

GeoPro-Net: Learning Interpretable Spatiotemporal Prediction Models Through Statistically-Guided Geo-Prototyping

Bang An^{2*}, Xun Zhou^{1*†}, Zirui Zhou¹, Ronilo Ragodos³, Zenglin Xu⁴, Jun Luo⁵

¹School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen

²Scheller College of Business, Georgia Institute of Technology

³Tippie College of Business, University of Iowa

⁴Artificial Intelligence Innovation and Incubation Institute, Fudan University

⁵Logistics and Supply Chain MultiTech R&D Centre

bang.an@scheller.gatech.edu, zhouxun2023@hit.edu.cn, 23s151098@stu.hit.edu.cn, ronilo-ragodos@uiowa.edu, zenglinxu@fudan.edu.cn, jluo@lscm.hk

Abstract

The problem of forecasting spatiotemporal events such as crimes and accidents is crucial to public safety and city management. Besides accuracy, interpretability is also a key requirement for spatiotemporal forecasting models to justify the decisions. Interpretation of the spatiotemporal forecasting mechanism is, however, challenging due to the complexity of multi-source spatiotemporal features, the non-intuitive nature of spatiotemporal patterns for non-expert users, and the presence of spatial heterogeneity in the data. Currently, no existing deep learning model intrinsically interprets the complex predictive process learned from multi-source spatiotemporal features. To bridge the gap, we propose GeoPro-Net, an intrinsically interpretable spatiotemporal model for spatiotemporal event forecasting problems. GeoPro-Net introduces a novel Geo-concept convolution operation, which employs statistical tests to extract predictive patterns in the input as “Geo-concepts”, and condenses the “Geo-concept-encoded” input through interpretable channel fusion and geographic-based pooling. In addition, GeoPro-Net learns different sets of prototypes of concepts inherently, and projects them to real-world cases for interpretation. Comprehensive experiments and case studies on four real-world datasets demonstrate that GeoPro-Net provides better interpretability while still achieving competitive prediction performance compared with state-of-the-art baselines.

Introduction

The spatiotemporal event forecasting problem aims at predicting where and when certain features or events will occur in the geographic space and time. The applications of spatiotemporal event forecasting such as crime prediction and traffic accident forecasting hold significant importance in various domains, including public safety (CPD 2020) and traffic management (Gov 2017). While it is important to make accurate predictions, understanding the predictive process of models in spatiotemporal event forecasting is, if not more, equally important to a wide range of urban stakeholders. This approach

assists in justifying government expenditure on unforeseen risks, thereby contributing to more informed decision-making in urban contexts.

Interpreting spatiotemporal event prediction models is a challenging task. Firstly, spatiotemporal event prediction often involves multi-source input, such as traffic volume, precipitation, and point-of-interests, with complex dependencies and semantics. Interpreting their interplay over space and time is a non-trivial task. Secondly, unlike other problems such as image classification, where a picture can be naturally understood by humans, visualizing high-dimensional spatiotemporal features with post-hoc analysis on a map (e.g., saliency map) is not readily understandable by most individuals without domain knowledge. Thirdly, spatial heterogeneity (Atluri, Karpatne, and Kumar 2018) commonly exists in spatiotemporal event datasets, which implies that the rationale behind the event occurrences might vary significantly from place to place. It poses difficulty in interpreting the modeling process within a heterogeneous space. These challenges call for an *intrinsically interpretable* model specifically designed for *spatiotemporal event forecasting problems*.

However, to the best of our knowledge, **there has not been any intrinsically interpretable deep learning model for the spatiotemporal event forecasting problem**. *On the spatiotemporal event model side*, early machine learning methods (Drukker, Prucha, and Raciborski 2013; Brunson, Fotheringham, and Charlton 1996; Jiang 2015) leveraged regression-based or tree-based models for interpretability but rely on manual feature augmentation, therefore, struggle to capture complex spatiotemporal dependencies in the data. Attention-based spatiotemporal event forecasting methods (Liu et al. 2020; Guo et al. 2019) generate visualizations of scores indicating relevance between input and target locations but fall short of offering insights into the underlying reasons for events at the feature level. *On the interpretable machine learning side* prototype-based explainers (Chen et al. 2019) are both intrinsically interpretable and naturally understandable by humans when used for image classification. However, visualizing multi-source spatiotemporal features in the same way makes it difficult to be understood. To strengthen the interpretability for abstract and complex features, concept-

*These authors contributed equally.

†Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

based models (Koh et al. 2020a) have been introduced to extract high-level information from inputs, making it easier to understand. However, training data must be manually annotated with predefined but interpretable concepts.

In this paper, we bridge the research gap by proposing the first intrinsically interpretable deep learning model for spatiotemporal event prediction, named **GeoPro-Net**. As a classifier, GeoPro-Net predicts if an event will occur at a given geo-location and time based on the spatiotemporal features observed around the location. To unravel the complexity of multi-source spatiotemporal features and transform them into comprehensible insights, we introduce a statistically guided retrieval process to identify multi-scale patterns such as local and global anomalies, thereby filtering noisy spatiotemporal information and alleviating the interpretation difficulty. Furthermore, we incorporate the idea of concepts (Koh et al. 2020a) and structuralize the captured spatiotemporal event patterns into geographic-based concepts (Geo-concepts) via a novel multi-scale hierarchical pooling mechanism. This enables non-expert users to easily understand the extracted and detailed spatial information. Lastly, we learn a set of prototypes, which are representative and predictive combinations of Geo-concepts associated with the two output classes. These prototypes generate predictions through a linear layer and can be projected onto cases in the training set, making predictions interpretable with real-world examples. Comprehensive experiments and case studies on four real-world event datasets in Chicago and New York City demonstrate the faithful interpretability and better prediction accuracy of GeoPro-Net in comparison with other baselines. Our main contributions are summarized below:

- To the best of our knowledge, this is the first prototype-based interpretable framework to learn from spatiotemporal event datasets and explain its reasoning process intrinsically.
- We propose a novel spatial conceptualization process to extract Geo-concepts, which are statistically interpretable multi-scale input feature patterns, followed by condensing concepts through channel fusion and geographic-based multi-scale hierarchical pooling.
- We propose to use a prototype learning framework to obtain a set of Geo-concepts associated with the occurrence or absence of events. By projecting these prototypes into training cases, we interpret predictions in the context of real-world scenarios.

