

Can Satellite-Observed Wildfire Incidents Alone Project Next-Year State Transition? A Case Study in British Columbia, Canada Using a Physics-Regularized Conditional Categorical Generative Model

Yimin Zhu¹, Zhengsen Xu¹, Lincoln Linlin Xu¹

¹ Department of Geomatics Engineering, University of Calgary, Calgary, Canada

Keywords: Wildfire State Transition, Generative Model, Historical Records, Conditional Discrete Diffusion Model, Uncertainty Quantification, Long Time Series

Abstract

Wildfire risk plays a critical role in natural disaster emergency response and climate change mitigation. Existing data-driven methods mainly focus on short-term wildfire forecasting (days to months), often relying on numerous driving factors and binary fire/no-fire representations. In this paper, we redefine wildfire risk from the perspective of next-year wildfire state transitions and propose a novel Physical-Aware Wildfire State Transition Discrete Diffusion Model (PA-WildfireSTDDM) that directly learns the high-dimensional distribution of wildfire risk using only historical wildfire records, with the following contributions: (1) We construct a 25-year (2000–2024) daily wildfire dataset for British Columbia (BC), Canada, derived from FIRMS at 10 km spatial resolution through spatial aggregation. Four wildfire state transition categories are defined: *Persistent no-fire*, *New ignition*, *Fire cessation*, and *Persistent fire*. (2) We propose a physically aware diffusion framework, where the prediction process is regularized using a physical and empirical wildfire transition kernel, improving physical consistency and long-term prediction reliability. (3) The proposed end-to-end model predicts the categorical distribution of wildfire state transitions conditioned on historical wildfire events and preceding wildfire states, reducing cumulative errors compared with iterative forecasting approaches. (4) Our framework generates high-confidence maps of next-year wildfire states using only long-term historical wildfire records, without requiring additional environmental or climate-driving factors, while effectively capturing complex and stochastic wildfire patterns. (5) The model formulates wildfire evolution as a discrete-time inhomogeneous stochastic process, enabling uncertainty quantification for next-year wildfire projections through Monte Carlo posterior sampling.

1. Introduction

Wildfires are an important component of terrestrial ecosystem ecology, but also a major natural hazard to societies (Knorr et al., 2016). Wildfires in British Columbia (BC), Canada, have been particularly severe (Daniels et al., 2024). From 2003 to 2022, BC experienced an average of 1,350 wildfires annually, burning approximately 284,000 hectares (Taylor et al., 2022).

Hence, there is a real need for novel wildfire warning and prediction technologies that enable better fire management, mitigation, and evacuation decisions. In particular, assessing the fire likelihood – the probability of wildfire burning in a specific location – would provide valuable insight for forestry and land management, disaster preparedness, and early-warning systems (Huot et al., 2021).

Traditional methods like empirical process or statistical models, although they are interpretable due to their reliance on well established relationships between fire occurrences and fire weather condition (Abatzoglou and Williams, 2016), but the imperfect parameterization of the emergent climate-fire relationships and inadequate representation of complex interactions between fires and their drivers, such as climate, fuel availability (Li et al., 2023, Littell, 2016). To overcome the inherent complexity of wildfire risk prediction, recent studies have employed advanced machine learning (ML) models along with critical social and environmental factors to predict fires (Li, 2024, Huot et al., 2021). Although they use some methods to make the ML models be more explainable and increase the model physical interpretability, but such complex driven factors and stochastic process of wildfire occurrence still hard to explain, and the lack of

uncertainty estimation of these model will effect the decision making because the model may be unwarranted confidence and raise latent risk.

To provide better understanding of wildfire risk prediction, we define the wildfire risk prediction task from the perspective of wildfire state transition of next year. The wildfire state transition types are treated as a set of discrete spatial units, and be depicted the discrete spatial units forward as stochastic process over time, and hence, proposing a novel approach that can directly capture the high-dimensional distribution of wildfire risk only through available historical wildfire events, and predict the next year wildfire state transition, with the following characteristics: A 25-year-long-term daily wildfire historical record for British Columbia (BC) province, Canada is built deriving from the Fire Information for Resource Management System (FIRMS) with 10 km spatial resolution, using spatial aggregation. We define four wildfire state transition types based on the presence or absence of fire in a three-year historical period versus the fourth year: *Persistent no-fire* (no fire in the four years), *New ignition* (only happens in the predicted year), *Fire cessation* (no fire happens in the fourth years), and *Persistent fire* (fire happens in the predicted year and also the preceding three years). (2) Our model is a physically aware method, where the prediction process is regularized by a physical and empirical transition kernel, contributing to transparent model interpretability, instead of pure data-driven methods. (3) The proposed model can generate a high confidence map of next year's wildfire state only through the long-term daily historical wildfire event without any other driving factors, and also correlate with the complex and stochastic wildfire pattern. (4) Our

model depicts the discrete wildfire state of each pixel forward as a discrete-time inhomogeneous stochastic process, making it well-suited for characterizing next year's wildfire state transition uncertainty in model projections by performing multiple posterior sampling through Monte Carlo.

