

How Powerful are Interest Diffusion on Purchasing Prediction: A Case Study of Taocode

Xuanwen Huang[†], Yang Yang^{*†}, Ziqiang Cheng[†], Shen Fan[§], Zhongyao Wang[§]
Juren Li[†], Jun Zhang[§], Jingmin Chen[§]
[†]Zhejiang University, [§]Alibaba Group
{xwhuang, yangya, petecheng, jrlee}@zju.edu.cn
{fanshen.fs, zhongyao.wangzy, zj157077, jingmin.cjm}@alibaba-inc.com

ABSTRACT

A *Taocode* is a kind of specially coded text-link on *Taobao.com* (the world’s biggest online shopping website), through which users can share messages about products with each other. Analyzing *Taocodes* can potentially facilitate understanding of the social relationships between users and, more excitingly, their online purchasing behaviors under the influence of *Taocode* diffusion. This paper innovatively investigates the problem of online purchasing predictions from an information diffusion perspective, with *Taocode* as a case study. Specifically, we conduct profound observational studies on a large-scale real-world dataset from *Taobao*, containing over 100M *Taocode* sharing records. Inspired by our observations, we propose *InfNet*, a dynamic GNN-based framework that models the information diffusion across *Taocode*. We then apply *InfNet* to item purchasing predictions. Extensive experiments on real-world datasets validate the effectiveness of *InfNet* compared with 8 state-of-the-art baselines.

CCS CONCEPTS

• **Information systems** → **Online shopping; Electronic commerce; Social networks.**

KEYWORDS

purchasing prediction, information diffusion, GNN, *Taocode*

ACM Reference Format:

Xuanwen Huang[†], Yang Yang^{*†}, Ziqiang Cheng[†], Shen Fan[§], Zhongyao Wang[§], Juren Li[†], Jun Zhang[§], Jingmin Chen[§]. 2021. How Powerful are Interest Diffusion on Purchasing Prediction: A Case Study of Taocode. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR ’21), July 11–15, 2021, Virtual Event, Canada*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3404835.3462898>

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR ’21, July 11–15, 2021, Virtual Event, Canada

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8037-9/21/07...\$15.00
<https://doi.org/10.1145/3404835.3462898>

1 INTRODUCTION

Online shopping has become an increasingly common practice for thousands of households. As reported by *Alibaba Group*¹, in the fiscal year of 2020, the mobile MAUs across the China retail marketplaces comes to 846 million, and the total GMV of *Taobao*² reaches 3,387 billions of CNY. Due to the extremely large scale of online retail marketplaces, it has become crucial for e-commerce platforms to better understand users’ online purchasing behaviors. One classical but challenging problem in this area is that of purchasing predictions [8]: given a user, along with a specific product or a category of goods, the question is how likely the user would be to buy this item.

Existing works have made great efforts in this field. Traditional methods mainly consider users and products by matrix factorization, statistical learning models [26], etc. Deep learning-based models adopt deep neural networks to integrate the diverse information in the e-commerce system [33, 36, 39, 44], such as users’ clicking, browsing and purchasing histories. Most of these previous frameworks utilize the historical information of both users and products; however, few of them have studied the impact of information diffusion [41, 42] across the social network on users’ online purchasing behaviors [21, 22, 43]. Taking a concrete example in practice, a user will be more likely to buy a product at a time when she actually needs it, or alternatively, when some close friends recommend that item to her (even though she might not have any urgent demand for it). Furthermore, such preferences for different products can spread across online social networks: a user who receives a recommendation from her friends may share this item with other users within her social circle [28], thus promoting further possible purchases, while the scale of sharing could explode as the message is propagated throughout the social relationships between users. We refer to this phenomena as diffusion of user interests in products, i.e., “*interest diffusion*” for short, which is likely to influence the purchasing behavior of users to some degree.

Although such social factors are expected to be highly relevant to users’ purchasing decisions, they have been rarely studied in previous works for two main reasons. First, interest diffusion across online e-commerce platforms is always *implicit* and *hard to observe*. While we can easily obtain users’ browsing, purchasing and product review histories, or even discussions on public social media platforms, the majority of recommendations and interest diffusion occur between acquaintances via private social communications,

¹See the FY 2020 annual report on <https://www.alibabagroup.com/en/ir/reports>.

