Mining Reaction and Diffusion Dynamics in Social Activities

Taichi Murayama
SANKEN, Osaka University
Japan
taichi@sanken.osaka-u.ac.jp

Yasuko Matsubara
SANKEN, Osaka University
Japan
yasuko@sanken.osaka-u.ac.jp

Yasushi Sakurai
SANKEN, Osaka University
Japan
yasushi@sanken.osaka-u.ac.jp

ABSTRACT

Large quantifies of online user activity data, such as weekly web search volumes, which co-evolve with the mutual influence of several queries and locations, serve as an important social sensor. It is an important task to accurately forecast the future activity by discovering latent interactions from such data, i.e., the ecosystems between each query and the flow of influences between each area. However, this is a difficult problem in terms of data quantity and complex patterns covering the dynamics. To tackle the problem, we propose FluxCube, which is an effective mining method that forecasts large collections of co-evolving online user activity and provides good interpretability. Our model is the expansion of a combination of two mathematical models: a reaction-diffusion system provides a framework for modeling the flow of influences between local area groups and an ecological system models the latent interactions between each query. Also, by leveraging the concept of physics-informed neural networks, FluxCube achieves high interpretability obtained from the parameters and high forecasting performance, together. Extensive experiments on real datasets showed that FluxCube outperforms comparable models in terms of the forecasting accuracy, and each component in FluxCube contributes to the enhanced performance. We then show some case studies that FluxCube can extract useful latent interactions between queries and area groups.

CCS CONCEPTS

• **Information systems** → **Data mining**; • **Mathematics of computing** → **Time series analysis**.

KEYWORDS

time series, data mininig, web data, reaction-diffusion systems, neural networks

ACM Reference Format:
Taichi Murayama, Yasuko Matsubara, and Yasushi Sakurai. 2022. Mining Reaction and Diffusion Dynamics in Social Activities. In *Proceedings of the 31st ACM Int'l Conference on Information and Knowledge Management (CIKM '22), Oct. 17–21, 2022, Atlanta, GA, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3511808.3557396

1 INTRODUCTION

The increasing volume of online user activity provides vital new opportunities to measure and understand the collective behavior of social and economic evolutions such as influenza prediction [17], the impact of individual performance [16], user activity modeling [36], and other problems [28, 44, 57]. A record of online user activity can play the role of a social sensor and offers important insights into people's decision-making. In particular, accurately forecasting the volume of online user activity and unraveling hidden patterns and interactions within them have many benefits. For example, marketers want to know the future popular volume of products and the relationship between multiple products and locations in online user attention to enable appropriate advertisement and new product development. They can avoid wasting human and material resources by accurately modeling future behavior.

The modeling of online user activity remains a challenging task due to the increasing data volume, which has multiple domains and a network of time series with mutual influences. For example, when we consider online user behavior analysis using web search activities as observations, which could be in the form of 3rd-order tensor (timestamp, location, keyword), it is important to deal with the following problems: (a) ***Capturing latent interactions and diffusion between observable data***: Many time series for individual locations and keywords are not independent. For example, the search volume per keyword in the same category may compete for user resources [35]. Also, the popularity of a keyword in one location diffuses and affects the search volume in another location [40]. In other words, we consider changes in the popularity of a keyword have the properties of flux phenomena, that can be expressed as a flow of influence. (b) ***Capturing latent temporal patterns behind observable data***: Many time series contain several patterns, such as trends and seasonality. Such patterns behind time series reveal the relationship between people's activities, such as Black Friday. However, it is extremely difficult to design an appropriate model for such patterns by hand without knowing their real characteristics in advance. The modeling method should therefore be fully automatic as regards estimating the hidden pattern for understanding data structures and saving human resources. (c) ***Accurate forecast***: Accurate time series forecasts give marketers useful insights into the future trends of their keywords and avoid wasting human and material resources. The key challenge is to achieve high forecasting accuracy while addressing the two aforementioned issues related to the interpretability of observable data.

This paper focuses on an important time series modeling and forecasting task, and we design FluxCube, which is an effective mining method that forecasts large collections of co-evolving online user activity and provides good interpretability. This model solves (a) and (b) in the above problems with the combination of mathematical models: a reaction-diffusion system [19] representing the change in space and time of chemical substances and the Lotka-Volterra population model [18, 34] describing the population dynamics of prey and predator. Also, it solves (c) with the idea

CIKM '22, October 17–21, 2022, Atlanta, GA, USA

ACM ISBN 978-1-4503-9236-5/22/10.
https://doi.org/10.1145/3511808.3557396

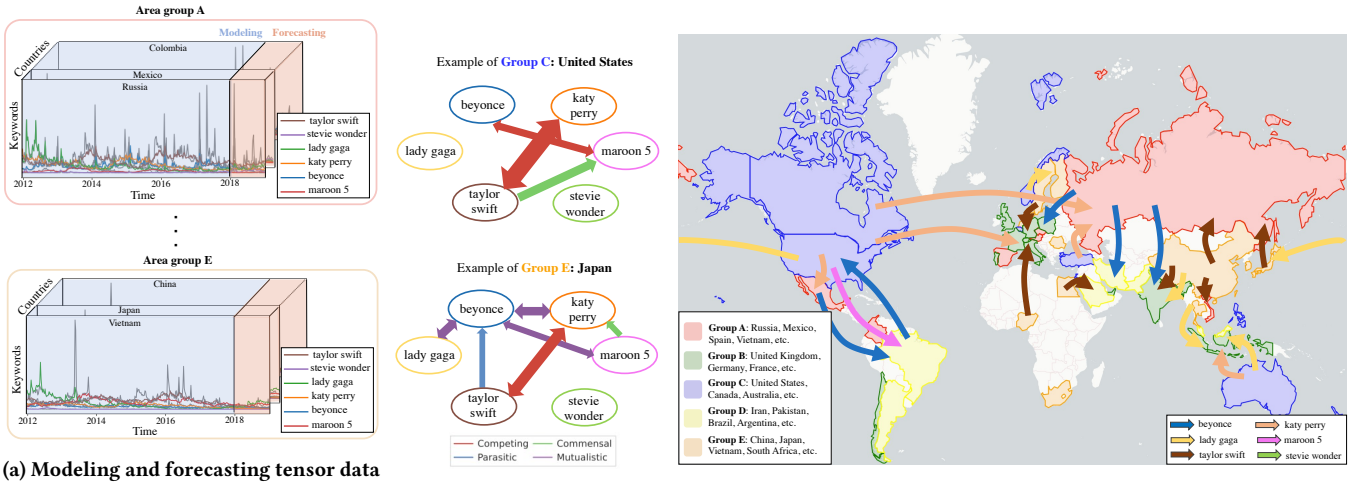

(a) Modeling and forecasting tensor data of each area group

(b) Latent interactions

(c) Diffusion process of each area group

Figure 1: Modeling power of FluxCube for online user activity data related to six artists (World#1): (a) Given the original data (gray lines), FluxCube captures the dynamics of social activities (blue). Futhermore, FluxCube provides a multiple steps ahead forecast (red). (b) Our method detects the hidden interactions between keywords in any country. The arrows indicate the direction and the type of influence. (c) It also automatically discovers similar area groups based on their keyword interactions and then finds the flow of influence of each keyword between area groups. The figure visualizes the flow of influence in 2015.

of physics-informed neural networks [46] to conduct supervised learning tasks while respecting any given laws of physics described by general nonlinear partial differential equations. Intuitively, the problem we wish to solve is as follows.