Related Work

Interpretable Spatiotemporal Event Forecasting Models traditionally rely on regression based methods (ex. Spatial Autoregressive Model (Brunsdon, Fotheringham, and Charlton 1996) and Geographically Weighted Regression (Drukker, Prucha, and Raciborski 2013)) and tree-based methods (ex. Spatial Decision Tree (Jiang 2015)) on small-scale datasets with limited features. Those methods are inherently explainable via their learned coefficients or tree structures. However, their accuracy is largely limited when complicated spatiotemporal dependencies are involved. Deep learning approaches with enormous parameters (Lv et al. 2019; Yao et al. 2018;

Zheyi Pan 2019; An et al. 2023; Meng et al. 2019; Zuo et al. 2021) has been proven effective with the superior performance. However, most of the existing spatiotemporal event forecasting deep learning approaches are either black-box models or only partially interpretable. For example, xGAIL (Pan et al. 2020) was proposed to use GAIL-based methods (Ho and Ermon 2016) to draw insights into the decision-making process of taxi drivers. Furthermore, attention-based methods (Guo et al. 2019; Liu et al. 2020; Tang, Xia, and Huang 2023; Ding et al. 2020) are used to interpret the predictive process of forecasting urban events using spatial attention matrix and saliency maps.

Prototype-based Interpretable Models (Snell, Swersky, and Zemel 2017; Chen et al. 2019) have brought more attention in recent years. Most existing prototype-based interpretable models (Chen et al. 2019; Nauta, van Bree, and Seifert 2021) are limited to image problems, where prototypes are projected as image tensors with RGB channels. Such prototypes are naturally human-understandable. For example, in a bird classification problem, people can easily distinguish different species parts by comparing birds' color, fur, paws, etc. However, such multi-dimensional prototype representations in urban incident problems are not directly interpretable from a human's perspective because a set of traffic speed, volume, and occupancy changing over time and space is not intuitive and self-explainable.

Concept-Based Interpretable Models label their inputs with high-level concepts. Koh et al. (Koh et al. 2020b) applied this approach by using bone spurs to predict arthritis in X-ray images. These concepts can either be predefined before training or learned during the training process. Essentially, concepts serve as extracted summaries of complex information. The systematic definition of necessary concepts is crucial, as having too few concepts may limit model performance, while an excessive number of concepts can compromise model interpretability.

Preliminaries

A *spatial-temporal field* $S \times T$ is a three-dimensional matrix, where $T = \{t_1, t_2, \dots, t_i\}$ is a study period divided into equal length intervals (e.g., hours, days) and $S = \{s_{(0,0)}, s_{(0,1)}, \dots, s_{(m,n)}\}$ is a $m \times n$ two-dimension spatial grid partitioned from the study area. Temporal features $F_T \in \mathbb{R}^{T \times f_t}$ (e.g., average temperature, day of week), spatial features $F_S \in \mathbb{R}^{m \times n \times f_s}$ (e.g., number of POIs), and spatio-temporal features $F_{ST} \in \mathbb{R}^{T \times m \times n \times f_{st}}$ (e.g., traffic volume) are mapped to $S \times T$ field. We denote the total number of features as $f = f_s + f_t + f_{st}$. The detailed list of features, processing steps, and symbol table in this work can be found in Appendix A.

Problem Definition: Given the socio-environmental features F_T, F_S, F_{ST} of a location l_m and its neighboring locations (e.g., within the distance of r grid cells) in time window $t \in T$, our problem is to **Predict** whether an event will occur in the future one-step interval t_{i+1} for location l_m . The **Objective** is to interpret the predictive process while minimizing prediction errors. In this spatiotemporal event forecasting problem, we work with a basic assumption commonly adopted by prior work that there exists spatial auto-

correlation over space, meaning nearby locations tend to have more correlated values (Tobler 1970).

Methodology

In this section, we present our GeoPro-Net model to intrinsically interpret the predictive process on the defined spatiotemporal event forecasting problems. Figure 1 demonstrates the proposed model architecture.

Statistically-guided Geo-Concept Encoding

In cognitive science, the simplicity principle states that the mind tends to seek the simplest available interpretation of observations over complex ones (Feldman 2016). It becomes crucial to extract the most representative information from extensive feature values and consolidate them into basic concepts to facilitate the interpretation of the modeling process. To address such challenges, we design a novel Statistically-guided Spatial Concept Encoding layer (SSCE) as demonstrated in the black dashed box of Figure 1.

Definition 1. Local Significance Test Ψ_L . Given a training sample $\mathcal{X} \in \mathbb{R}^{m \times n \times t \times f}$, and a test window $\Lambda \in \mathbb{R}^{d \times d}$, which is a spatial, temporal, or spatiotemporal sub-region of \mathcal{X} , $\Psi_L(\Lambda, F_i) \rightarrow \{True, False\}$ is a local statistical test function for feature F_i on whether the distribution in Λ is **different** from the distribution of F_i in the entire \mathcal{X} with α -level statistical significance with a higher or lower expected value, where α is a hyperparameter.

For example, a local significance test Ψ_L can be defined as a Poisson likelihood ratio test (Jung, Kulldorff, and Klassen 2007) on the number of POIs in a $d \times d$ region of an input sample against the distribution of the number of POIs in all the $d \times d$ sub-regions in the entire sample. An output of “True” for the test suggests that this $d \times d$ region might be a local hotspot of POIs near the location to be classified. Similarly, we can define a Global Significance Test as follows.

