likelihood, although their effects remain secondary to biophysical and thermal factors.

**Performance Ceilings** Our benchmark reveals persistent performance ceilings among existing wildfire prediction architectures under the 10 day input configuration. Even the best performing models, including CrossLinear (recall

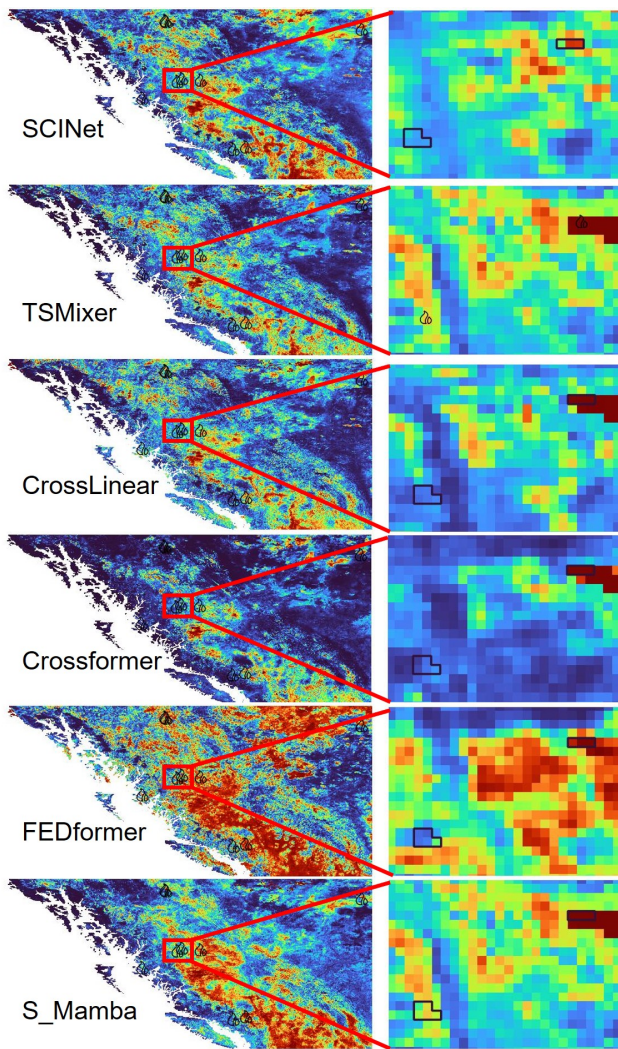


Figure 4: Visualization of predictions from different time series forecasting models on August 15, 2024.

87.56%), Crossformer (87.49%), and FEDformer (91.19%), plateau below 92% recall, while precision remains within a narrow 83–88% range. These results underscore fundamental challenges in wildfire risk prediction, such as extreme class imbalance, strong spatiotemporal heterogeneity, and complex nonlinear interactions among drivers.

Precision varies only slightly across architectures, suggesting that current representations still struggle to capture causal ignition signals and to suppress false positives. Transformer based architectures, especially Crossformer and FEDformer, achieve higher recall and stability than CNN or Mamba based models, reflecting their advantage in capturing long range temporal dependencies. Linear models such as TSMixer and CrossLinear remain competitive and efficient when temporal correlations are strong. Notably, S\_Mamba achieves balanced performance across all metrics (precision 84.21%, recall 86.44%, F1 85.31%, PR-AUC 0.9483), confirming its ability to model fine scale tem-

poral dynamics with good computational efficiency.

We further assessed the impact of incorporating spatial positional embedding, which encodes pixel level geographic information. As shown in Table 2, adding spatial context consistently improves accuracy across architectures. Transformer based models exhibit the largest gains in recall and F1 score, indicating that spatial cues enhance their ability to capture regional patterns and neighborhood dependencies. Linear and Mamba based models also benefit, with S\_Mamba showing clear improvement in recall (3.659%) and F1 (from 2.15%). Overall, spatial embedding provides valuable contextual information that improves model robustness and accuracy in wildfire risk prediction.

**Qualitative Comparison between Different Models** As shown on the left side of Figure 4, most wildfire occurrences correspond closely to areas classified as high risk. Combined with the quantitative comparison above, this result demonstrates that existing time series models can effectively predict next day wildfire risk. However, as illustrated in the zoomed in view on the right, model performance limitations remain evident, particularly in predicting new ignitions and small scale fires.