Informal Problem 1. ***Given*** *a tensor $\mathcal{X}^c$ up to the time point t_c, which consists of elements at L locations for K keywords, i.e., $\mathcal{X}^c = \{x_{tij}\}_{t,i,j=1}^{t_c,L,K}$,*

- ***find*** *trends and seasonal patterns*
- ***find*** *area groups of similar dynamics*
- ***find*** *latent interections between keywords*
- ***find*** *the flow of influence between area groups*
- ***forecast*** *l_f-step ahead future values, i.e., $\mathcal{X}^f = \{x_{tij}\}$ where $(t = t_c + l_f; i = 1, ..., L; j = 1, ..., K)$*

1.1 Preview of our results

Figures 1 and 2 show the results obtained with FluxCube for modeling online user activity data related to six artists (World#1). Specifically, our method captures the following properties:

- **Long-term trends**: Figure 1 (a) shows tensor data that consists of weekly online search volumes in relation to artists in fifty countries. FluxCube automatically captures the trends and the hidden interactions between keywords and area groups in the modeling term (blue part) and realizes the accurate long-term forecasting over two years (red part).
- **Latent interaction between keywords**: Figure 1 (b) shows the latent interaction between each keyword. FluxCube uncovers four types of interaction (competing, commensal, parasitic, mutualistic) between the keywords, where the arrow colors and widths show the interaction type and its connection intensity, from the tensor data in each location. For example, while the relationship between maroon 5 and beyonce is competing in the United States, it is mutualistic in Japan. Also, the competing relationship between taylor swift and katy perry is common in the two countries.
- **Diffusion process of each area group**: FluxCube finds the location clustering and the flow of influence between each area group, as shown in Figure 1 (c). This figure shows that the colors of the countries on the map indicate area groups, and the colors and directions of the arrows indicate the keyword type and the influence flow, respectively. FluxCube discovers the most suitable clustering result: Europe in group B, North America in group C, and East Asia in group D. The influences of each keyword on each group, which are represented by arrows, have time-varying intensities, which FluxCube also captures, as shown in Figures 2 (a) and (b). There are constant influence flows of maroon 5 and katy perry from group C to D, while there is an influence flow of taylor swift in the early 2010s from group C to E, which has disappeared. Our model provides the flow of influence in an intuitive form.
- **Seasonality**: Figure 2 (c) shows the seasonalities extracted from the search volume data. Two kinds of annual patterns resulting from "Grammy Awards" and "New Year holidays" are found from search volumes in the United States and Brazil. It is important to discover such periodic patterns if we are to accurately forecast and model the dynamics.

2 RELATED WORK

In this section, we briefly describe investigations related to our research. The details of related work are separated into three categories: time series forecasting and modeling, web data modeling, and a physics-informed neural network.

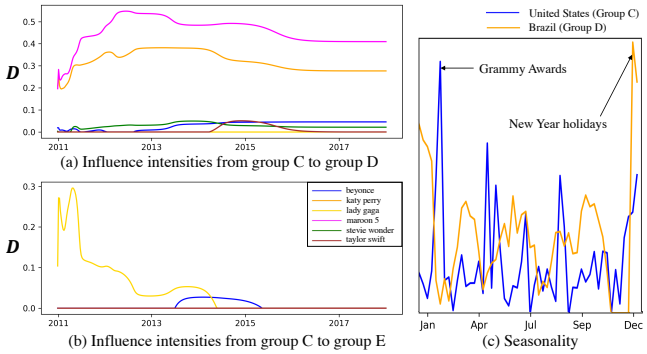

Figure 2: Latent pattern mining of FluxCube for search volumes related to six artists in the US: (a) and (b) show that it automatically detects time-varying influence intensities related to each artist from group C. (c) shows that it automatically detects seasonality patterns about keyword "beyonce" in the US and Brazil.

Time series forecasting and modeling: There is a lot of interest in time series forecasting and modeling. Conventional approaches to time series forecasting and modeling are based on statistical models such as auto regression (AR), state space models (SSM), Kalman filters (KF), and their extensions [9, 11, 30]. Some models combining these classical methods with dimension reduction, such as TRMF [55] and SMF [20] have shown useful results in the field of data mining. However, these methods are limited as regards the interpretability of observation data because of the dimension reduction. In recent years, many models have been based on neural networks thanks to their rapid development. These networs include convolutional neural networks (CNNs) [2, 4, 41] and recurrent neural networks (RNNs) [26, 47, 52], which capture the temporal variation. In particular, time series forecasting models using Transformer, developed in the field of a natural language preocessing [54], achieve high prediction performance [12, 31, 33]. For example, Informer [58], one of the Transformer-based models, achieves a successful outcome in long-ahead forecasting. Deep learning-based models guarantee high forecasting performance; however, they are still problematic in terms of their interpretability. We achieve high interpretability as well as a high predictive performance by our mathematical model fused with deep learning.

Web data modeling: Online user activity data, such as web search query data and social media posts, are expected to capture human movements as social sensors [1, 6, 49] and are utilized in real-world applications such as influenza forecasting [17], talent flow forecasting [56], and item popularity prediction [39]. This potential for real-world application has led to a huge interest in uncovering and modeling the movement of web data. The author of [3, 43] investigated the competing dynamics of two keywords on the web and found factors impacting final states based on the infection model. Various other modeling studies on the web have been conducted as follows: competing dynamics of membership-based websites [48], modeling the content and viewer dwell time on social media platforms [27], modeling the online activities into return and exploration across online communities [21, 29], extracting the relationships between search queries and external events [22, 23], and modeling online user behavior patterns in the knowledge community [37]. CubeCast [24] and EcoWeb [35] are focused on modeling the dynamics of web search queries. CubeCast is an online algorithm designed to capture useful time series segmentation and seasonality patterns. EcoWeb is a modeling method for capturing keyword interactions from the dynamics of web search queries based on biological mathematics. These studies revealed the interaction of search queries in a single location, but did not explore how keywords interact between locations. Our proposed model discovers the flow of influence between locations from the time series of each query.

Table 1: Symbols and definitions

Symbol	Definition
t_c, t_f	Modeling and forecasting time point
$\mathcal{X}^c, \mathcal{X}^f$	Modeling and forecasting tensor data
l_c, l_f	Time interval for modeling and forecasting
L, K	Number of locations and keywords
d_l	Number of area groups
a_j	Intrinsic growth rate of keyword j
b_j	Carrying capacity of keyword j
$c_{jj'}$	Intra/inter-keywords interaction strength from j'-th keyword to j-th keyword
$\boldsymbol{D}^t_{mnk}$	Influence flow of the keyword k from location n to m
$\boldsymbol{S}$	Seasonality latent pattern matrix
p	Period of seasonality
α, β	Weights of regularization terms for loss function
$\boldsymbol{o}^i$	Compression representations of item interactions in location i
$g^1, ...g^{d_l}$	Area groups
Θ'_{d_l}	Parameter sets related to d_l

Physics-informed neural network: Recently, neural networks have found applications in the context of numerical solutions of partial differential equations and dynamical systems because a neural network is a universal approximator [46]. These models, which are called physics-informed neural networks, can solve forward problems [45], where the approximate solutions of governing equations are obtained, as well as inverse problems, where parameters involved in the governing equation are obtained from the observation data [13]. We try to utilize the idea to represent some of the parameters in our partial differential equations flexibly with a neural network.

3 OVERVIEW

This section provides an overview of our proposed model, namely FluxCube, for mining and forecasting co-evolving online user activity data. We first introduce related notations and definitions and then describe the characteristics of FluxCube.

3.1 Problem definition

Table 1 lists the main symbols that we use throughout this paper. We consider online user activity data to be a 3rd-order tensor, which is denoted by $\mathcal{X} \in \mathbb{R}^{T \times L \times K}$, where T is the number of time points, and L and K are the numbers of locations and keywords, respectively. The element x_{tij} in $\mathcal{X}$ corresponds to the search volume at time t in the i-th location of the j-th keyword. Our overall aim is to realize the long-term forecasting of a tensor $\mathcal{X}$ while extracting the hidden

patterns and latent interactions. We define $\mathcal{X}^c = \{x_{tij}\}_{t,i,j=1}^{t_c,L,K}$ as observable tensor data, whose length is denoted by l_c. Similarly, let $\mathcal{X}^f = \{x_{tij}\}_{t,i,j=t_c,1,1}^{t_f,L,K}$ denotes a partial tensor from t_c to t_f for forecasting. Here, we define l_c and l_f as a time interval for modeling and forecasting. Finally, we formally define our problem as follows.