Definition 2. Global Significance Test Ψ_G . Given a training sample $\mathcal{X} \in \mathbb{R}^{m \times n \times t \times f}$, and a test window Λ , which is a spatial, temporal, or spatiotemporal sub-region of \mathcal{X} , $\Psi_G(\Lambda, F_i) \rightarrow \{True, False\}$ is a global statistical test function for feature F_i on whether the distribution in Λ is statistically **different** from the global distribution of F_i in the entire training set with a higher or lower expected value, with α -level significance.

Here, the number of dimensions of the window Λ depends on the type of feature F_i , i.e., one, two, or three dimensional for temporal, spatial, and spatio-temporal features. The choice of the statistical tests can be determined by the user based on the nature of the features. Typically, in the urban event forecasting problem, Poisson distributions (Haight 1967) are often preferred to estimate the distribution of events, such as counts of taxi-up and traffic volume, which are spatiotemporal features in our dataset. In this paper, we demonstrate the model using a Poisson likelihood ratio test (Vaheidian et al. 2019) for discrete count features. Other choices of statistical tests such as the student t-test can be used for data commonly assumed to follow a normal distribution. For features without a known or preferred distribution, we use

the non-parametric Kolmogorov–Smirnov test as an example (Wiesen 2019). The Kolmogorov–Smirnov (K-S) test is based on the empirical distribution function. It evaluates the hypothesis that the data comes from a specific distribution against the alternative hypothesis that the data do not follow this specific distribution.

$$E_N = \frac{n(i)}{N} \quad (1)$$

where $n(i)$ is the number of points less than Y_i . The value of Y increases with higher i . The Kolmogorov–Smirnov test statistic is defined as,

$$\max_{1 \leq i < N} (F(Y_i) - \frac{i-1}{N}, \frac{i}{N} - F(Y_i)), \quad (2)$$

where Y is the value and F is the theoretical cumulative distribution of the distribution being tested. By performing two-tail tests we can test whether a feature value in a sample is significantly higher or significantly lower than the specified baseline distribution. Note the distribution assumptions vary over different problems, and they are user-defined choices that are replaceable in our solution framework.

Now we are ready to define Geo-concepts and the concept convolution operation in GeoPro-Net.

Definition 3. Geo-concept Θ . Given an input sample $\mathcal{X} \in \mathbb{R}^{m \times n \times t \times f}$, a Geo-concept Θ of \mathcal{X} is a tuple $\langle \Lambda, \Psi, F_i \rangle$, where Λ is a sub-region of the spatiotemporal area of \mathcal{X} , Ψ is a global or local test on feature F_i in Λ , which has an output of “True”. Geo-concept is designed to focus on spatiotemporal interoperability, thus the temporal dimension will be flattened by the mean to reduce the total number of generated concepts.

Geo-concepts encode statistically significant patterns of input features at both local and global scales in each input sample. We design a Statistically guided Spatial Concept Encoding (SSCE) to extract these concepts and convert the input sample into a multi-channel map tensor that records the positions in the input where these Geo-concepts are valid.

The process of transforming input features into Geo-concept encoding is done through a set of “Concept Convolution” operations. From the left part of the black dashed box in Figure 1, the Statistically-guided Spatial Concept Encoding (SSCE) starts with scanning the samples from the training set across the study area using multiple user-defined window sizes and tests, which is similar to the workflow of the convolutional operation, except that there are no “parameters” to train. Instead, users need to specify the choice of statistical tests and the scan window dimensions.

A concept convolution operation with concept $\Theta = \langle \Lambda_\Theta, \Psi_\Theta, F_\Theta \rangle$ on an input dataset \mathcal{X} can be defined as follows, where $I()$ is an indicator function with an output of 0 or 1:

$$ConceptConv_{\Theta} * \mathcal{X}(\Lambda_\Theta, F_\Theta) = I(\Psi_\Theta(\Lambda_\Theta, F_\Theta)) == True \quad (3)$$

With the use of sliding windows of Λ , a concept convolution can be applied continuously over the entire input spatiotemporal area. Since multiple window sizes can be used, each input area could be scanned multiple times to capture

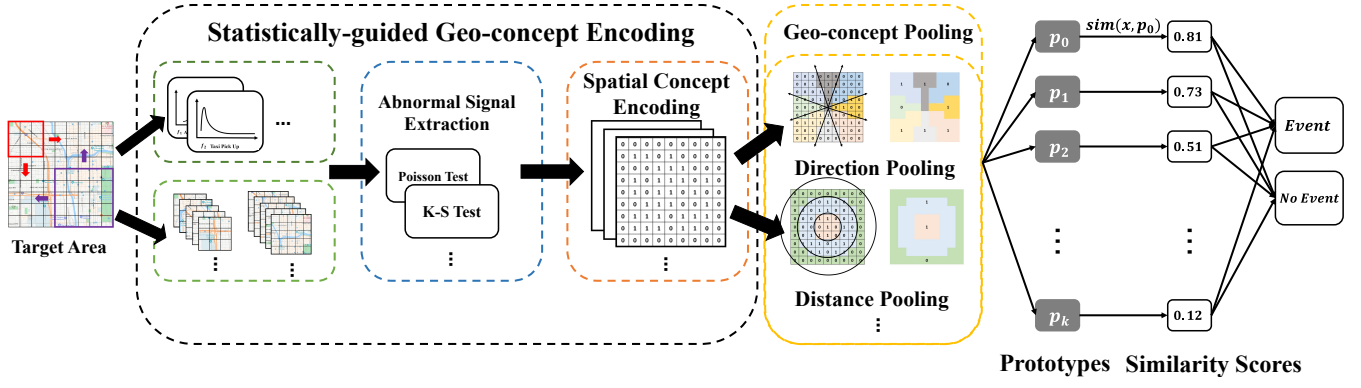


Figure 1: The overall architecture of GeoPro-Net. On the left side, signals are extracted through statistical tests applied to values over features, and then they are mapped into Geo-concepts within the study area. The obtained concepts are further selected via geographical pooling by selected approaches in the middle of the figure. On the right side, a prototype-based framework is integrated to learn the relationships between the occurrence of events and different sets of concept prototypes

patterns of various scales. Taking the spatial features F_s of an input sample \mathcal{X} with $m \times n$ spatial dimensions as an example, the original input \mathcal{X} will be encoded into a tensor with dimensions in $\mathbb{R}^{m \times n \times f_s \times 2 \times 4}$, where \mathbb{R}^4 indicates the global and local tests Ψ on feature F_i in Λ , which has an output of “True”. To scan locations on the edges of the input area, padding is applied to these locations with zeros. Spatiotemporal features are encoded in the same way except flatten \mathbb{R}^t by averaging overtime in the beginning.