2. Data and Method

2.1 Dataset Illustration

Historical wildfire data are from the MOD14A and MYD14A V6 dataset, provided by NASA LP DAAC at the USGS EROS Center (Giglio, 2015). We performed spatial aggregation at 10 km resolution and added up the number of wildfire records, filtering out those events with low confidence, to improve dataset quality. Figure 1 and Figure 2 show the 25-year wildfire state transition dataset. From Figure 1, we can see that the historical wildfire events (heatmap in Figure 1) follow a complex mixture distribution with multiple peaks along the latitude and longitude directions. The wildfire state transition of each corresponding three-year period also represents a multimodal distribution. Figure 2 shows the accumulated wildfire state transition counts across four types in 25 years. Based on this, we calculate the physical transition kernel, which serves as the constraint when performing the model inference.

2.2 Wildfire State Transition Discrete Diffusion Model

2.3 Discrete Diffusion Model

Our proposed model consists of a forward diffusion process to gradually corrupt the clean wildfire state transition categorical mask $\mathbf{x}_0 \in \{0, 1, \dots, K-1\}^{H \times W}$, yielding a sequence of increasingly noisy mask tokens $\mathbf{x}_t \in \{0, 1, \dots, K-1\}^{H \times W}$, $t \in \{1, 2, \dots, T\}$ via a fixed Markov chain $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ with total T timesteps, and a reverse diffusion process to gradually denoise the noisy masks towards the desired clean mask distribution by learning the conditional distribution. The forward diffusion process can be defined as:

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathbf{v}^T(\mathbf{x}_t)\bar{\mathbf{Q}}_t\mathbf{v}(\mathbf{x}_0), \bar{\mathbf{Q}}_t = \mathbf{Q}_t\mathbf{Q}_{t-1}\dots\mathbf{Q}_1 \quad (1)$$

where $\mathbf{v}(\mathbf{x}_t)$ is a K -length one-hot column vector. $\mathbf{Q}_t \in \mathbb{R}^{K \times K}$ demotes the transition matrix, where the i, j element in $\mathbf{Q}_t$ represents the probability from $\mathbf{x}_{t-1} = j$ to $\mathbf{x}_t = i$, i.e., $q(\mathbf{x}_t = i|\mathbf{x}_{t-1} = j)$. Following (Gu et al., 2022), we introduce and additional category denoted as M , and the transition matrix $\mathbf{Q}_t \in \mathbb{R}^{(K+1) \times (K+1)}$ is defined as follows:

$$\mathbf{Q}_t = \begin{bmatrix} \alpha_t + \beta_t & \beta_t & \beta_t & \dots & 0 \\ \beta_t & \alpha_t + \beta_t & \beta_t & \dots & 0 \\ \beta_t & \beta_t & \alpha_t + \beta_t & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \gamma_t & \gamma_t & \gamma_t & \dots & 1 \end{bmatrix} \quad (2)$$

The reverse process gradually denoises the noisy masks and tries to restore the desired clean mask $\mathbf{x}_0$ from noisy mask $\mathbf{x}_t$ using the posterior distribution $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$, which can be derived based on the forward diffusion process as:

$$\begin{aligned} q_{\text{d3pm}}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) &= \frac{q(\mathbf{x}_t|\mathbf{x}_{t-1})q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)} \\ &= \frac{(\mathbf{v}^T(\mathbf{x}_t)\mathbf{Q}_t\mathbf{v}(\mathbf{x}_{t-1}))(\mathbf{v}^T(\mathbf{x}_{t-1})\bar{\mathbf{Q}}_{t-1}\mathbf{v}(\mathbf{x}_0))}{\mathbf{v}^T(\mathbf{x}_t)\bar{\mathbf{Q}}_t\mathbf{v}(\mathbf{x}_0)} \end{aligned} \quad (3)$$

To plug in the prior physical transition, we modify the posterior term $q_{\text{d3pm}}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ in Equation 3, forming a hybrid conditional posterior as follows:

$$q_{\text{hyb}}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = (1 - \lambda_t)q_{\text{d3pm}}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) + \lambda_t q_{\text{phys}}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \quad (4)$$