Problem 1 (l_f-step ahead forecasting). ***Given*** *a tensor* $\mathcal{X}^c = \{x_{tij}\}_{t,i,j=1}^{t_c,L,K}$, *up to the time point* t_c; ***Forecast****:* l_f*-step ahead future values* $\hat{\mathcal{X}}^f = \{\hat{x}_{tij}\}_{t,i,j=t_c,1,1}^{t_f,L,K}$, *where* $t_f = t_c + l_f$.

3.2 Reaction-diffusion system

Our model for capturing the latent interactions between keywords and area groups is inspired by a reaction-diffusion system. A reaction-diffusion system is a mathematical model corresponding to physical phenomena such as the change in space and time of chemical substances, represented as partial differential equations. The system can also describe the dynamic processes of non-chemical areas such as biology, geology, and physics [8, 14, 15, 42]. The general form of a reaction-diffusion system can be described with the following equations:

$$\frac{\partial \boldsymbol{u}}{\partial t} = f(\boldsymbol{u}, t) + \mathcal{D}\Delta \boldsymbol{u} \tag{1}$$

where $\boldsymbol{u}$ represents the concentration values in a chemical substance, t represents the present time, and $\mathcal{D}$ is the diffusion coefficient. The first term on the right hand side, $f(\boldsymbol{u}, t)$, represents the reaction term, which accounts for all local reactions, and the second term, $\mathcal{D}\Delta\boldsymbol{u}$, represents the diffusion term, which accounts for how the substance spreads to other locations.

Our model is an extended reaction-diffusion system for modeling and forecasting online user activity data. For example, the related search volumes in the data interact with each other and compete for user resources. The interaction between keywords in a location can be represented as a reaction term in the reaction-diffusion system. Also, a vogue for any keyword in a location can affect that in another location. Such an influence flow of any keyword between locations can be represented as a diffusion term. Thus, we assume that the interactions between area groups and keywords can be represented by a diffusion-reaction system.

3.3 Capabilities

To capture all the components outlined in the introduction, we present our model, namely FluxCube which can model the latent patterns and interactions underlying online user activity 3rd-order tensor data. So, how can we build our model so that it models the user activity data while capturing the latent patterns and interactions and achieving accurate forecasting? Specifically, our model should have the following three capabilities:

- **Latent interaction and diffusion**: User activity data evolves naturally over time depending on many latent factors, such as user preferences and customs for online web-search activities. We need to capture the latent interactions of keywords and the influence flow of locations from the real-time series data. To explicitly capture such interactions, we propose using a reaction-diffusion system [19] including the Lotka-Volterra population model [34], which is expressed in differential equations.

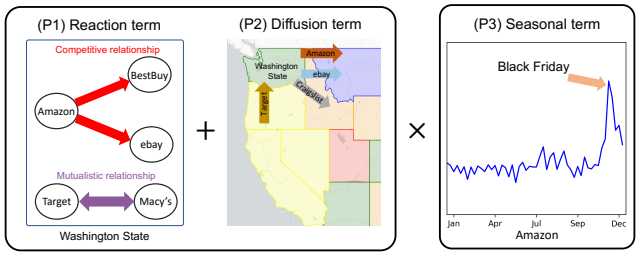

Figure 3: Conceptual image of our model targeting ecommerce data (US#1) in Washington state. Our model consists of three parts: (P1) Reaction term represents the interaction between keywords, (P2) Diffusion term represents the influence flow of each keyword between locations, and (P3) Seasonal term represents the seasonality of each keyword.

- **Seasonality**: We should also note that online user activity has certain annual patterns. As Figure 2 (c) shows, there are annual patterns such as a huge spike in New Year holidays in online user activity data. Capturing seasonalities realizes our suitable modeling.
- **Time-varying components**: User activity data change significantly due to external events such as the influence of other locations. The variables in the differential equation cannot easily explain the relationships between locations that vary greatly with time and external events, for which we require a flexible modeling method. To achieve this, we leverage the idea of physics-informed neural networks [46]. Intuitively, our model combines a neural network with the characteristics of universal approximation [32] and differential equations with high interpretability.

4 FLUXCUBE

Here, we describe the formula of FluxCube. We represent the base form of our model applied to online user activity data, mainly web search volumes, which is inspired by the reaction-diffusion system, Eq. (1), in the following:

$$\frac{\partial x_{tij}}{\partial t} = f(x_{tij}|\mathcal{X}^c_{t,i,:}) + g(x_{tij}|\mathcal{X}^c_t, t) \tag{2}$$

$$\hat{x}_{t+1ij} = F(\frac{\partial x_{tij}}{\partial t} + x_{tij}) \tag{3}$$

The first term in Eq. (2) serves as the reaction term for extracting the interaction between keywords in location i (**P1**). The second term in Eq. (2) serves as the diffusion term for extracting the influence flow of keyword j between locations (**P2**). Finally, we set Eq. (3) to capture seasonality (**P3**). A conceptual drawing of our modeling method is shown in Figure 3.

Next, we introduce each component of the above form of our model in steps.

4.1 Reaction term (P1)

We first introduce the reaction term in detail. The term aims to model the interaction between keywords within a location.

Observation of the search volume of each keyword on the web shows that the keywords compete for user attention. No keyword

can survive on the web if no one pays attention to that topic. Besides, the number of instances of user attention on the web is finite. This relationship between keywords and user attention on the web is very similar to the relationship between species and food resources in the jungle, e.g., squirrel monkeys and spider monkeys compete for fruit and no species can survive without resources. We employ the Lotka-Volterra population model [18, 34], which represents the growth of the species considering the interaction between them in biological mathematics, as the reaction term.

The reaction term in Eq. (2) for keyword j is represented below:

$$f(x_{tij}|\mathcal{X}^c_{t,i,:}) = a_j x_{tij}\left(1 - \frac{\sum_{j'} c_{jj'} x_{tij'}}{b_j}\right), \tag{4}$$

where, $a_j > 0$, $b_j > 0$ and $c_{jj} = 1$. The size of each parameter depends on the number of keywords K: $\boldsymbol{a} \in \mathbb{R}^K$, $\boldsymbol{b} \in \mathbb{R}^K$, and $\boldsymbol{C} \in \mathbb{R}^{K\times K}$. Each parameter is interpreted as:

- a_j: intrinsic growth rate of keyword j;
- b_j: carrying capacity of keyword j;
- $c_{jj'}$: intra/inter-keyword interaction strength from the j'-th keyword to the j-th keyword, which is each value in $\boldsymbol{C}$

The signs of the off-diagonal elements of $\boldsymbol{C}$ provide an interpretation of the latent relationship between the two keywords:

- $c_{jj'} > 0, c_{j'j} > 0$: a competitive relationship
- $c_{jj'} < 0, c_{j'j} > 0$: a parasitic relationship
- $c_{jj'} < 0, c_{j'j} = 0$: a commensal relationship
- $c_{jj'} < 0, c_{j'j} < 0$: a mutualistic relationship

We expect these characteristics to be available for modeling online user activity on the web and extracting the relationships between keywords.

4.2 Diffusion term (P2)

We then introduce the diffusion term in detail. The term aims to model the influence flow of any keyword between locations.