Formally, the encoder E can be defined as a function over input $\mathcal{X} \in \mathbb{R}^{m \times n \times (f_{st} + f_s)}$: $E(\mathcal{X}) \rightarrow C \in \{0, 1\}^{m \times n \times (f_{st} + f_s) \times \omega \times 4}$ for F_S and F_{ST} , which generates a binary concept map tensor C , ω is the number of different concept convolution window sizes for each feature. The last dimension of C indicates four combinations of global vs local tests and higher or lower expected values in the test. Temporal features are encoded in the same way but removing spatial dimensions $\mathbb{R}^{m \times n}$.

The use of multiple concept window sizes could capture patterns of varying scales. These concepts may overlap and contain redundant information as a result. We design a Channel Fusion Operation and Multi-scale geo-pooling on extracted raw Geo-concepts to address these issues, which is detailed in the next subsection.

Geo-Concept Aggregation

Geo-concept Channel Fusion In GeoPro-Net, multiple sets of concepts are generated on each input sample by corresponding scanning windows, which increases the amount of extracted concepts and poses challenges for interpretation. To address these issues, we design a Channel Fusion Operation on the different channels of the concept encoding C to consolidate and refine the concept encoding into more concise but still interpretable representations. Especially, for each window size of the same test (i.e., channels), we learn a weight matrix w_i with the same dimensions as each raw concept encoding channel and combine the channels as a

weighted some of the features. Formally,

$$C'_{F,\Theta} = \sum_{i=1}^{\omega} C_{F,\Theta,\Lambda_i} * w_i, \text{ where } \sum_{i=1}^{\omega} w_i = \mathbf{1}^{m \times n} \quad (4)$$

Here C_{F,Θ,Λ_i} represents the encoded input channel for feature F convoluted by concept Θ with a scanning window of Λ_i , which is combined into a single channel for feature F and Θ , denoted as $C'_{F,\Theta}$, and $*$ is the Hadamard product. Eventually, we obtain a fused concept tensor $C' \in \mathbb{R}^{(f_s + f_{st}) \times m \times n \times 4}$.

Geo-concept Pooling. To further enhance the semantic representation of extracted concepts and reduce the encoding dimensions, we design a Geo-concept pooling layer with user-defined pooling strategies. The dashed yellow box in Figure 1 illustrates the pooling process. Depending on the target problem, the study area can be partitioned into sub-regions, and the mean or average of encoding C' is calculated from each pooled sub-region. In this problem, urban events are highly associated with nearby conditions (Tobler 1970). Therefore, we partition the study area into near, middle, and far sub-regions based on the Euclidean distance between each grid cell and the target location at the center of the sample. Meanwhile, we partition the study area into different directions based on the degree between the coordinates of each location and the central target location. Finally, the concept tensor is reshaped to $C' \in \mathbb{R}^{f' \times q \times 4}$, where $f' = (f_s + f_{st}) \times q + f_t$ and q represents the number of different pooling regions and 4 indicates the statistical test types. The pooling layer maintains the interpretability of the model. A positive value in a pooled region indicates the presence of at least one statistically significant pattern, while the magnitude of the pooled values indicates the overall strength of such presences. For example, an average-pooled value of 0.2 over a sub-region with 10 grid cells indicates at least two instances of statistically significant patterns in this sub-region.

Concept Prototype Layer

In GeoPro-Net, the prototype layer is illustrated on the right part of Figure 1. The prototype layer is a learnable parameter

$P' \in \mathbb{R}^{K \times f' \times q \times 4}$, where each prototype is a vector and K is the number of prototypes in our model. K is a hyperparameter and can be adjusted. The concept prototypes share the same dimension as pooled Geo-concept encoding and will be learned to be close to a subset of pooled concepts during training. The prototype learning is related to case-based classifications (Priebe 2003)(Chen et al. 2019)(Bien and Tibshirani 2011), and K representative prototypes are learned to be compared with unseen testing samples. Given a Geo-concept encoded testing sample $C \in \mathbb{R}^{f' \times q \times 4}$, the squared L^2 distances to a set of prototypes $P' \in \mathbb{R}^{K \times f' \times q \times 4}$ are inverted into K similarity scores,

$$Sim(C, p_k) = \frac{1}{\|C - p_k\|_2^2} \quad (5)$$

where p is the k^{th} learned prototype vector. Lastly, all similarity scores are connected to classifications, distinguishing the occurrence of an event versus no event, through a fully connected layer. The coefficients of connections indicate the linear correlation between prototypes and their corresponding class.