However, $\mathbf{x}_0$ is unknown during inference, the proposed model is used to learn the conditional distribution $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, F(H))$ at the reverse process during training to approximate the posterior distribution $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$. The learned conditional distribution would be applied to unseen wildfire state transitions at the inference (sampling) stage. $F(H)$ demotes the embedded features of historical wildfire records, and l is the wildfire state transition type in the last year of the year need be predicted. Figure 3 shows the forward and reverse processes of the proposed WildfireSTDDM. Instead of directly learn the predicting the conditional distribution $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, F(H), l)$, we let the model predict the noiseless wildfire state transition categorical type mask distribution $p_\theta(\mathbf{x}_{0,t}|\mathbf{x}_t, F(H), l)$ at each timestep t for better generation quality (Yang et al., 2025), which is formulated as follows:

$$p_\theta(\mathbf{x}_{0,t}|\mathbf{x}_t, F(H), l) = \mathcal{D}_{WFST}(\mathbf{x}_t, t, F(H), l) \quad (5)$$

where, $\mathcal{D}_{WFST}$ represent the proposed WildfireSTDDM model. The predicted distribution is enforced to approximate the clean wildfire state transition categorical mask $q(\mathbf{x}_0)$. Based on the $p_\theta(\mathbf{x}_{0,t}|\mathbf{x}_t, F(H), l)$, we can compute the reverse transition distribution from mask $\mathbf{x}_t$ to $\mathbf{x}_{t-1}$ by

$$\begin{aligned} p_\theta^{\text{hyb}}(\mathbf{x}_{t-1}|\mathbf{x}_t, F(H), l) \\ = \sum_{\mathbf{x}_{0,t}} q_{\text{hyb}}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_{0,t}) p_\theta(\mathbf{x}_{0,t}|\mathbf{x}_t, F(H), l) \end{aligned} \quad (6)$$

2.4 Overview of WildfireSTDDM

Figure 4 shows the overview of the proposed WildfireSTDDM, which mainly consists of three parts, i.e., Encoder, HR-WS Ali Block (Historical Records-Wildfire State Alignment Block), and the Decoder part.

Since the discrete daily historical wildfire records contain a lot of background pixels (the left Ocean part, and the right Alberta province, -1 is set as the background pixels), and many zero-value pixels (No Fire area), we design a robust feature extractor before feeding it into the embedding layer. The "Conv" in Figure 4 (b) contains four 2D convolutional layers for extracting features, an adaptive weighted layer for adaptively fusing features, and a normalization layer for normalizing each feature map to be zero mean and unit variance, yielding the output robust image features, which can be formulated as follows:

$$f_i^{\text{out}}(HR) = (\sum_j w_{ij} (f_j^{\text{in}}(HR) \odot f_j^{\text{in}}(HR)))^{\frac{1}{2}} \quad (7)$$

where $f_j^{\text{in}} \in \mathbb{R}^{H \times W \times D}$ with D feature maps $\{f_j^{\text{in}}(HR)\}_{j=1}^D$. w_{ij} is a learnable parameter that adaptively updates the channel-wise weights, follows non-negative and sum to one. $\odot$ denotes the element-wise product, which will amplify the feature from wildfire state transition type (New ignition, Fire cessation, and Persistent fire) area, because these areas occur much more wildfires than other places, meaning that more weights need to be assigned to these interested area.




Figure 1. Illustration of the 3-year wildfire incidents datasets derived from FIRMS, from 2000 to 2024, and the corresponding next year wildfire state transition type (1, 2, 3, 4 means persistent no-fire (black), new ignition (red), fire cessation (green), persistent fire (blue), respectively). The curve on the top and right side shows the marginal probability density function (PDF) along the longitude and latitude. (a)-(o) are used as the training dataset, the rest as the testing datasets. Each dataset with shape (1097, 116, 250), 1097, 116, and 250 represent the total days in three years, height (latitude), and width (longitude), respectively.




Figure 2. Illustration of the transition direction across training date, from 2000 to 2017. (*i, j*) subfigure shows that the transition count spatially from state type *i* to state type *j*. The pixel-wise physical transition probability for each ecozone is also shown below. The global transition matrix is used in our PA-WildfireSTDDM.

The Embedding layer is used to perform the spatial aggregation, meanwhile increasing the feature dimension, the self-attention and linear projection layer are used to perform self-refinement.