Can partial differential equations and mathematical models adequately represent the relationships between locations in online user activity on the web? This task is hard to achieve. The reaction-diffusion system in chemistry can represent the change in space and time of chemical substances by a constant such as the diffusion coefficient, based on the observation that the spread of a chemical substance is constant without external influences. However, on the web, interactions of any keyword between locations are not constant due to external factors. To capture the time change of the interaction, a dynamic linear model with time-varying coefficients could be utilized, but it is difficult to take account of complex phenomena and rapid changes. Here, we utilize a neural network with universal approximation as part of our mathematical model to represent the changing interactions between locations over time. This corresponds to a kind of physics-informed neural network [46]. Concretely, we represent the diffusion term in Eq. (2) for keyword j at location i as below:

$$\boldsymbol{D}^t = ReLU\left(RNN(1:t)\right) \tag{5}$$

$$\boldsymbol{E}^t_i = \boldsymbol{D}^t_i \odot \mathcal{X}_t, \qquad (i = 1, ..., d_l) \tag{6}$$

$$g(x_{tij}|\mathcal{X}^c_t, t) = \sum_{i'}^{d_l} \boldsymbol{E}^t_{i,i',j} \tag{7}$$

where $\boldsymbol{D}^t \in \mathbb{R}^{d_l\times d_l\times K}$ and $\boldsymbol{E}^t \in \mathbb{R}^{d_l\times d_l\times K}$, and $\boldsymbol{D}^t_{i,i,:} = 0$. $\boldsymbol{D}^t_i \in \mathbb{R}^{d_l\times K}$ and $\boldsymbol{E}^t_i \in \mathbb{R}^{d_l\times K}$ represent the i-th vectors of $\boldsymbol{D}^t$ and $\boldsymbol{E}^t$, respectively. We express the flow of attention of each user between locations using a recurrent neural network (RNN) [51], which captures time-varying parameters using time $1:t$ as input in Eq. (5). The RNN enables our model to be more expressive by allowing time-varying parameters along with temporal dependencies in the covariates. The output of RNN in Eq. (5) is transformed to a tensor form: $\boldsymbol{D}^t$ is a form of 3rd-order tensor and $\boldsymbol{D}^t_{mnk}$ indicates the degree to which location n contributes to the popularity of the keyword k in location m. In other words, $\boldsymbol{D}^t_{mnk}$ represents the influence intensities of the keyword k from n to m, which is provided in an intuitive form for us. We show some examples of $D^t(t = 1, ..., t_c)$ in Figure 3 (a) and (b). A total of the influences that keyword j at location i at time t receives from other locations is expressed by multiplying the observed data $\mathcal{X}$ with the output in Eq. (5) ($\boldsymbol{D}$), as shown in Eq. (6) and (7). By applying a neural network to a part of our mathematical model, we expect to achieve both flexible modeling and high explainability.

4.3 Seasonality (P3)

We finally consider seasonal/cyclic patterns in user activity data by extending our model. Each keyword always has certain users' attention; however, the users change their behavior dynamically, according to various seasonal events, e.g., Christmas Day and Black Friday. We need to detect hidden seasonal activities with our model. More specifically, we model the user activity value at the next step considering the seasonality pattern from our reaction-diffusion system in Eq. (3), as below:

$$F\left(\frac{\partial x_{tij}}{\partial t} + x_{tij}\right) = \left(\mathbf{1} + [\boldsymbol{S}_{t \bmod p}]_{i,j}\right) \odot \left(\frac{\partial x_{tij}}{\partial t} + x_{tij}\right) \tag{8}$$

where $\boldsymbol{S} \in \mathbb{R}^{p\times L\times K}$, $\boldsymbol{S} \geq 0$, p is period of seasonality, and $\odot$ indicates the element-wise product. $\boldsymbol{S}$ acts as a projection matrix for latent seasonal patterns, such as the seasonal term in Figure 3.

4.4 Loss function

To train FluxCube, we use the mean squared error (MSE) loss with regularization terms as below:

$$\|\mathcal{X}^c - \hat{\mathcal{X}}\|^2 + \alpha \sum \|\boldsymbol{D}\|^2 + \beta \sum \|\boldsymbol{S}\|^2 \tag{9}$$

where α and β are hyper-parameters for controling the weights of the regularization terms, and $\hat{\mathcal{X}}$ is our modeling value. The first term is MSE loss that penalizes the difference between modeling and observation values, and the second and third terms are regularization terms that encourage the suppression of the adoption of large values for two parameters: the influence between locations ($\boldsymbol{D}$) and the latent seasonal patterns ($\boldsymbol{S}$).

5 OPTIMIZATION ALGORITHM

This section presents our optimization algorithm for learning FluxCube while grouping locations with similar user activity characteristics.

Thus far, we have proposed FluxCube for modeling and forecasting online user activity data with multiple locations. There are two challenges in fitting our model to data with many locations: (a)

the learning can be inefficient and unstable because there are large candidate parameters that must be estimated and converged, and (b) the interactions between locations become more complex and results less interpretable. Here, we propose an algorithm to learn FluxCube where we cluster locations that have observed data with similar characteristics as the same area group.

5.1 Automatic location clustering

Our hypothesis for clustering online user activity data as regards location is simple. It is that locations with similar interactions between items represented by the reaction term (P1) can be considered as the same area group. The clustering method refers to [10], according to the number of area group d_l we cluster the locations, the number of which is L, using Uniform Manifold Approximation and Projection (UMAP) [38] and K-means: The equations are as below:

$$\boldsymbol{o}^i = UMAP(\boldsymbol{a}^i, \boldsymbol{b}^i, \boldsymbol{C}^i) \qquad (i = 1, ..., L) \tag{10}$$

$$g^1, ..., g^{d_l} = \text{K-means}(\boldsymbol{o}, d_l) \tag{11}$$

where $\boldsymbol{o}^i \in \mathbb{R}^2$, $\boldsymbol{o} \in \mathbb{R}^{L\times 2}$ is the compression representation of the keyword interactions and L is the number of locations. We perform clustering using three parameters $(\boldsymbol{a}^i, \boldsymbol{b}^i, \boldsymbol{C}^i)$ that represent the interaction of keywords at location i in Eq. (4) in the reaction term. These three parameters are compressed to the two-dimensional vector $\boldsymbol{o}^i$ by UMAP. We then cluster the compression representation $\boldsymbol{o}$ by K-means into d_l area groups $(g^1, ..., g^{d_l})$.

5.2 Finding the appropriate number of area groups

We describe how to find the appropriate number of area groups d_l. The situation where our model with a small number of parameters including d_l can adequately explain the observed data is the best. Here, we apply the minimum description length (MDL) principle [5, 50, 53] to find the appropriate d_l. The MDL principle enables us to determine the nature of good summarization by minimizing the sum of the data encoding cost and the model description cost. The cost is described as below:

$$< \mathcal{X}^c | \Theta' > + < \Theta'_{d_l} > + C \tag{12}$$

where $< \mathcal{X}^c | \Theta' >$ represnts the cost of describing the data $\mathcal{X}^c$ given the model parameters Θ', $< \Theta'_{d_l} >$ shows the cost of describing Θ'_{d_l}, which is a parameter set that varies with d_l, and C is a constant value, which is not affected by d_l. In short, it follows the assumption that the more we compress data, the more we can learn about the underlying pattern. We begin by defining the two components of the total cost more concretely.

Data encoding cost: We can encode the observation data $\mathcal{X}^c$ using Θ' based on Huffman coding [50]. The coding scheme assigns a number of bits to each value in $\mathcal{X}^c$, which is the negative log-likelihood under a Gaussian distribution with mean μ and variance σ^2, i.e.,

$$< \mathcal{X}^c | \Theta' > = \sum_{t,i,j=1}^{t_c,L,K} -\log_2 p_{\mu,\sigma}(x^c_{t,i,j} - \hat{x}^c_{t,i,j}) \tag{13}$$

where, $\hat{x}^c_{t,i,j} \in \hat{\mathcal{X}}^c$ is the modeling value of $x^c_{t,i,j} \in \mathcal{X}^c$ by our proposed model.