Optimization

The training of GeoPro-Net requires diversifying and disambiguating the learned prototypes to ensure representability and interpretability. This section outlines three regularization terms and a prototype projection process to achieve this goal. Related prototype-based models (Chen et al. 2019)(Yao et al. 2019) demonstrate the effectiveness of diversifying and disambiguating prototypes by using such regularization. Specifically, the *Dlv* (Diversity) term encourages model learning different prototypes, *Sep* (Separation) term pushes encoded concepts vector to stay away from the prototypes not of its class, and *Clst* (Cluster) encourages encoded concepts vector to be close to the prototype of its class.

$$Dlv = \min_{p_i \in P, p_j \in P} \frac{1}{|P|^2} \|p_i - p_j\|_2^2, \quad (6)$$

$$Sep = -\frac{1}{n} \sum_{i=1}^n \min_{p_j \notin P_{y_i}} \min_{z \in (f(x_i))} \|z - p_j\|_2^2, \quad (7)$$

$$Clst = \frac{1}{n} \sum_{i=1}^n \min_{p_j \in P_{y_i}} \min_{z \in (f(x_i))} \|z - p_j\|_2^2. \quad (8)$$

This approach simplifies interpretation by presenting only an appropriate number of significant concepts for each prototype, making the information more comprehensible and interpretable. The overall loss function for training our model is defined as minimizing

$$\mathcal{L} = \frac{1}{n} \sum_{i=1}^n \text{CrsEnt}(g(x_i), y_i) + \lambda_1 Dlv + \lambda_2 Sep + \lambda_3 Clst \quad (9)$$

where the $g()$ is the learned model, and CrsEnt represents the cross-entropy loss,

Complexity Analysis GeoPro-Net achieves practical scalability with highly efficient inference. The initial preprocessing involves calculating local and global feature distributions

via Local Significance Test Ψ_L and Global Significance Test Ψ_G across multiple window sizes ω over a spatial grid of size $m \times n$, with complexity $O(\omega \times m \times n \times t \times f)$ per feature F_i . Though computationally costly during training, the historical baseline only needs to be calculated once and can be reused during inference, making Concept Convolution more efficient. Once training data is encoded, geo-concepts achieve through Channel Fusion and Geo-concept Pooling, which refines the concept map C by consolidating channels $\theta' = \sum \omega \times C(F, \theta, \Lambda)$ and aggregating pooled sub-regions, resulting in complexity of $O(f' \times m \times n)$. The resulting Geo-concept tensor, a list of vectors, enables efficient prototype matching. Geo-concept encodings are matched with learned prototypes P , yielding in an efficient complexity of $O(K \times f' \times q)$.

Concept Prototype Projection is a post-training process. To interpret the learned concept prototypes, they are projected back into encoded Geo-concepts from the training set. Each prototype represents a real scenario by demonstrating a sample of encoded concepts. Specifically, we find the sample of encoded concepts with maximum similarity score to every prototype. This set of extracted concepts can be visualized in an easily digestible manner by listing only the most significant ones

Experiments

We perform comprehensive experiments using four real-world traffic accident and crime datasets from Chicago¹ and New York City². The experiments and case studies demonstrate that GeoPro-Net achieves better interpretability than interpretable competitors and higher performance than Black-box baselines. The results of experiments in New York City and case studies on other interpretable baselines can be found in the supplementary materials.

Experiment Settings

For the Chicago dataset, the time frame spans from 2019 to 2021, with the initial 18 months serving as the training set and the entirety of 2021 as the testing set. The area of Chicago is a grid with the size of 64×80 . We evaluate the performance of the models using the following quantitative measures: Cross Entropy (CrsEnt), Accuracy (ACC), Precision (Pre), Recall, and F1.

Baselines

First, we consider popular machine-learning methods including **ConvLSTM** (Shi et al. 2015), and **DCRNN** (Li et al. 2018). Secondly, we compare with recent methods such as **Hetero-ConvLSTM** (Yuan, Zhou, and Yang 2018), **GSNet** (Wang et al. 2021), and **DSTPP** (Yuan et al. 2023). Lastly, we compare with interpretable machine learning baselines including **ProtoPNet** (Chen et al. 2019), **ViT** (Dosovitskiy et al. 2020) **AGL-STAN** (Sun et al. 2022), and **ASTGCN** (Guo et al. 2019)

¹ <https://data.cityofchicago.org/Transportation/Traffic-Crashes-Crashes/85ca-t3if>

² <https://opendata.cityofnewyork.us/>

| NYC | CRIME | | | | | ACCIDENT | | | | |
|----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| | CRSENT | ACC | PRECISION | RECALL | F1 | CRSENT | ACC | PRECISION | RECALL | F1 |
| CONVLSTM | .62±2.6% | .76±0.8% | .51±2.5% | .19±0.8% | .33±4.4% | .60±1.6% | .73±7.7% | .40±3.8% | .19±1.0% | .25±0.9% |
| DCRNN | .53±1.6% | .77±3.1% | .66±1.1% | .27±0.6% | .40±1.8% | .56±1.8% | .75±4.1% | .49±2.0% | .26±0.7% | .36±1.1% |
| HETEROCONVLSTM | .52±2.6% | .78±3.5% | .65±1.6% | .28±1.0% | .42±1.1% | .55±1.5% | .74±2.0% | .50±2.3% | .27±0.6% | .35±1.3% |
| DSTPP | .57±0.7% | .76±1.1% | .60±0.5% | .22±0.6% | .36±0.7% | .58±0.9% | .72±0.8% | .44±2.0% | .25±1.2% | .29±1.7% |
| ViT | .54±1.3% | .77±2.6% | .65±1.9% | .26±0.5% | .39±1.4% | .58±1.5% | .74±3.7% | .46±2.7% | .27±0.9% | .35±0.3% |
| GSNET | .56±0.3% | .77±2.4% | .53±1.4% | .25±0.5% | .35±0.7% | .57±1.4% | .73±0.7% | .46±1.8% | .24±0.8% | .31±1.3% |
| AGLSTAN | .56±0.9% | .77±1.1% | .56±0.5% | .25±0.8% | .37±1.1% | .57±0.3% | .74±1.1% | .47±0.4% | .26±0.9% | .33±1.7% |
| ASTGCN | .57±0.3% | .76±0.4% | .55±1.0% | .26±0.4% | .37±0.6% | .57±0.3% | .74±0.9% | .48±0.5% | .27±0.4% | .34±0.6% |
| PROTOPNET | .58±0.0% | .77±0.4% | .59±0.1% | .22±0.2% | .36±0.3% | .59±0.9% | .73±0.5% | .43±0.1% | .20±0.1% | .28±0.3% |
| GEOPRO-NET | .53±0.5% | .77±0.4% | .64±0.4% | .24±0.0% | .39±0.3% | .57±0.6% | .74±0.7% | .47±0.6% | .26±0.2% | .33±0.3% |
| GEOPRONET | .51±0.3% | .79±0.1% | .68±0.4% | .27±0.2% | .43±0.2% | .54±0.2% | .76±0.2% | .50±0.6% | .28±0.2% | .37±0.4% |