On the other hand, Figure 4 shows that different models produce distinct spatial patterns in their predictions, particularly in false positives and high-risk regions, even though their quantitative metrics in Table 2 differ only marginally. This observation suggests two key implications. First, ensemble predictions that combine different model outputs may further improve overall accuracy. Second, quantitative comparisons based solely on sampled test data may be biased and fail to fully capture the true predictive capability of each model.

## Conclusion

We present BCWildfire, a 25 year multimodal wildfire benchmark dataset covering 240 million hectares across British Columbia and adjacent northern regions. The dataset integrates 38 key wildfire driving factors, including meteorological, fuel, topographic, and human activity variables, harmonized at 1 km spatial and daily temporal resolution. Based on this resource, we establish the first unified time series forecasting benchmark for wildfire risk prediction.

Our experiments show that advanced architectures such as Crossformer, FEDformer, and S\_Mamba achieve strong predictive performance but remain constrained by the stochastic nature of ignition events and class imbalance between fire and non fire samples. SHAP analysis reveals that wildfire risk is shaped by recent fire activity, soil moisture transitions, surface energy flux, snow cover, vegetation productivity, and terrain effects, reflecting physically meaningful ignition mechanisms. Incorporating spatial positional embedding further improves accuracy across models, emphasizing the importance of spatial context in capturing regional fire dynamics.

Future work will extend the temporal horizon of modeling to better exploit the dataset’s long-term information and to examine how extended temporal drivers influence wildfire occurrence and evolution.

## Acknowledgments

This work was supported in part by the Natural Sciences and Engineering Research Council of Canada Discovery Grant (Grant No. RGPIN-2019-06744) and in part by MITACS in partnership with BMO. Zhengsen Xu acknowledges support from the China Scholarship Council.