Algorithm 1 Optimization

Input: Input tensor $\mathcal{X}^c$
Output: Learned FluxCube with interpretable parameters

```
/* Initialize
d_l ← 1
g^{d_l} ← {1, ..., L}
MinCost ← inf
while improving the cost do
    if d_l > 1 then
        /* Automatic location clustering (Section 5.1)
        g^1, ..., g^{d_l} ← Clustering(d_l, a_{d_l−1}, b_{d_l−1}, C_{d_l−1})
    end if
    /* FluxCube training (Section 4)
    FluxCube_{d_l} ← Training(X^c, g^1, ..., g^{d_l})
    /* Calculate the encoding cost (Section 5.2)
    Cost_{d_l} ← f(FluxCube_{d_l}, X^c)    //f(·) : Equation(12)
    /* Update d_l
    if Cost_{d_l} < MinCost then
        MinCost ← Cost_{d_l}
        d_l ← d_l + 1
    else
        return FluxCube_{d_l}
    end if
end while
```

Model description cost: The model cost is the number of bits needed to describe the model. If we use a more powerful model architecture, the total cost becomes higher. We focus only on parameters related to d_l, i.e., $\boldsymbol{D}$ in the diffusion term, because of our objective of selecting the appropriate d_l. Parameters unrelated to d_l can be treated as constant value C, which we can ignore in the cost calculation. Our model cost is represented as below:

$$< \Theta'_{d_l} > = < d_l > + < D > \tag{14}$$

$$< d_l > = \log^*(d_l) \tag{15}$$

$$< D > = |\boldsymbol{D}| \cdot (\log(t_c) + 2 \cdot \log(d_l) + \log(K) + c_F) + \log^*(|\boldsymbol{D}|) \tag{16}$$

Here, $\log^*$ is the universal code length for integers, $|\cdot|$ describes the number of non-zero elements and c_F denotes the floating point cost (32 bits).

5.3 Optimization

Summarizing the descriptions so far, we have shown the overall procedure of learning FluxCube while grouping locations with similar user activity characteristics in Algorithm 1. The basic idea of the algorithm is that we search FluxCube to minimize the encoding cost in Eq. (12) while increasing the number of area groups d_l.

6 EXPERIMENTS AND RESULTS

6.1 Experimental settings

6.1.1 Datasets. We used the two kinds of online user activity data from different target areas on GoogleTrends, which contained

Table 2: Datsets

ID	Dataset	Query
US#1	E-commerce	Amazon/Apple/BestBuy/Costco/Craigslist/Ebay/ Homedepot/Kohls/Macys/Target/Walmart
US#2	VoD	AppleTV/ESPN/HBO/Hulu/Netflix/Sling/ Vudu/YouTube
US#3	Sweets	Cake/Candy/Chocolate/Cookie/Cupcake/ Gum/Icecream/Pie/Pudding
US#4	Facilities	Aquarium/Bookstore/Gym/Library/Museum/ Theater/Zoo
World#1	Music	Beyonce/KatyPerry/LadyGaga/Maroon5/ StevieWonder/TaylorSwift
World#2	SNS	Facebook/LINE/Slack/Snapchat/Twitter/ Viber/WhatsApp
World#3	Apparel	Gap/H&M/Primark/Uniqlo/Zara

weekly web search volumes collected for about 10 years, from January 1, 2011, to December 31, 2020.

- **US data** includes web search volumes for 50 states of the US. Four datasets of this type were constructed:
- **World data** includes three kinds of web search volumes for the top 50 countries ranked by GDP score.

The queries of our datasets are described in Table 2.

6.1.2 Comparative models. We used the following comparative models, which are state-of-the-art algorithms for modeling and forecasting time series:

- **EcoWeb** [35], which is intended for modeling online user activity data as well as our model, is a mathematical model constructed on the basis of differential equations.
- **SMF** [20] is a matrix factorization model that takes into account seasonal patterns. It is mainly used for time series forecasting and anomaly detection.
- **DeepAR** [52] has an encoder-decoder structure that employs an auto-regression RNN modeling probabilistic distribution in the future. Here, we selected a two-stack LSTM as the structure and its suitable number of hidden units in the validation period.
- **Gated Recurrent Unit (GRU)** [7] is a RNN-based model for time series forecasting. We employed an encoder-decoder architecture based on the GRU for the multi-step-ahead forecast. We also applied a dropout rate of 0.5 to the connection of the output layer.
- **Informer** [58], which is a transformer-based model based on ProbSparse self-attention and self-attention distilling, is known for its remarkable efficiency in long time series forecasting. We select a two-layer stack for both encoder and decoder and set 26 (half of the prediction sequence length) as the token length of the decoder.

6.1.3 Setting. We forecast the online user activity data using FluxCube and other comparative models for the validation of the prediction performance.

Dataset preprocessing: We set search volumes in the datasets from January 1, 2011, to December 31, 2017, as the training and modeling period (seven years = t_c), then from January 1, 2018, to December 31, 2018, as the validation period (one year), and then from January 1, 2019, to December 31, 2020, as the test period (two years). We then normalized their search volumes in the range 0 to 1.

Evaluation metrics: We evaluate the predictive performance of these models at three points, 13 weeks (3 months), 26 weeks (half a year), and 52 weeks (a year) ahead. Two evaluation metrics were utilized to compare the forecast performance levels of each model: root mean squared error (RMSE) and mean absolute error (MAE). A lower value in these metrics indicates better forecasting accuracy.

Hyper-parameter: All the parameters in neural network based models were updated in the Adam update rule [25]. We set 52 as the input and forecast length and MSE loss as the loss function for DeepAR, GRU, and Informer. We then conducted a grid search for parts of parameters such as hidden layers in comparative models. In the training of FluxCube, we selected the sizes of the RNN hidden layers in Eq. (5) as (16, 32, 64) in the validation period. We also set p in Eq. (8) as 52 to capture annual patterns from weekly datasets, the hyperparameters α and β in Eq. (9) as 0.1, and the number of epochs as 2,000 with early stopping.

6.2 Results

The experimental results are presented in Table 3. These results indicate that FluxCube outperformed most of the comparative models, confirming the benefits of our model architecture and algorithm. EcoWeb, which is a linear model that captures only interactions between items to be built with the same motivation as FluxCube, was unstable for long-ahead forecasting. SMF, which focused on the extraction of seasonal patterns, outputs a good result of forecasting online user activity data with strong seasonality, compared to EcoWeb. RNN-based models such as DeepAR and GRU were superior baseline models and exhibited almost the same performance as Informer in the 13 week forecast, but in some cases, they showed unstable results in long-ahead forecasting such as 52 weeks ahead. These results indicate that it is not easy to improve the accuracy of long-ahead forecasts. On the other hand, Informer, which is built for long-term forecasting, achieved the highest accuracy among the comparative models. Informer showed a trend whereby the accuracy with the 52 weeks ahead forecast is not much worse than that with the 13 weeks ahead forecast.

FluxCube achieved the highest accuracy in 5 out of 7 datasets for RMSE and 6 out of 7 for MAE compared to other models. Our model shows high predictive performance as well as high interpretability by explicitly modeling the online user activity data with capturing the three hidden components: keyword interaction, location diffusion, and seasonality.