| CHICAGO | CRIME | | | | | ACCIDENT | | | | |
|----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| | CRSENT | ACC | PRECISION | RECALL | F1 | CRSENT | ACC | PRECISION | RECALL | F1 |
| CONVLSTM | .59±4.4% | .75±6.3% | .46±7.2% | .17±0.6% | .21±0.8% | .61±5.6% | .72±3.1% | .30±1.5% | .18±1.6% | .23±1.1% |
| DCRNN | .53±1.4% | .78±4.1% | .52±2.3% | .23±0.9% | .25±0.6% | .54±2.9% | .74±4.6% | .39±2.7% | .27±1.1% | .34±1.6% |
| HETEROCONVLSTM | .50±1.4% | .79±1.2% | .57±2.0% | .24±0.4% | .28±1.3% | .54±1.4% | .73±2.0% | .40±4.4% | .29±1.4% | .32±2.2% |
| DSTPP | .55±1.9% | .76±1.8% | .53±0.8% | .21±0.9% | .24±0.8% | .58±2.2% | .72±2.9% | .36±1.9% | .21±1.2% | .27±1.5% |
| ViT | .53±0.9% | .77±2.7% | .54±1.7% | .23±1.3% | .24±1.1% | .55±2.8% | .73±3.2% | .39±2.3% | .26±1.6% | .33±0.9% |
| GSNET | .51±1.6% | .78±1.5% | .54±0.9% | .23±0.7% | .28±0.5% | .56±2.1% | .73±3.5% | .36±1.4% | .22±1.0% | .28±1.3% |
| AGLSTAN | .53±1.2% | .78±1.8% | .52±0.5% | .24±0.9% | .27±1.1% | .56±1.1% | .72±0.3% | .35±2.1% | .21±1.2% | .28±0.8% |
| ASTGCN | .53±1.4% | .78±2.4% | .53±0.8% | .23±0.5% | .26±0.8% | .57±0.9% | .73±0.3% | .34±1.3% | .23±0.2% | .30±0.7% |
| PROTOPNET | .56±0.5% | .78±0.7% | .50±0.4% | .19±0.2% | .23±0.2% | .59±0.5% | .72±0.6% | .38±0.3% | .21±0.3% | .27±0.4% |
| GEOPRONET* | .53±0.8% | .78±0.1% | .50±0.4% | .21±0.2% | .24±0.2% | .56±0.5% | .73±0.6% | .36±0.3% | .24±0.2% | .29±0.3% |
| GEOPRONET | .49±0.3% | .80±0.7% | .55±0.8% | .25±0.4% | .30±0.3% | .53±0.5% | .76±0.5% | .41±0.1% | .29±0.3% | .34±0.3% |

%: $\times 10^{-3}$

*: REPLACE TWO CONCEPT CONVOLUTION WINDOWS BY A SINGLE 3×3 WINDOW

Table 1: Performance Comparison

Evaluation on Model Performance

In table ??, we can observe that GeoPro-Net achieves better performance compared with its Blackbox competitors in most measurements from four datasets. Specifically, ConvLSTM performs worst due to its preference for precipitation now-casting problems rather than urban event datasets used in the experiments. The state-of-art Blackbox models such as DCRNN and Hetero-ConvLSTM benefit from complex network design and perform well compared to interpretable baselines. ProtoPNet and ViT, designed for image classification problems, struggles with handling the spatiotemporal features, resulting in comparatively worse performance. GeoPro-Net, with its simple prototype-based network structure, demonstrates stable performance and outperforms other baselines. We conduct an ablation study on the channel fusion part. GeoPro-Net with two concept convolution windows performs better than the model with one window.

Ablation Study

We examine the effects of removing spatial pooling or replacing with max pooling in the GeoPro-Net. Table. 2 shows that GeoPro-Net achieves optimal performance with spatial concept pooling. Replacing spatial pooling by max pooling slightly lowers the performance and loses the capability to interpret based on geographic information. Removing the pooling layer results in a dramatic decrease in its performance due to noisy information.

| | CRIME | | | | |
|------|-----------------|-----------------|-----------------|-----------------|-----------------|
| | CRSENT | ACC | PRECISION | RECALL | F1 |
| No" | .67±2.5% | .71±1.4% | .39±2.4% | .17±0.5% | .23±1.2% |
| Max" | .51±0.8% | .80±0.7% | .52±0.8% | .23±0.5% | .28±1.2% |
| Spa* | .49±0.3% | .80±0.7% | .55±0.8% | .25±0.4% | .30±0.3% |

| | ACCIDENT | | | | |
|------|-----------------|-----------------|-----------------|-----------------|-----------------|
| | CRSENT | ACC | PRECISION | RECALL | F1 |
| No" | .69±3.7% | .67±3.6% | .28±0.4% | .21±0.6% | .26±0.8% |
| Max" | .55±0.5% | .73±1.1% | .39±0.6% | .28±0.2% | .33±0.5% |
| Spa* | .53±0.5% | .76±0.5% | .41±0.1% | .29±0.3% | .34±0.3% |

*: SPAITAL CONCEPT POOLING ' : $\times 10^{-3}$
": Pooling

Table 2: Ablation Study

GeoPro-Net Interpretability

In this section, We identified an accident case that happened on January 15, 2021, Chicago. Figure 2 illustrates the explanations by GeoPro-Net for the cases in Chicago. In this particular case with an event, the associated concepts include high traffic volumes and crowd flows, which are similar to the concepts of positive prototypes. The higher similarity score