The tailored designed HR-WS-Ali Block successively refines the noisy wildfire state transition mask embeddings based on conditional robust daily wildfire records feature under different discrete time step scales, followed by the decoder part to recover the original spatial size using upsampling methods for getting the noiseless wildfire state transition type. Specifically, this transformer block iteratively interacts among the patch-level embeddings of noisy wildfire state transition mask, robust daily wildfire records feature by using *N* multihead attention blocks, yielding the refined embeddings of noisy wildfire state transition mask, which can be formulated as follows:

$$\begin{aligned}
 f_{LWS} &= f_{LWS} + MHA(LN(f_{LWS}), f_{HR}) \\
 f_{WS} &= f_{WS} + MHA(AdaLN(f_{WS}, t)) \\
 f_{WS} &= f_{WS} + MHA(AdaLN(f_{WS}, t), f_{LWS}) \\
 f_{WS} &= f_{WS} + MHA(AdaLN(f_{WS}, t), f_{HR}) \\
 f'_{WS} &= f_{WS} + MLP(LN(f_{WS}))
 \end{aligned} \quad (8)$$

where LN, AdaLN, and MHA represent the LayerNorm, AdaptiveLayerNorm (also Time-Conditional Layer Normalization), and multi-head attention block, respectively. The feature alignment across different features can model the dependencies between historical records and wildfire state, then allowing to generate final refined noisy mask embeddings f'_{WS} .




Figure 3. Illustration of forward (green arrow) and reverse (orange arrow) discrete diffusion process of WildfireSTDDM, where gray pixel represents the additional category.

2.5 Model Training

The training objective function of reverse model is to let the learned conditional distribution $p_\theta(x_{0,t}|x_t, F(H), l)$ approximate the wildfire state transition mask $q(x_0)$ using cross entropy loss, and defined as:

$$\mathcal{L} = \mathbb{E}_{q(x_0)} \mathbb{E}_{q(x_{t-1}|x_t, x_{0,t})} [-\log p_\theta(x_{0,t}|x_t, F(H), l)] \quad (9)$$

2.6 Inference for Next Year Wildfire State Transition

Once the model training is finished, WildfireSTDDM can generate the next year's wildfire state transition type, as the orange arrow shown in Figure 3. Specifically, starting from uniform distribution with value 5 (mask token), the $F(H)$ is fed into WildfireSTDDM as conditions, and iteratively refine the noisy wildfire state transition masks in the chain of $x_T, x_{T-1}, \dots, x_1, x_0$ and model the conditional distribution $p_\theta^{\text{hyb}}(x_{t-\Delta t}|x_t, F(H), l)$. Here, we make our model skip some timesteps in the reverse process during model inference to reduce the time cost.

Since our method models the wildfire state transition as a stochastic process, which is suitable for uncertainty quantification through multiple Monte Carlo sampling, which is useful for environmental monitoring, because a single prediction may lead to some unwarranted results and raise risk (Ji and Tang, 2025, Li et al., 2024).

3. Results and Discussions

3.1 Results

We include both deterministic and stochastic Markov chain models as comparative baselines, which is straightforward because the wildfire evolution is inherently a state-transition problem $s_{t+1} \sim p(s_{t+1}|s_t)$. In order to distinguish from x_t , here we use the s_t to represent the wildfire state at year t . The deterministic Markov model represents a persistence-dominated transition rule using the maximum-probability state, while the stochastic Markov model samples transitions from the empirical transition matrix to account for aleatoric variability. These baselines provide physically interpretable references and allow us to distinguish improvements arising from uncertainty modeling and spatial-temporal learning from those due solely to learned transition statistics. Another baseline is the Naïve model, which assumes complete temporal persistence and predicts the wildfire state at year $t + 1$ to be identical to the observed state at year t (Beck et al., 2025).

To provide tolerance to spatial misalignment across different models, we evaluate performance under three spatial buffer sizes (1, 2, and 3 pixels). Table 1 presents the quantitative evaluation results for the wildfire state New Ignition. Precision for the Naïve baseline without a spatial buffer is not reported, as this baseline yields zero precision, recall, and F1 score when the




Figure 4. Illustration of our proposed WildfireSTDDM, which consists of Encoder, HR-WS Ali Block (Historical Records-Wildfire State Alignment Block), Decoder part. Decoder part is implemented by a simple upsampling method, which has not been shown in the Figure.

buffer size is 0 or 1. This behavior arises from the strict definition of the New Ignition state, where no spatial tolerance is allowed for prediction errors. When the spatial buffer size is 0 or 1, the deterministic Markov baseline achieves higher Precision than both the stochastic Markov model and PA-WildfireSTDDM. However, this apparent advantage is largely attributable to the evaluation protocol, in which the spatial buffer is applied only to the ground truth. Moreover, the ground-truth wildfire maps are spatially aggregated to a 10 km resolution during data pre-processing, ensuring high confidence in wildfire occurrence and reducing the likelihood of omitting small fires. Under these conditions, Precision alone does not adequately reflect model performance; instead, Recall provides a more meaningful assessment.