Ablation study: To demonstrate the effectiveness of all the components of FluxCube, we performed ablation experiments with three ablation models, namely, Constant model, and Full model w/o P2, and w/o P3. **Constant model** replaces the RNN in Eq. (5) in FluxCube with time-invariant parameters. This reveals the benefit of time-varying components. **Full model w/o P2** and **Full model w/o P3** are the removals of the diffusion term and the seasonal term in FluxCube respectively. Figure 4 shows the RMSE performances of 52 weeks ahead forecasting in each dataset. Full model

Table 3: Forecasting performance comparison

		13 weeks		26 weeks		52 weeks	
Dataset	Model	RMSE	MAE	RMSE	MAE	RMSE	MAE
US#1	EcoWeb	0.1470	0.0950	0.1554	0.1082	0.1654	0.1197
	SMF	0.0869	0.0620	0.0910	0.0654	0.1012	0.0674
	DeepAR	0.1003	0.0634	0.1302	0.0907	0.1385	0.1014
	GRU	0.1723	0.1175	0.1924	0.1374	0.2059	0.1525
	Informer	0.1477	0.1045	0.1375	0.0985	0.1575	0.1111
	FluxCube	**0.0478**	**0.0257**	**0.0574**	**0.0323**	**0.0631**	**0.0365**
US#2	EcoWeb	0.1440	0.1133	0.1981	0.1621	0.1920	0.1684
	SMF	0.0621	0.0445	0.0713	0.0522	0.0760	0.0529
	DeepAR	0.1471	0.1026	0.1781	0.1314	0.1906	0.1474
	GRU	0.1518	0.1171	0.1619	0.1231	0.1683	0.1440
	Informer	0.1277	0.0878	0.1292	0.0876	0.1436	0.1012
	FluxCube	**0.0245**	**0.0130**	**0.0276**	**0.0156**	**0.0310**	**0.0181**
US#3	EcoWeb	0.1555	0.1208	0.1730	0.1384	0.1754	0.1369
	SMF	0.0276	0.0186	0.0281	0.0170	0.0281	0.0190
	DeepAR	0.1107	0.0753	0.1267	0.0833	0.1309	0.0908
	GRU	0.1300	0.0869	0.1368	0.9843	0.1368	0.0939
	Informer	0.1322	0.0954	0.1311	0.0946	0.1279	0.0914
	FluxCube	**0.0200**	**0.0121**	**0.0222**	**0.0136**	**0.0238**	**0.0148**
US#4	EcoWeb	0.0847	0.0573	0.1348	0.0950	0.1511	0.1182
	SMF	0.0905	0.0762	0.1100	0.0751	0.1206	0.1077
	DeepAR	0.0927	0.0682	0.1662	0.1119	0.2168	0.1639
	GRU	0.1199	0.0872	0.1737	0.1223	0.2319	0.1764
	Informer	0.1014	0.0720	0.0992	0.0690	0.1055	0.0762
	FluxCube	**0.0495**	**0.0256**	**0.0610**	**0.0358**	**0.0794**	**0.0503**
World#1	EcoWeb	0.1259	0.0831	0.1422	0.1034	0.2101	0.1460
	SMF	0.0936	0.0783	0.0901	0.0602	0.1087	0.0787
	DeepAR	0.0900	0.0636	0.0929	0.0681	0.1395	0.0972
	GRU	0.0633	0.0452	0.0718	0.0501	0.0823	0.0572
	Informer	0.0704	0.0423	0.0719	0.0416	0.0738	0.0446
	FluxCube	**0.0454**	**0.0274**	**0.0477**	**0.0286**	**0.0546**	**0.0331**
World#2	EcoWeb	0.0908	0.0300	0.1089	0.0570	0.1353	0.0742
	SMF	0.0841	0.0436	0.0799	0.0454	0.0826	0.0480
	DeepAR	0.0374	**0.0098**	**0.0585**	0.0199	0.0643	0.0209
	GRU	0.0401	0.0159	0.0588	**0.0174**	0.0739	0.0254
	Informer	**0.0371**	0.0159	0.0595	0.0196	**0.0642**	**0.0208**
	FluxCube	0.0704	0.0271	0.0711	0.0304	0.0831	0.0351
World#3	EcoWeb	0.0523	0.0208	0.0626	0.0200	0.1080	0.0293
	SMF	0.0206	0.0111	0.0289	0.0160	0.0254	0.0195
	DeepAR	0.0211	0.0110	0.0275	0.0099	0.0613	0.0214
	GRU	0.0191	0.0090	0.0217	**0.0096**	0.00235	0.0115
	Informer	0.0223	0.0105	0.0226	0.0105	**0.0214**	0.0108
	FluxCube	**0.0176**	**0.0085**	**0.0214**	**0.0096**	0.0221	**0.0100**

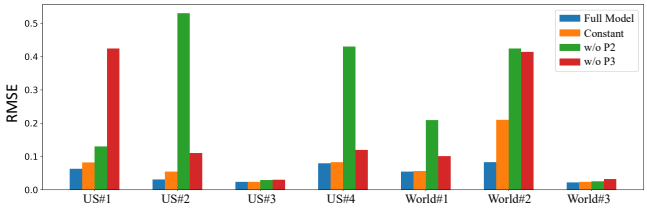

Figure 4: Experimental ablation results for RMSE performance of 52 weeks ahead forecasting.

outperformed the other ablation models in all datasets. Full model w/o P2 and w/o P3, which exclude the diffusion term and seasonal term, respectively, exhibited significantly poorer forecasting performance. Also, Constant showed a reduced forecasting performance, suggesting the effectiveness of the time-varying parameters generated by RNN. Each component in FluxCube is useful for modeling online user activity data.

7 CASE STUDY

Here, we describe the analytical results obtained with FluxCube. The result of World#1 has already been presented in Section 1 (i.e., Figure 1 and 2.) Figure 5 (a) and (b) show the FluxCube modeling and forecasting results for US#1 and World#2 tensor data, respectively. The figure contains three results obtained by FluxCube modeling each set of data: the modeling and forecasting results of the tensor data, the latent interactions between keywords for each country obtained from $\boldsymbol{C}$, and the location clustering results and the flow of influences between each area group obtained from $\boldsymbol{D}$.

Case study of US#1: Figure 5 (a-1) shows that our proposed model successfully captures and forecasts latent dynamic patterns for multiple countries and keywords. Concretely, our proposed model can capture the seasonality patterns in amazon and apple, as well as the decreasing trends of craigslist. Figure 5 (a-2) based on the reaction term in FluxCube indicates the network graphs summarizing latent interactions between keywords with different colors. For example, in New York State, it shows that amazon competes with bestbuy and ebay, and target and macys are a mutualistic relationship. Figure 5 (a-3) shows the results of clustering each state into multiple groups. The influence of macys from group A, into which FluxCube groups the northeast region of the US, affects other areas.

Case study of World#2: Figure 5 (b-1) shows that FluxCube captures the increasing trend of slack and snapchat in Spain and the decreasing trend of facebook in Japan. On the other hand, it failed to adequately capture the trend of slack, which increased rapidly in popularity in Japan after 2018. As shown in Figure 5 (b-2), FluxCube discovers many commensal and mutualistic relationships between social media platforms: e.g., in Spain, facebook has mutualistic relationships with several platforms such as line, viber, and twitter. We find that many social media tend to be symbiotic rather than compete for online users. Figure 5 (b-3) shows that our model groups countries such as Japan and Saudi Arabia, where twitter is mainly used, into group B and countries such as the United States and the United Kingdom, where facebook is mainly used, into group F. Our results about the influence interactions between area groups show that the influence of facebook flows from group F to other area groups. Also, the influence of viber founded in Israel flows from group E, including countries in the Middle East, to other area groups.

As our case studies show, FluxCube provides the ability to describe complex interactions that can reveal underlying relationships among keywords and locations, and is suitable for modeling and forecasting online user activity data.