²the biggest online shopping website in the world that belongs to Alibaba Group.

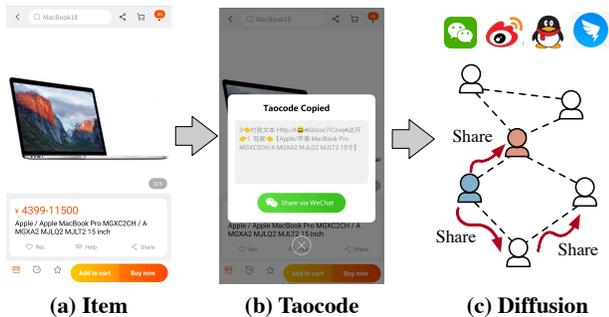


Figure 1: An illustration of Taocode diffusion process. When users browse an item on Taobao (a), they can share its homepage to friends by generating the Taocode message (b), and sending it via social media. The Taocode then spread over the social network (c).

which are difficult to discover and analyze. Indeed, previous works have explored the impact of product sharing on users’ shopping behavior through analysis of email data [17, 20]; however, the social cost of using emails to share goods is much higher than using instant social tools, meaning that these studies may be biased in real-world scenarios. Second, interest diffusion across social networks is very complicated to model, since diffusion trajectories change naturally at any time with the spread of social interactions, and on online shopping platforms, user interest in different products could exhibit various diffusion patterns. Modeling the dynamics of interest diffusion remains a challenging research problem.

Excitingly, *Taocode*, a kind of specially coded text-link on Taobao, provides us with an appropriate testbed to observe and capture the social influences and product interest diffusion between users in the online shopping scenario. *Taocode* is essentially a hyper-link containing the information of some specific product on Taobao. Fig. 1 illustrates how the *Taocode* is created and spread. It is convenient for users to create and share a *Taocode* by clicking specific options in Taobao when they browse an item, and anyone who receives the *Taocode* can browse this product by opening the embedded hyper-link. More importantly, since *Taocodes* spread within Taobao, the diffusion process of each sharing message can be recorded in its entirety; we can trace the *explicit sharing path* of the item on the diffusion flow, and obtain the details of which user shares which product with whom at what time. This provides us with a great opportunity to model the dynamics of product sharing diffusion.

This paper presents a large-scale study on the effect of product interest diffusion on online shopping platforms, taking *Taocodes* sharing on Taobao as an example. It is pioneering work from the perspective of user interest diffusion to predict item-level purchases. We sample a total of 100M *Taocode* sharing records, and then construct the *interest diffusion networks* from those *Taocode* records at different time steps: there is a unique diffusion network within each single time span, where each node represents a user, and the directed edges between two nodes denote the relationships of *Taocode* sharing. Moreover, the dynamics of the network structures reveal the message diffusion across the edges, which may reflect the purchasing preferences of users regarding different groups of items. Given a specific query of a user-item pair, our goal is to estimate how likely the user is to buy this product based on purchasing and

Taocode diffusion histories, by utilizing the sequence of interest diffusion networks during the given time span.

As discussed above, modeling interest diffusion networks remains challenging: it is difficult for deep sequential models to handle geometric data, while most existing GNNs may be unable to deal with the dynamic edge weights at different times. To deeply understand how *Taocode* messages are spread over the social network, we first conduct empirical observations across the entire *Taocode* diffusion (Sec. 3), which provides several comprehensive insights: 1) *Taocodes* sharing behaviors are indeed highly correlated with users’ purchasing decisions, in that users who send or receive more *Taocodes* are likely to make more purchases; 2) the influence of *Taocode* diffusion varies across different product categories, e.g., items with higher price index are more likely to be purchased after being shared via *Taocodes*; 3) the graph structures of interest diffusion networks reflect users’ purchasing preferences: for instance, users whose neighbors share *Taocodes* more actively tend to have higher purchasing rates; 4) temporal factors play an important role in *Taocode* diffusion, i.e., the time interval between *Taocode* senders and sharing the item has implicit impacts on the receivers.