Notably, even without any spatial buffer (or equivalently with a buffer size of 1), PA-WildfireSTDDM consistently outperforms all comparison methods by a large margin in both Recall and F1 score. While ensemble prediction based on the stochastic Markov model improves upon the deterministic Markov chain, it remains insufficient to capture the inherent complexity and randomness of wildfire state transitions. In contrast, our model can be interpreted as an enhanced Markov framework that explicitly models a time-inhomogeneous stochastic process, thereby yielding a more faithful representation of wildfire dynamics. As the spatial buffer size increases, the Naïve baseline appears competitive; however, this static strategy fundamentally assumes that new wildfires emerge primarily around previously ignited regions. Such an assumption oversimplifies the underlying pro-

cesses governing wildfire occurrence and fails to capture complex spatial-temporal dynamics. Meanwhile, the stochastic Markov model exhibits decreasing Recall and F1 scores with larger buffer sizes, indicating that reliance on a time-homogeneous physical transition kernel P_{phy} is insufficient to describe inhomogeneous wildfire transitions. This limitation leads to increasing spatial discrepancies between predictions and observed wildfire events. In contrast, PA-WildfireSTDDM consistently produces predictions that are spatially closer to the ground truth, regardless of whether a spatial buffer is applied.

Figure 5 illustrates the wildfire state transition predictions produced by each comparative model. The deterministic Markov chain tends to over-predict Persistent Fire relative to New Ignition, which can be attributed to the imbalance in the column-wise probabilities of the physical transition matrix P_{phy} . Although the stochastic Markov model partially accounts for randomness in wildfire state transitions, it operates in a purely pixel-wise manner and neglects spatial dependence, leading to scattered and weakly aggregated prediction patterns. The Naïve baseline assumes no state change between consecutive years. While this assumption is reasonable for relatively stable states such as Persistent No-Fire and Persistent Fire, it is inadequate for the New Ignition state. Given the strong influence of dynamic climate conditions and fuel load, this static assumption inevitably introduces pronounced spatial bias and fails to capture emerging wildfire patterns. In contrast, PA-WildfireSTDDM explicitly incorporates local spatial context during training by treating each 4×5 patch as a basic processing unit. This design enables the model to learn local spatial correlations, resulting in predictions with clearer and more coherent aggregation structures. The corresponding zoomed-in visual comparisons are presented in Figure 6, which further support this analysis. As shown in the first row of Figure 6, PA-WildfireSTDDM produces more spatially aggregated and structurally consistent predictions, particularly for the New Ignition state. These predictions align more closely with the ground truth than those of the other three methods, as exemplified by the highlighted region on the right (yellow circle).

The second zoomed-in example in Figure 6 further demonstrates this aggregation behavior, underscoring the practical importance of accurately identifying such regions. Notably, this region is located adjacent to Alberta and the Northwest Territories—is of particular significance during the extreme wildfire year of 2023 and had not experienced severe burning in previous years (see Figure 1). In this challenging scenario, PA-WildfireSTDDM yields more spatially coherent New Ignition predictions with reduced mixing into the Fire Cessation state compared with the stochastic Markov model. Although discrepancies with the ground truth remain, the proposed model more clearly delineates emerging wildfire activity, highlighting its superior ability to capture complex, spatially structured wildfire dynamics.

The third zoomed-in example in Figure 6 shows the same region as the second zoomed-in example, but after an extreme wildfire event. Compared with other methods, our model effectively suppresses spurious new ignitions (red) in the post-wildfire period, while preserving coherent fire-cessation (green) and persistent no-fire (black) patterns, indicating a more realistic representation of post-wildfire dynamics.

From a statistics and probability perspective, we build some metrics to evaluate the performance of the proposed WildfireSTDDM. One is measuring the similarity of the probability density




Figure 5. Wildfire state transition prediction from 2019 to 2024 (left to right) on (a) Deterministic Markov, (b) Stochastic Markov, (c) Naïve, (d) Our proposed, and (e) Ground truth. Three yellow boxes show the location of ROI.




Figure 6. Comparison on ROI region level from 2021 (first row) and 2023 (second row) on (a) Deterministic Markov, (b) Stochastic Markov, (c) Naïve, (d) Our proposed, and (e) Ground truth. ROI from Figure 5.

function of each wildfire state transition type; MSE, KL divergence, and Wasserstein distance are used to measure the closeness. Another is to visualize and compare the generated ensemble samples with the ground truth, examining the correlation structure of generative samples spatially. Some related results are presented in Table 2. As we can see from this table, based on the distribution closeness calculation, our predicted results are similar with the ground truth on MSE and KL divergence along the longitude and latitude, new ignition and persistent fire type be more harder to be predicted than persistent no-fire and fire cessation on these three metrics.