8 CONCLUSION

In this study, we attempted to model and forecast large time-evolving online user activity data such as web search volumes. To achieve this, we proposed an effective modeling and forecasting method,

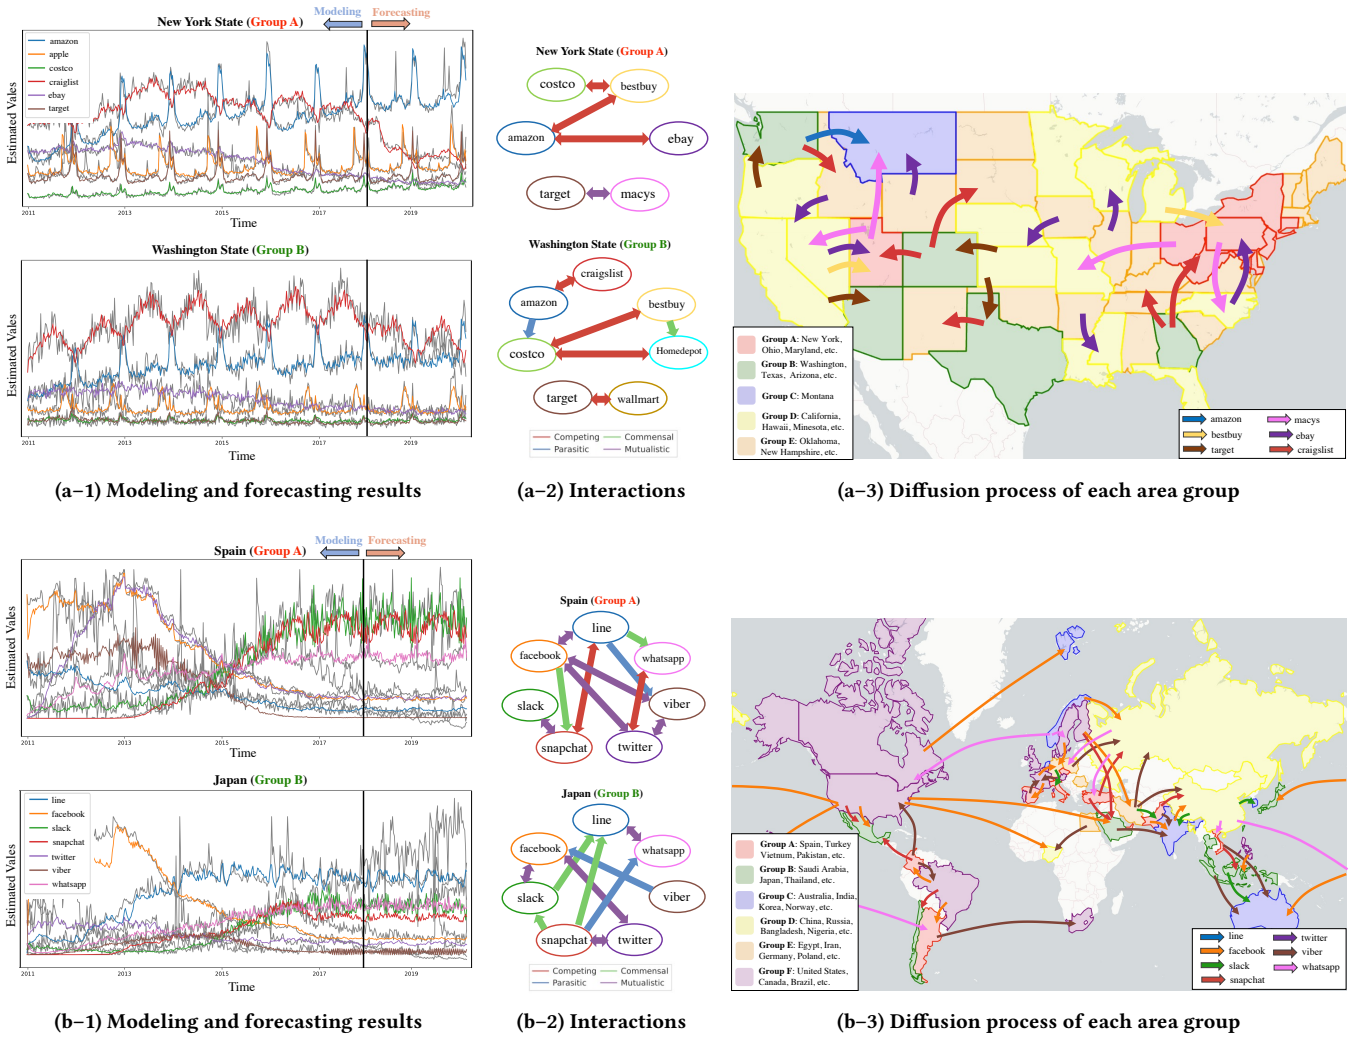

(a–1) Modeling and forecasting results (a–2) Interactions (a–3) Diffusion process of each area group

(b–1) Modeling and forecasting results (b–2) Interactions (b–3) Diffusion process of each area group

Figure 5: FluxCube modeling and forecasting results for US#1 (a) and World#2 (b) data: (a/b-1) shows the modeling and forecasting results of FluxCube given the observational data. It represents observational data in gray and the modeling and forecasting results for each keyword in various colors. (a/b-2) shows the network graphs summarizing four latent interactions between keywords in each country. FluxCube uncovers four types of interaction, each of which is represented by arrows. Competing and mutualistic relationships are bidirectional, while commensal and parasitic relationships are one-directional: source species benefit from destination species' growth. (a/b-3) shows the clustering results with countries/states divided into multiple area groups and the flow of influences between each area group.

namely FluxCube based on reaction-diffusion and ecological systems. Our method can recognize trends, seasonality and interactions in input observations by extracting their latent dynamic systems. Throughout the search volume forecasting experiments in multiple GoogleTrends datasets, the proposed model achieved the highest performance by capturing the latent dynamics.

In addition to its predictive performance, FluxCube provides the latent interactions and the influence flows hidden behind observational data in a human-interpretable form. For example, in Washington State, costco competes with bestbuy and homedepot for online users, as shown by our case studies. Uncovering such latent interactions behind the observational data assists to our decision-making.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees for their valuable comments and helpful suggestions. This work was supported by JSPS KAKENHI Grant-in-Aid for Scientific Research Number JP20H00585, JP21H03446, NICT 21481014, MIC/SCOPE 192107004, JST-AIP JPMJCR21U4, and ERCA-Environment Research and Technology Development Fund JPMEERF20201R02.

REFERENCES

[1] Eiji Aramaki, Sachiko Maskawa, and Mizuki Morita. 2011. Twitter catches the flu: detecting influenza epidemics using Twitter. In *Proceedings of the 2011 Conference on empirical methods in natural language processing*. 1568–1576.

[2] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271* (2018).

[3] Alex Beutel, B Aditya Prakash, Roni Rosenfeld, and Christos Faloutsos. 2012. Interacting viruses in networks: can both survive?. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. 426–434.

[4] Anastasia Borovykh, Sander Bohte, and Cornelis W Oosterlee. 2017. Conditional time series forecasting with convolutional neural networks. *arXiv preprint arXiv:1703.04691* (2017).

[5] Deepayan Chakrabarti, Spiros Papadimitriou, Dharmendra S Modha, and Christos Faloutsos. 2004. Fully automatic cross-associations. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. 79–88.

[6] Serina Chang, Mandy L Wilson, Bryan Lewis, Zakaria Mehrab, Komal K Dudakiya, Emma Pierson, Pang Wei Koh, Jaline Gerardin, Beth Redbird, David Grusky, et al. 2021. Supporting covid-19 policy response with large-scale mobility-based modeling. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2632–2642.

[7] Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. 2018. Recurrent neural networks for multivariate time series with missing values. *Scientific reports* 8, 1 (2018), 1–12.

[8] C Cosner. 2008. Reaction–diffusion equations and ecological modeling. In *Tutorials in mathematical biosciences IV*. Springer, 77–115.

[9] Joel Janek Dabrowski, Ashfaqur Rahman, Andrew George, Stuart Arnold, and John McCulloch. 2018. State space models for forecasting water quality variables: an application in aquaculture prawn farming. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 177–185.

[10] Kareem Darwish, Peter Stefanov, Michaël Aupetit, and Preslav Nakov. 2020. Unsupervised user stance detection on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 141–152.

[11] James Durbin and Siem Jan Koopman. 2012. *Time series analysis by state space methods*. Vol. 38. OUP Oxford.

[12] Chenyou Fan, Yuze Zhang, Yi Pan, Xiaoyue Li, Chi Zhang, Rong Yuan, Di Wu, Wensheng Wang, Jian Pei, and Heng Huang. 2019. Multi-horizon time series forecasting with temporal attention learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2527–2535.

[13] Zhiwei Fang and Justin Zhan. 2019. A physics-informed neural network framework for PDEs on 3D surfaces: time independent problems. *IEEE Access* 8 (2019), 26328–26335.