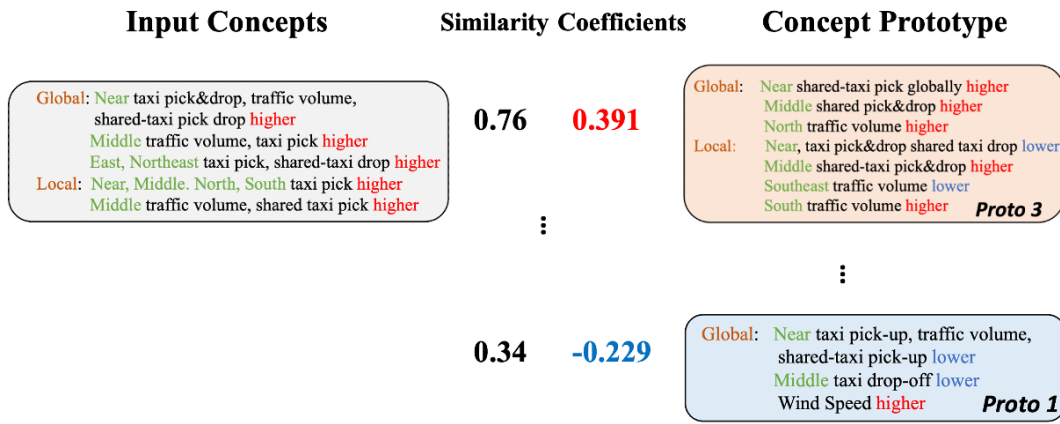


Figure 2: Grey boxes represent the encoded concepts for this given case, and red and blue boxes represent learned concept prototypes with positive and negative correlations to the occurrence of events.




Figure 3: Mapping Similar Samples to Every Prototype over Space

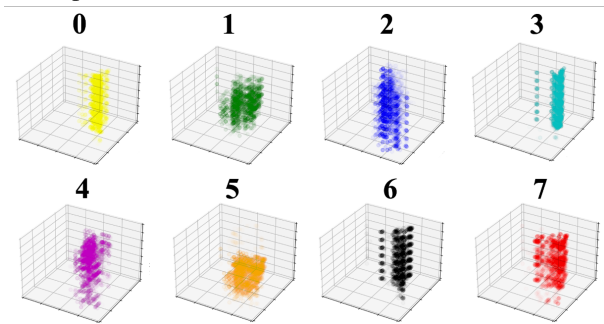


Figure 4: Mapping Similar Samples to Every Prototype over Space and Time

multiplied by higher coefficients contributes to the prediction of an event. Conversely, blue prototypes with concepts, including low traffic volume and speed indicate a safe environment, contributing less to the occurrence of an event. To understand how our learned prototypes can capture heterogeneous patterns over time and space, We map the top 1% most similar samples to each learned prototype in the study area as illustrated in Figure 3. Each prototype captures different patterns represented by hot regions over the heterogeneous space. We can visualize the heterogeneous patterns through the distribution of learning prototypes, and the positive and negative correlations between prototypes and the occurrence of events become self-explainable. To understand how prototypes are learned to be close to different samples over time, we plot the top 1% most similar samples over space and different weekdays as shown in Figure 4. Prototype P_6 covers different days of the week. Therefore, this not only proves that GeoPro-Net can capture heterogeneous patterns over time, but also that GeoPro-Net has strong interpretability by mapping diverse concept prototypes to the study area.

Conclusion

The spatiotemporal event forecasting is essential for public safety and city management. While deep neural networks have superior prediction accuracy, they struggle to intrinsically interpret the complex spatial-temporal features involved. Presenting predicted scores alone is insufficient for public understanding and future urban planning. GeoPro-Net, our proposed novel solution, addressed these challenges by introducing statistical tests on spatiotemporal event data to extract meaningful concepts. It enhanced interpretability through channel fusion and geographic-based pooling, condensing concepts for clearer insights. GeoPro-Net gained better interpretability by learning prototypes of concepts and projecting them onto real-world events scenarios. GeoPro-Net demonstrated superior interpretability and competitive performance compared to other baselines.