As shown in Figure 7, the low-dimensional manifold of no-fire remains the same across the test dataset, and the proposed model can capture the persistent no-fire area (first column) with high consistency, meaning that the model provides with high confidence and low uncertainty when inferring the no-fire state type. The manifold feature of the other three wildfire state transition types represents different patterns, meaning that the dynamic changes of wildfire development. In all, the statistical distribution of fire cessation state transition type is clearly modeled by WildfireSTDDM; more dense orange dots are mostly located on the darker blue areas. By combining these two points, our model can better discriminate the negative samples. For the new ignition type (second column), the overall manifold structure is well captured, for example, (a), (d), (f). Note that Fig-

Table 1

Prediction performance on *New Ignition* among Deterministic Markov Baseline, Stochastic Markov Baseline, Naïve Baseline, and our proposed model with different spatial buffer size (0, 1, 2, 3). The **Best** and **second** are marked with color.

Wildfire state transition Year	Buffer Size	0/1				2				3			
		Deterministic Markov	Stochastic Markov	Naïve	PA-WildfireSTDDM	Deterministic Markov	Stochastic Markov	Naïve	PA-WildfireSTDDM	Deterministic Markov	Stochastic Markov	Naïve	PA-WildfireSTDDM
2019	Precision	0.1667	0.0356	-	0.0517	0.2500	0.1068	0.0705	0.1765	0.4583	0.2313	0.1774	0.3173
	Recall	0.0106	0.0265	-	0.0767	0.0238	0.1243	0.2169	0.4074	0.0397	0.2063	0.4074	0.4630
	F1	0.0199	0.0303	-	0.0618	0.0435	0.1149	0.1065	0.2159	0.0730	0.2181	0.2472	0.3765
2020	Precision	0.1452	0.0463	-	0.0375	0.2581	0.1142	0.1270	0.1203	0.4355	0.2253	0.2540	0.2426
	Recall	0.0407	0.0679	-	0.0860	0.0860	0.1674	0.1900	0.2670	0.1448	0.3167	0.3982	0.4796
	F1	0.0636	0.0550	-	0.0522	0.1290	0.1358	0.1522	0.1659	0.2173	0.2633	0.3101	0.3222
2021	Precision	0.2247	0.1019	-	0.0937	0.3708	0.2413	0.1946	0.2618	0.5506	0.3780	0.4480	0.4162
	Recall	0.0303	0.0576	-	0.0833	0.0682	0.1409	0.0742	0.2242	0.1242	0.2855	0.1667	0.4121
	F1	0.0534	0.0736	-	0.0924	0.1152	0.1779	0.1075	0.2416	0.2027	0.3317	0.2429	0.4141
2022	Precision	0.1322	0.0530	-	0.0514	0.2066	0.1082	0.0924	0.1505	0.3554	0.2141	0.2182	0.2800
	Recall	0.0518	0.0777	-	0.0874	0.0971	0.1845	0.1748	0.2654	0.1586	0.3010	0.3657	0.4693
	F1	0.0744	0.0630	-	0.0661	0.1321	0.1364	0.1209	0.1921	0.2193	0.2502	0.2733	0.3507
2023	Precision	0.2515	0.2015	-	0.1231	0.4371	0.3612	0.2104	0.2822	0.6287	0.5700	0.4628	0.4773
	Recall	0.0307	0.0599	-	0.0475	0.0774	0.1592	0.0657	0.1519	0.1271	0.2725	0.1454	0.2747
	F1	0.0547	0.0923	-	0.0685	0.1316	0.2210	0.1002	0.1975	0.2115	0.3687	0.2212	0.3487
2024	Precision	0.0878	0.0345	-	0.0648	0.1394	0.1164	0.1169	0.1370	0.2424	0.2309	0.2367	0.2519
	Recall	0.0360	0.0406	-	0.0748	0.0577	0.1517	0.2927	0.2244	0.1346	0.3419	0.5513	0.3803
	F1	0.0511	0.0373	-	0.0750	0.0816	0.1317	0.1671	0.1701	0.1731	0.2756	0.3312	0.3030

Table 2

The distribution characteristic between the predicted wildfire state transition and real data.