Motivated by these findings, along with the advances of sequence- and graph-level attentions, we propose *InfNet* to integrate both the structural and temporal information of dynamic attributed diffusion networks. The key hints derived from our observations are that, on one hand, a user’s preference for different items is highly related with her own and her neighbors’ interest diffusion patterns, while on the other hand, those patterns are linked with time and are affected by neighborhoods’ dynamic product interests at different time steps. In other words, the structural and temporal dynamics of interest diffusion change in an interrelated fashion, and should thus be incorporated into a unified framework. Accordingly, we adopt a multi-level attention mechanism to model the interest diffusion patterns, i.e., two graph attentions on diffusion networks, and a self-attention-based encoder on diffusion sequences. Extensive experiments applying purchasing predictions to real-world datasets on Taobao demonstrate the effectiveness of our proposed model: *InfNet* significantly outperforms 8 state-of-the-art baselines, ranging from session-based models to GNN-based social recommendation frameworks. We also conduct ablation studies to determine how different levels of attention work, and visualize several cases to illustrate the diffusion patterns of products extracted by *InfNet*. Overall, the contributions of this paper can be described as follows:

- We present a large-scale study on the effect of product sharing in online shopping scenarios, and investigate the problem of online purchasing predictions from an information diffusion perspective. More specifically, we take *Taocode*, a specially coded type of text link of Taobao, as a case study.
- Comprehensive observations over 100M *Taocode* sharing records illustrate several implicit characteristics of product interest diffusion across the social network on Taobao.
- Motivated by our observational insights, we design an end-to-end framework, named *InfNet*, to make item purchasing predictions. Experimental results validate the effectiveness of the way, in which we adopt multi-level sequential and graph attentions to capture the users’ preferences for different products and the interest diffusion dynamics.

2 PROBLEM FORMULATION

To clearly illustrate the proposed framework focusing on the purchasing predictions from the perspective of information diffusion, we first give the detailed and formal definitions of interest diffusion based on *Taocode* sharing and the problem formulation.

We refer to the set of users and products in our dataset from *Taobao* as \mathcal{V} and \mathcal{P} respectively, and further denote $t_i \in T = \{t_1, t_2, \dots, t_n\}$ as the time step; here, n is the number of the observed time step, and each $t_i = (t_b^i, t_e^i)$ represents a specific period of time, starting at t_b^i and ending at t_e^i .

DEFINITION 1. *Taocode diffusion.* Each *Taocode* sharing message represents a diffusion record, which we formulate as a 4-tuple $d \in \mathcal{D}_t = \{(u, v, p, t_d) \mid u, v \in \mathcal{V}, p \in \mathcal{P}, t = (t_b, t_e) \in T\}$, satisfying $t_b \leq t_d \leq t_e$. Here, u and v denote the receiver and sender respectively of the *Taocode* message, p is the item that is shared, and t_d is the time at which the *Taocode* is received by u from v .

DEFINITION 2. *Interest diffusion network* is a directed attributed network $\mathcal{G}_{\mathcal{D}} = (\mathcal{V}, \mathcal{E}, \mathcal{H}(\mathcal{D}))$ constructed from a *Taocode* diffusion set \mathcal{D} , where each node $u \in \mathcal{V}$ represents an online shopping user with feature vector \mathcal{H}_u , while the edge $e \in \mathcal{E}$ indicates the interest diffusion messages between users, with the features H_e denoting the edge attributes. Formally, there exists an edge e_{uv} between two users u and v , if the tuple (u, v, p^*, t^*) is an element in \mathcal{D} for some item p^* and timestamp t^* , and the attribute $H_{e_{uv}}$ of the edge e_{uv} contains the detailed information of the product p^* , e.g., *PI*, categories, etc.