References

- An, B.; Zhou, X.; Zhong, Y.; and Yang, T. 2023. SpatialRank: Urban Event Ranking with NDCG Optimization on Spatiotemporal Data. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*.
- Atluri, G.; Karpatne, A.; and Kumar, V. 2018. Spatio-Temporal Data Mining: A Survey of Problems and Methods. In *Association for Computing Machinery*, volume 51. New York, NY, USA.
- Bien, J.; and Tibshirani, R. 2011. Prototype Selection for Interpretable Classification. *The annals of applied statistics*, 5(4): 2403–2424.
- Brunsdon, C.; Fotheringham, A. S.; and Charlton, M. E. 1996. Geographically Weighted Regression: A Method for Exploring Spatial Nonstationarity. *Geographical analysis*, 28(4): 281–298.
- Chen, C.; Li, O.; Tao, D.; Barnett, A.; Rudin, C.; and Su, J. K. 2019. This Looks Like That: Deep Learning for Interpretable Image Recognition. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- CPD. 2020. Chicago Police Department June 25. <https://home.chicagopolice.org/cpd-expands-smart-policing-technology-to-support-strategic-deployment-and-cta-safety/>.
- Ding, Y.; Zhu, Y.; Feng, J.; Zhang, P.; and Cheng, Z. 2020. Interpretable spatio-temporal attention LSTM model for flood forecasting. *Neurocomputing (Amsterdam)*, 403: 348–359.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *The Ninth International Conference on Learning Representations*.
- Drukker, D. M.; Prucha, I. R.; and Raciborski, R. 2013. Maximum Likelihood and Generalized Spatial Two-Stage Least-Squares Estimators for a Spatial-Autoregressive Model with Spatial-Autoregressive Disturbances. *The Stata journal*, 13(2): 221–241.
- Feldman, J. 2016. The simplicity principle in perception and cognition: The simplicity principle. *Wiley interdisciplinary reviews. Cognitive science*, 7: 330–340.
- Gov. 2017. Tennessee Highway Patrol Predictive Analytics Software. <https://www.tn.gov/safety/news/2017/10/10/tennessee-highway-patrol-provides-predictive-analytics-software-to-all-tennessee-sheriffs-offices.html>.
- Guo, S.; Lin, Y.; Feng, N.; Song, C.; and Wan, H. 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 922–929.
- Haight, F. A. 1967. Handbook of the Poisson distribution / [by] Frank A. Haight.
- Ho, J.; and Ermon, S. 2016. Generative Adversarial Imitation Learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, 4572–4580. Red Hook, NY, USA: Curran Associates Inc.
- Jiang, Z. 2015. Learning Spatial Decision Trees for Land Cover Mapping. In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, 1522–1529. IEEE.
- Jung, I.; Kulldorff, M.; and Klassen, A. C. 2007. A spatial scan statistic for ordinal data. *Statistics in medicine*, 26(7): 1594–1607.
- Koh, P. W.; Nguyen, T.; Tang, Y. S.; Mussmann, S.; Pierson, E.; Kim, B.; and Liang, P. 2020a. Concept Bottleneck Models. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 5338–5348. PMLR.
- Koh, P. W.; Nguyen, T.; Tang, Y. S.; Mussmann, S.; Pierson, E.; Kim, B.; and Liang, P. 2020b. Concept Bottleneck Models. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 5338–5348. PMLR.
- Li, Y.; Yu, R.; Shahabi, C.; and Liu, Y. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *The Sixth International Conference on Learning Representations*.
- Liu, L.; Zhen, J.; Li, G.; Zhan, G.; He, Z.; Du, B.; and Lin, L. 2020. Dynamic Spatial-Temporal Representation Learning for Traffic Flow Prediction. *IEEE Transactions on Intelligent Transportation Systems*.
- Lv, Y.; Duan, Y.; Kang, W.; Li, Z.; and Wang, F.-Y. 2019. Traffic Flow Prediction With Big Data: A Deep Learning Approach. *IEEE Transactions on Intelligent Transportation Systems*.
- Meng, L.; Zhao, B.; Chang, B.; Huang, G.; Sun, W.; Tung, F.; and Sigal, L. 2019. Interpretable Spatio-Temporal Attention for Video Action Recognition. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, 1513–1522. IEEE.
- Nauta, M.; van Bree, R.; and Seifert, C. 2021. Neural Prototype Trees for Interpretable Fine-Grained Image Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14933–14943. United States.
- Pan, M.; Huang, W.; Li, Y.; Zhou, X.; and Luo, J. 2020. xGAIL: Explainable Generative Adversarial Imitation Learning for Explainable Human Decision Analysis. In *Assoc Computing Machinery*, 1334–1343.
- Priebe, M. D. J., and Carey E. 2003. Classification using class cover catch digraphs. In *Classification using class cover catch digraphs*. Journal of classification.
- Shi, X.; Chen, Z.; Wang, H.; Yeung, D.-Y.; Wong, W.-k.; and Woo, W.-c. 2015. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical Networks for Few-shot Learning. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

- Sun, M.; Zhou, P.; Tian, H.; Liao, Y.; Xie, H.; and Pimenidis. 2022. Spatial-Temporal Attention Network for Crime Prediction with Adaptive Graph Learning. In *Artificial Neural Networks and Machine Learning - ICANN 2022*, volume 13530 of *Lecture Notes in Computer Science*, 656–669. Switzerland: Springer.
- Tang, J.; Xia, L.; and Huang, C. 2023. Explainable Spatio-Temporal Graph Neural Networks. *32nd ACM International Conference on Information and Knowledge Management (CIKM' 23)*.
- Tobler, W. R. 1970. A Computer Movie Simulating Urban Growth in the Detroit Region. *Economic geography*, 46: 234–240.
- Vahedian, A.; Zhou, X.; Tong, L.; Street, W. N.; and Li, Y. 2019. Predicting urban dispersal events: A two-stage framework through deep survival analysis on mobility data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 5199–5206.
- Wang, B.; Lin, Y.; Guo, S.; and Wan, H. 2021. GSNet: Learning Spatial-Temporal Correlations from Geographical and Semantic Aspects for Traffic Accident Risk Forecasting. In *2021 AAAI Conference on Artificial Intelligence (AAAI'21)*.
- Wiesen, C. 2019. In *Learn to use the Kolmogorov-Smirnov test in Stata with the Cardiac catheterization diagnostic dataset (2018) / Christopher Wiesen*. London: SAGE Publications, Ltd.
- Yao, H.; Wu, F.; Ke, J.; Tang, X.; Jia, Y.; Lu, S.; Gong, P.; Ye, J.; and Zhenhui, L. 2018. Deep Multi-View Spatial-Temporal Network for Taxi Demand Prediction. In *The Thirty-Second AAAI Conference on Artificial Intelligence*.
- Yao, M.; Xu, P.; Qu, H.; and Liu, R. 2019. Interpretable and Steerable Sequence Learning via Prototypes.
- Yuan, Y.; Ding, J.; Shao, C.; Jin, D.; and Li, Y. 2023. Spatio-Temporal Diffusion Point Processes. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3173–3184.
- Yuan, Z.; Zhou, X.; and Yang, T. 2018. Hetero-convlstm: A deep learning approach to traffic accident prediction on heterogeneous spatio-temporal data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 984–992.
- Zheyi Pan, W. W., Yuxuan Liang. 2019. Urban Traffic Prediction from Spatio-Temporal Data Using Deep Meta Learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19*. Association for Computing Machinery.
- Zuo, S.; Jiang, H.; Li, Z.; Zhao, T.; and Zha, H. 2021. Transformer Hawkes Process. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.