Marginal PDF	Wildfire Ignition Type	FIRMS 2016-2018		
Longitude	Persistent no-fire	1.67E-08	6.56E-04	1.32
	New ignition	4.53E-07	2.34E-02	2.34
	Fire cessation	1.49E-07	8.97E-03	1.39
	Persistent fire	4.41E-07	9.59E-03	2.01
Latitude	Persistent no-fire	2.00E-07	1.64E-03	1.48
	New ignition	1.90E-06	1.53E-02	1.84
	Fire cessation	2.46E-06	1.36E-02	2.05
	Persistent fire	2.01E-06	1.49E-02	1.76
FIRMS 2017-2019				
Longitude	Persistent no-fire	1.97E-08		
	New ignition	2.61E-07	1.46E-02	1.25
	Fire cessation	4.13E-07	1.43E-02	3.19
	Persistent fire	2.32E-06	7.72E-02	7.05
Latitude	Persistent no-fire	1.23E-07	8.34E-04	0.82
	New ignition	1.65E-05	1.46E-01	5.53
	Fire cessation	1.65E-06	1.12E-02	1.63
	Persistent fire	1.75E-05	1.41E-01	5.63
FIRMS 2018-2020				
Longitude	Persistent no-fire	2.15E-08	6.27E-04	0.41
	New ignition	1.04E-06	4.16E-02	5.51
	Fire cessation	1.08E-06	2.29E-02	5.08
	Persistent fire	5.55E-07	4.90E-02	2.43
Latitude	Persistent no-fire	4.14E-08	3.05E-04	0.29
	New ignition	2.71E-06	4.60E-02	4.26
	Fire cessation	1.13E-06	8.57E-03	2.62
	Persistent fire	2.24E-06	7.76E-02	2.05
FIRMS 2019-2021				
Longitude	Persistent no-fire	2.43E-08	5.55E-04	0.87
	New ignition	1.42E-06	1.10E-01	7.13
	Fire cessation	1.64E-07	1.00E-02	1.13
	Persistent fire	1.60E-06	5.99E-02	6.51
Latitude	Persistent no-fire	1.06E-07	7.23E-04	0.59
	New ignition	3.11E-06	2.75E-02	3.33
	Fire cessation	1.17E-06	1.21E-02	1.16
	Persistent fire	2.41E-05	1.36E-01	6.38
FIRMS 2020-2022				
Longitude	Persistent no-fire	1.01E-07	2.42E-03	1.67
	New ignition	4.07E-06	7.92E-02	4.1
	Fire cessation	5.96E-07	1.73E-02	3.17
	Persistent fire	7.35E-07	5.52E-02	3.14
Latitude	Persistent no-fire	4.91E-07	2.79E-03	1.98
	New ignition	4.09E-05	3.66E-01	20.6
	Fire cessation	1.17E-06	1.20E-02	1.56
	Persistent fire	7.75E-06	7.43E-02	4.61
FIRMS 2021-2023				
Longitude	Persistent no-fire	1.14E-07	2.80E-03	1.94
	New ignition	1.88E-06	5.93E-02	5.21
	Fire cessation	3.95E-07	9.74E-03	1.7
	Persistent fire	6.31E-06	9.24E-02	5.06
Latitude	Persistent no-fire	3.90E-07	2.26E-03	1.81
	New ignition	1.70E-05	1.38E-01	12.29
	Fire cessation	2.35E-06	2.16E-02	4.22
	Persistent fire	3.47E-05	2.97E-01	19.14




Figure 7. Statistical spatially coverage of ensemble predicted results (blue kernel density) and ground truth (orange dots).

Each column of (a)-(f) represents the persistent no-fire, new ignition, fire cessation, and persistent fire. The generated results are based on 50 times Mento Carlo sampling, and the kernel density estimation is performed on the voting results through Mento Carlo results according to the uncertainty.

truth. Some similar results can also be found in the persistent fire state type, the last column.

3.2 Discussions and Limitations

In this paper, we investigate the potential of predicting next-year wildfire state transitions using only historical wildfire records, through a generative diffusion framework. The results demon-

Figure 7 (f) is the prediction results of 2023, which is an extreme wildfire year in Canada. Our model can also preserve the correlation structure of generative samples spatially with the ground

strate that our model outperforms the naive persistence baseline, as well as both deterministic and stochastic models, particularly in predicting new ignition events during anomalous wildfire years.

To the best of our knowledge, this is the first attempt to model wildfire state transitions in this manner. However, several limitations remain. First, we have not yet explored a comprehensive set of evaluation metrics, which limits our ability to fully assess the strengths and weaknesses of the proposed approach. Second, due to the novelty of this task, there are currently no well-established comparative methods, making it difficult to construct a robust benchmark for evaluation. In addition, differences in overpass times between MODIS Terra (mid-morning) and MODIS Aqua (early afternoon) may lead to underestimation of fire activity.