## References

- Ager, A. A.; Palaiologou, P.; Evers, C. R.; Day, M. A.; Ringo, C.; and Short, K. 2019. Wildfire exposure to the wildland urban interface in the western US. *Applied geography*, 111: 102059.
- Ali, S. H.; Zhang, S.; Mhatre, S.; and Xiao, T. 2024. Advancing wildfire predictive models: A novel dataset for next-day wildfire spread. In *2024 Conference on AI, Science, Engineering, and Technology (AIxSET)*, 69–76. IEEE.
- Alonso, L.; Gans, F.; Karasante, I.; Ahuja, A.; Prapas, I.; Kondylatos, S.; Papoutsis, I.; Panagiotou, E.; Mihail, D.; Cremer, F.; Weber, U.; and Carvalhais, N. 2023. SeasFire Cube: A Global Dataset for Seasonal Fire Modeling in the Earth System.
- Bhowmik, R. T.; Jung, Y. S.; Aguilera, J. A.; Prunicki, M.; and Nadeau, K. 2023. A multi-modal wildfire prediction and early-warning system based on a novel machine learning framework. *Journal of environmental management*, 341: 117908.
- Bowring, S. P.; Li, W.; Mouillot, F.; Rosan, T. M.; and Ciaia, P. 2024. Road fragment edges enhance wildfire incidence and intensity, while suppressing global burned area. *Nature Communications*, 15(1): 9176.
- Chen, C.; Petty, K.; Skabardonis, A.; Varaiya, P.; and Jia, Z. 2001. Freeway performance measurement system: mining loop detector data. *Transportation research record*, 1748(1): 96–102.
- Coen, J. L.; Cameron, M.; Michalakes, J.; Patton, E. G.; Riggan, P. J.; and Yedinak, K. M. 2013. WRF-Fire: coupled weather-wildland fire modeling with the weather research and forecasting model. *Journal of Applied Meteorology and Climatology*, 52(1): 16–38.
- Daniels, L. D.; Dickson-Hoyle, S.; Baron, J. N.; Copes-Gerbitz, K.; Flannigan, M. D.; Castellanos-Acuna, D.; Hoffman, K. M.; Bourbonnais, M.; Wilkinson, S. L.; Roeser, D.; et al. 2024. The 2023 wildfires in British Columbia, Canada: impacts, drivers, and transformations to coexist with wildfire. *Canadian Journal of Forest Research*, 55: 1–18.
- Das, A.; Kong, W.; Leach, A.; Mathur, S.; Sen, R.; and Yu, R. 2024. Long-term Forecasting with TiDE: Time-series Dense Encoder. arXiv:2304.08424.
- Deng, J.; Hong, B.; Wang, W.; and Gu, G. 2025. Daily Wildfire Risk Prediction by Mining Global and local spatio-temporal dependency. *Earth Science Informatics*, 18(3): 1–20.
- Di Giuseppe, F.; McNorton, J.; Lombardi, A.; and Wetterhall, F. 2025. Global data-driven prediction of fire activity. *Nature Communications*, 16(1): 2918.
- Eddin, M. H. S.; Roscher, R.; and Gall, J. 2023. Location-aware adaptive normalization: a deep learning approach for wildfire danger forecasting. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–18.
- Friedl, M.; and Sulla-Menashe, D. 2019. MCD12Q1 MODIS/Terra+Aqua Land Cover Type Yearly L3 Global 500m SIN Grid V006. Date Accessed: 2025-10-24.
- Gerard, S.; Zhao, Y.; and Sullivan, J. 2023. Wildfire-SpreadTS: A dataset of multi-modal time series for wildfire spread prediction. *Advances in Neural Information Processing Systems*, 36: 74515–74529.
- Giglio, L.; and Justice, C. 2015. MOD14A1 MODIS/Terra Thermal Anomalies/Fire Daily L3 Global 1km SIN Grid V006. Date Accessed: 2025-10-24.
- Hatamizadeh, A.; Tang, Y.; Nath, V.; Yang, D.; Myronenko, A.; Landman, B.; Roth, H. R.; and Xu, D. 2022. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 574–584.
- Huot, F.; Hu, R. L.; Goyal, N.; Sankar, T.; Ihme, M.; and Chen, Y.-F. 2022. Next day wildfire spread: A machine learning dataset to predict wildfire spreading from remote-sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–13.
- Jain, P.; Barber, Q. E.; Taylor, S. W.; Whitman, E.; Castellanos Acuna, D.; Boulanger, Y.; Chavardès, R. D.; Chen, J.; Englefield, P.; Flannigan, M.; et al. 2024. Drivers and Impacts of the Record-Breaking 2023 Wildfire Season in Canada. *Nature Communications*, 15(1): 6764.
- Kondylatos, S.; Camps-Valls, G.; and Papoutsis, I. 2025. Uncertainty-Aware Deep Learning for Wildfire Danger Forecasting. arXiv:2509.25017.
- Kondylatos, S.; Prapas, I.; Camps-Valls, G.; and Papoutsis, I. 2023. Mesogeos: A multi-purpose dataset for data-driven wildfire modeling in the Mediterranean. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Kondylatos, S.; Prapas, I.; Ronco, M.; Papoutsis, I.; Camps-Valls, G.; Piles, M.; Fernández-Torres, M.-Á.; and Carvalhais, N. 2022. Wildfire danger prediction and understanding with deep learning. *Geophysical Research Letters*, 49(17): e2022GL099368.
- Lai, G.; Chang, W.-C.; Yang, Y.; and Liu, H. 2018. Modeling long-and short-term temporal patterns with deep neural networks. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, 95–104.
- Liu, M.; Zeng, A.; Chen, M.; Xu, Z.; LAI, Q.; Ma, L.; and Xu, Q. 2022. SCINet: Time Series Modeling and Forecasting with Sample Convolution and Interaction. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2024. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. In *The Twelfth International Conference on Learning Representations*.