[14] Richard FitzHugh. 1961. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical journal* 1, 6 (1961), 445–466.

[15] Srinjoy Ganguly, Upasana Neogi, Anindya S Chakrabarti, and Anirban Chakraborti. 2017. Reaction-diffusion equations with applications to economic systems. In *Econophysics and Sociophysics: Recent Progress and Future Directions*. Springer, 131–144.

[16] Kiran Garimella and Robert West. 2019. Hot streaks on social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 13. 170–180.

[17] Jeremy Ginsberg, Matthew H Mohebbi, Rajan S Patel, Lynnette Brammer, Mark S Smolinski, and Larry Brilliant. 2009. Detecting influenza epidemics using search engine query data. *Nature* 457, 7232 (2009), 1012–1014.

[18] James W Haefner. 2005. *Modeling Biological Systems:: Principles and Applications*. Springer Science & Business Media.

[19] Elizabeth E Holmes, Mark A Lewis, JE Banks, and RR Veit. 1994. Partial differential equations in ecology: spatial interactions and population dynamics. *Ecology* 75, 1 (1994), 17–29.

[20] Bryan Hooi, Kijung Shin, Shenghua Liu, and Christos Faloutsos. 2019. Smf: Drift-aware matrix factorization with seasonal patterns. In *Proceedings of the 2019 SIAM International Conference on Data Mining*. SIAM, 621–629.

[21] Tianran Hu, Yinglong Xia, and Jiebo Luo. 2019. To return or to explore: Modelling human mobility and dynamics in cyberspace. In *The World Wide Web Conference*. 705–716.

[22] Shubhra Kanti Karmaker Santu, Liangda Li, Yi Chang, and ChengXiang Zhai. 2018. Jim: Joint influence modeling for collective search behavior. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 637–646.

[23] Shubhra Kanti Karmaker Santu, Liangda Li, Dae Hoon Park, Yi Chang, and ChengXiang Zhai. 2017. Modeling the influence of popular trending events on user search behavior. In *Proceedings of the 26th International Conference on World Wide Web Companion*. 535–544.

[24] Koki Kawabata, Yasuko Matsubara, Takato Honda, and Yasushi Sakurai. 2020. Non-Linear Mining of Social Activities in Tensor Streams. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2093–2102.

[25] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[26] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. 2018. Modeling long-and short-term temporal patterns with deep neural networks. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 95–104.

[27] Hemank Lamba and Neil Shah. 2019. Modeling dwell time engagement on visual multimedia. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1104–1113.

[28] Jure Leskovec, Lars Backstrom, and Jon Kleinberg. 2009. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. 497–506.

[29] Haoyang Li, Peng Cui, Chengxi Zang, Tianyang Zhang, Wenwu Zhu, and Yishi Lin. 2019. Fates of Microscopic Social Ecosystems: Keep Alive or Dead?. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 668–676.

[30] Lei Li, B Aditya Prakash, and Christos Faloutsos. 2010. *Parsimonious linear fingerprinting for time series*. Technical Report. CARNEGIE-MELLON UNIV PITTSBURGH PA SCHOOL OF COMPUTER SCIENCE.

[31] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyou Zhou, Wenhu Chen, Yu-Xiang Wang, and Xifeng Yan. 2019. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. In *Proceedings of Advances in Neural Information Processing Systems*. 5243–5253.

[32] Shiyu Liang and R Srikant. 2017. Why deep neural networks for function approximation?. In *5th International Conference on Learning Representations, ICLR 2017*.

[33] Bryan Lim, Sercan O Arik, Nicolas Loeff, and Tomas Pfister. 2019. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *arXiv preprint arXiv:1912.09363* (2019).

[34] SS Antman JE Marsden, L Sirovich S Wiggins, L Glass, RV Kohn, and SS Sastry. 2002. Interdisciplinary Applied Mathematics. (2002).

[35] Yasuko Matsubara, Yasushi Sakurai, and Christos Faloutsos. 2015. The web as a jungle: Non-linear dynamical systems for co-evolving online activities. In *Proceedings of the 24th international conference on world wide web*. 721–731.

[36] Yasuko Matsubara, Yasushi Sakurai, Christos Faloutsos, Tomoharu Iwata, and Masatoshi Yoshikawa. 2012. Fast mining and forecasting of complex time-stamped events. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. 271–279.

[37] Charalampos Mavroforakis, Isabel Valera, and Manuel Gomez-Rodriguez. 2017. Modeling the dynamics of learning activity on the web. In *Proceedings of the 26th International Conference on World Wide Web*. 1421–1430.

[38] Leland McInnes, John Healy, and James Melville. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426* (2018).

[39] Swapnil Mishra, Marian-Andrei Rizoiu, and Lexing Xie. 2016. Feature driven and point process approaches for popularity prediction. In *Proceedings of the 25th ACM international on conference on information and knowledge management*. 1069–1078.

[40] Maya Okawa, Tomoharu Iwata, Yusuke Tanaka, Hiroyuki Toda, Takeshi Kurashima, and Hisashi Kashima. 2021. Dynamic Hawkes Processes for Discovering Time-evolving Communities' States behind Diffusion Processes. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 1276–1286.

[41] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499* (2016).

[42] John E Pearson. 1993. Complex patterns in a simple system. *Science* 261, 5118 (1993), 189–192.

[43] B Aditya Prakash, Alex Beutel, Roni Rosenfeld, and Christos Faloutsos. 2012. Winner takes all: competing viruses or ideas on fair-play networks. In *Proceedings of the 21st international conference on World Wide Web*. 1037–1046.

[44] Julia Proskurnia, Przemyslaw Grabowicz, Ryota Kobayashi, Carlos Castillo, Philippe Cudré-Mauroux, and Karl Aberer. 2017. Predicting the success of online petitions leveraging multidimensional time-series. In *Proceedings of the 26th International Conference on World Wide Web*. 755–764.

[45] Maziar Raissi and George Em Karniadakis. 2018. Hidden physics models: Machine learning of nonlinear partial differential equations. *J. Comput. Phys.* 357 (2018), 125–141.

[46] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics* 378 (2019), 686–707.

[47] Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. 2018. Deep state space models for time series forecasting. In *Proceedings of Advances in neural information processing systems*. 7785–7794.

[48] Bruno Ribeiro. 2014. Modeling and predicting the growth and death of membership-based websites. In *Proceedings of the 23rd international conference on World Wide Web*. 653–664.

[49] Manoel Horta Ribeiro, Kristina Gligorić, Maxime Peyrard, Florian Lemmerich, Markus Strohmaier, and Robert West. 2021. Sudden attention shifts on wikipedia during the covid-19 crisis. In *Proc. Int. AAAI Conf. Web Soc. Media*, Vol. 15. 208–219.

[50] Jorma Rissanen. 1978. Modeling by shortest data description. *Automatica* 14, 5 (1978), 465–471.

[51] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. 1986. Learning representations by back-propagating errors. *nature* 323, 6088 (1986), 533–536.

[52] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting* 36, 3 (2020), 1181–1191.

[53] Chun Tang and Aidong Zhang. 2003. Mining multiple phenotype structures underlying gene expression profiles. In *Proceedings of the Twelfth International Conference on Information and knowledge management*. 418–425.

[54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of Advances in neural information processing systems*. 5998–6008.

[55] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S Dhillon. 2016. Temporal regularized matrix factorization for high-dimensional time series prediction. *Advances in neural information processing systems* 29 (2016).

[56] Le Zhang, Hengshu Zhu, Tong Xu, Chen Zhu, Chuan Qin, Hui Xiong, and Enhong Chen. 2019. Large-scale talent flow forecast with dynamic latent factor model?. In *The World Wide Web Conference*. 2312–2322.

[57] Yuchen Zhao, Neel Sundaresan, Zeqian Shen, and Philip S Yu. 2013. Anatomy of a web-scale resale market: a data mining approach. In *Proceedings of the 22nd international conference on World Wide Web*. 1533–1544.

[58] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of AAAI*.