In future work, we will investigate the impact of such data limitations on wildfire prediction performance. While this study relies on the MOD14A1 and MYD14A1 products, which are standard satellite-based fire datasets, other reliable data sources are available for Canada, such as the Canadian Fire Spread Dataset (Barber et al., 2024). Furthermore, higher-resolution fire perimeter data can be obtained from the Canadian Wildland Fire Information System (including Canadian National Fire Database and the National Burned Area Composite). In future studies, we plan to integrate these data sources to construct a more comprehensive dataset for wildfire state transition analysis. Additionally, we aim to incorporate more driving factors into the prediction framework and to systematically evaluate uncertainty in both burned and unburned areas under varying environmental and societal conditions.

4. Conclusion

In this paper, a new generative discrete diffusion model with physical regularization sampling processing is presented. We think of the wildfire prediction as a generative task and define a new prediction task by defining four wildfire transition types. This model not only provides the prediction of wildfire state transition in the next year, but also provides the uncertainty quantification, which can help decision makers to better understand the model behavior. Although it is a first attempt to view the wildfire prediction in a different way, in the future, we will build a more comprehensive, on-hand, and available dataset to have a better understanding of wildfire state transition.

References

- Abatzoglou, J. T., Williams, A. P., 2016. Impact of anthropogenic climate change on wildfire across western US forests. *Proceedings of the National Academy of Sciences*, 113(42), 11770-11775.
- Barber, Q. E., Jain, P., Whitman, E., Thompson, D. K., Guindon, L., Parks, S. A., Wang, X., Hethcoat, M. G., Parisien, M.-A., 2024. The Canadian fire spread dataset. *Scientific data*, 11(1), 764.
- Beck, N., Dovern, J., Vogl, S., 2025. Mind the naive forecast! a rigorous evaluation of forecasting models for time series with low predictability: N. Beck et al. *Applied Intelligence*, 55(6), 395.
- 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*.
- Giglio, L., J. C., 2015. MOD14A1 MODIS/Terra Thermal Anomalies/Fire Daily L3 Global 1km SIN Grid V006 [Data set]. *NASA Land Processes Distributed Active Archive Center*.
- Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B., 2022. Vector quantized diffusion model for text-to-image synthesis. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10696–10706.
- Huot, F., Hu, R. L., Ihme, M., Wang, Q., Burge, J., Lu, T., Hickey, J., Chen, Y.-F., Anderson, J., 2021. Deep learning models for predicting wildfires from historical remote-sensing data.
- Ji, C., Tang, H., 2025. Towards reliable land cover mapping under domain shift: An overview and comprehensive comparative study on uncertainty estimation. *Earth-Science Reviews*, 263, 105070.
- Knorr, W., Arneth, A., Jiang, L., 2016. Demographic controls of future global fire risk. *Nature Climate Change*, 6(8), 781-785.
- Li, F., Z. Q. Y. K. J. F. P. A. L. P. R. V. C. . C. M., 2024. Projecting Large Fires in the Western US With an Interpretable and Accurate Hybrid Machine Learning Method. *Earth's Future*, 12(10).
- Li, F., Zhu, Q., Riley, W. J., Zhao, L., Xu, L., Yuan, K., Chen, M., Wu, H., Gui, Z., Gong, J., Rander-son, J. T., 2023. AttentionFire_v1.0: interpretable machine learning fire model for burned-area predictions over tropics. *Geoscientific Model Development*, 16(3), 869–884. <https://gmd.copernicus.org/articles/16/869/2023/>.
- Li, L., Carver, R., Lopez-Gomez, I., Sha, F., Anderson, J., 2024. Generative emulation of weather forecast ensembles with diffusion models. *Science Advances*, 10(13), eadk4489.
- Littell, J. S., P. D. L. R. K. L. L. Y. . L. C. H., 2016. A review of the relationships between drought and forest fire in the United States. *Global Change Biology*, 22(7), 2353–2369.
- Taylor, S., Daniels, L., Copes-Gerbitz, K., Forbes, K., 2022. 1.5 Wildfires, in *Resilient Pathways Report: Co-creating new Knowledge for Understanding Risk and Resilience in BC*. Open File 8910, 102–129.
- Yang, H., Hua, W., Xu, Z., Sun, J., 2025. Domain-Generalized Discrete Diffusion Model for Cross-Domain Medical Image Segmentation. *IEEE Transactions on Medical Imaging*, 1-1.