- Muñoz-Sabater, J.; Dutra, E.; Agustí-Panareda, A.; Albergel, C.; Arduini, G.; Balsamo, G.; Boussetta, S.; Choulga, M.; Harrigan, S.; Hersbach, H.; et al. 2021. ERA5-Land: A state-of-the-art global reanalysis dataset for land applications. *Earth system science data*, 13(9): 4349–4383.
- Myneni, R.; Knyazikhin, Y.; and Park, T. 2015. MCD15A3H MODIS/Terra+Aqua Leaf Area Index/FPAR 4-Day L4 Global 500m SIN Grid V006. Date Accessed: 2025-10-24.
- NASA; METI; AIST; Spacesystems, J.; and the U.S./Japan ASTER Science Team. 2019. ASTER Global Digital Elevation Model V003. Data set. Accessed: 2025-12-24.
- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *The Eleventh International Conference on Learning Representations*.
- OpenStreetMap contributors. 2021. OpenStreetMap Database. <https://www.openstreetmap.org>. PostgreSQL via API. Data mining by Overpass Turbo. Available under the Open Database Licence.
- Parisien, M.-A.; Barber, Q. E.; Bourbonnais, M. L.; Daniels, L. D.; Flannigan, M. D.; Gray, R. W.; Hoffman, K. M.; Jain, P.; Stephens, S. L.; Taylor, S. W.; et al. 2023. Abrupt, climate-induced increase in wildfires in British Columbia since the mid-2000s. *Communications Earth & Environment*, 4(1): 309.
- Prapas, I.; Kondylatos, S.; and Papoutsis, I. 2022. FireCube: A Daily Datacube for the Modeling and Analysis of Wildfires in Greece.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.
- Shams Eddin, M. H.; Roscher, R.; and Gall, J. 2023. Location-Aware Adaptive Normalization: A Deep Learning Approach for Wildfire Danger Forecasting. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–18.
- Singla, S.; Diao, T.; Mukhopadhyay, A.; Eldawy, A.; Shachter, R.; and Kochenderfer, M. 2020. WildfireDB: A Spatio-Temporal Dataset Combining Wildfire Occurrence with Relevant Covariates. In *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*.
- Smith, J. T.; Allred, B. W.; Boyd, C. S.; Davies, K. W.; Jones, M. O.; Kleinhesselink, A. R.; Maestas, J. D.; and Naugle, D. E. 2023. Where there's smoke, there's fuel: dynamic vegetation data improve predictions of wildfire hazard in the Great Basin. *Rangeland Ecology & Management*, 89: 20–32.
- Stephens, S. L.; Burrows, N.; Buyantuyev, A.; Gray, R. W.; Keane, R. E.; Kubian, R.; Liu, S.; Seijo, F.; Shu, L.; Tolhurst, K. G.; et al. 2014. Temperate and boreal forest mega-fires: characteristics and challenges. *Frontiers in Ecology and the Environment*, 12(2): 115–122.
- Vermote, E. 2015. MOD09CMG MODIS/Terra Surface Reflectance Daily L3 Global 0.05Deg CMG V006. Date Accessed: 2025-10-24.
- Vermote, E.; and Wolfe, R. 2015. MOD09GA MODIS/Terra Surface Reflectance Daily L2G Global 1km and 500m SIN Grid V006. Date Accessed: 2025-10-24.
- Wan, Z.; Hook, S.; and Hulley, G. 2021. MODIS/Terra Land Surface Temperature/Emissivity Daily L3 Global 1km SIN Grid V061. Date Accessed: 2025-10-24.
- Wang, Y.; Wu, H.; Dong, J.; Qin, G.; Zhang, H.; Liu, Y.; Qiu, Y.; Wang, J.; and Long, M. 2024. TimeXer: Empowering Transformers for Time Series Forecasting with Exogenous Variables. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Wang, Z.; Kong, F.; Feng, S.; Wang, M.; Yang, X.; Zhao, H.; Wang, D.; and Zhang, Y. 2025. Is mamba effective for time series forecasting? *Neurocomputing*, 619: 129178.
- Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2022. Timesnet: Temporal 2d-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186*.
- Wu, H.; Xu, J.; Wang, J.; and Long, M. 2021. Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 22419–22430. Curran Associates, Inc.
- Xu, Z.; Li, J.; Cheng, S.; Rui, X.; Zhao, Y.; He, H.; Guan, H.; Sharma, A.; Erxleben, M.; Chang, R.; and Xu, L. L. 2025. Deep learning for wildfire risk prediction: Integrating remote sensing and environmental data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 227: 632–677.
- Zhang, Y.; and Yan, J. 2023. Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting. In *The eleventh international conference on learning representations*.
- Zhao, Y.; Gerard, S.; and Ban, Y. 2024. TS-SatFire: A Multi-Task Satellite Image Time-Series Dataset for Wildfire Detection and Prediction. *arXiv:2412.11555*.
- Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; and Zhang, W. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11106–11115.
- Zhou, P.; Liu, Y.; Liang, J.; Song, Q.; and Li, X. 2025a. CrossLinear: Plug-and-Play Cross-Correlation Embedding for Time Series Forecasting with Exogenous Variables. *arXiv preprint arXiv:2505.23116*.
- Zhou, T.; Ma, Z.; Wen, Q.; Wang, X.; Sun, L.; and Jin, R. 2022. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, 27268–27286. PMLR.
- Zhou, Y.; Kong, R.; Xu, Z.; Xu, L.; and Cheng, S. 2025b. Comparative and interpretative analysis of CNN and transformer models in predicting wildfire spread using remote sensing data. *Journal of Geophysical Research: Machine Learning and Computation*, 2(2): e2024JH000409.