DEFINITION 3. *Dynamic interest diffusion network* is a sequence of interest diffusion networks $\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_n\}$ constructed from *Taocode* diffusion sets $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_n\}$, where i indicates the time step t_i . For each time step t , we collect all *Taocode* diffusions \mathcal{D}_t that occurs during the given time span (t_b, t_e) from the data source, and add edges into the network \mathcal{G}_t as described in Definition. 2.

Problem definition. Considering a dynamic interest diffusion network \mathcal{G} , along with a query $q = (v, p, t_i)$ that $v \in \mathcal{V}, p \in \mathcal{P}, t_i \in T$, the problem we aim to solve is that of determining whether the user v will purchase product p within the time span t_i . Since the total number of items on an online shopping platform (*Taobao.com* in our experiment) is very large, to simplify the problem and investigate the influence of *Taocode* diffusion, we implement a restriction that v should receive *Taocode* messages containing the product p from some other users during the time span t_{i-1} . We do not place limitations on that whether the user v has received the *Taocode* before t_{i-1} ; thus, both the short- and long-term impact of *Taocode* diffusion are expected to be modeled to achieve better performance. The additional benefit is that the number of possible queries is significantly reduced, which makes it feasible to conduct comparison experiments between selected baselines.

3 OBSERVATIONAL STUDIES

In this section, we provide an in-depth examination of a large-scale dataset of *Taocode* sharing records from three different perspectives; this is done to facilitate a better understanding of *Taocode* diffusion’s intrinsic properties, along with its influences on the users’ purchases.

3.1 Dataset descriptions

The dataset for our observational studies is described as follows: we sample records of *Taocode* sharing between *Taobao* users during a specific period of time in 2020. The entire dataset contains over 100 million *Taocode* sharing records. Each record consists of basic information, including: 1) who creates and shares the *Taocode*, namely the *sender*; 2) who receives the *Taocode*, namely the *receiver*; 3) when the receiver opens the *Taocode* link, namely the *timestamp*, and 4) the details of the item that this *Taocode* represents, including its *Price Index* (PI)³ and the category it belongs to.

To determine the correlations between sharing and purchasing, we extract the purchasing logs related to the *Taocode* sharing records. In fact, there is no explicit evidence to indicate that a user buys an item as a direct result of receiving the *Taocode*; thus, we adopt the implicit influences of *Taocode* by collecting the purchasing records of the receiver and the item that each *Taocode* message contains, and assume that *Taocode* sharing is the main reason that leads to the purchases if the receiver buys that item *after* receiving a *Taocode* within a predefined time span. Note that for the purchasing predictions, we collect another group of *Taocode* samples, and the above-mentioned dataset is only used for the observations, feature extractions for users and items, along with the *Taocode* diffusion network construction. We will introduce the “prediction dataset” for the purchasing predictions in Sec. 5.

3.2 *Taocode*’s effect on purchases

The first research question is: **Does *Taocode* diffusion have impacts on users’ purchasing behaviors? (Q1)** To directly answer the question, we introduce the *conversion index* (CI):

DEFINITION 4. *The Conversion Index (CI) of an item p is the proportion of its purchasing records to the total, i.e.,*

$$f\left(\frac{\#(\text{purchasing records of } p)}{\#(\text{all records of } p)}\right)$$

f is a monotonically increasing scalar function that maps the origin proportion to another value for data privacy issues. It can be formulated under different conditions: (1) *CI of browsing*, where records are browsing histories; (2) *CI of Taocode*, where records are *Taocode* messages. We further denote the ratio of *Taocode* sharing *CI* and browsing *CI*, i.e., $\frac{CI(\text{Taocode})}{CI(\text{sharing})}$, as *CI Lift*, to reveal the effect of *Taocode* sharing on purchasing behaviors. If an item has a high *CI Lift*, it is much more likely that it will be purchased after being shared via *Taocode*.

We compare the averaged *CI* of both the *Taocode* sharing and product browsing records for all products. The first of our findings is that items shared via *Taocode* are much more likely to be purchased: a 31.0% higher *CI* on average (p-value<0.01) than that with browsing records alone. Meanwhile, we group the items by their leaf categories, count the averaged *PI* among each group, and compare their *CI Lift*. As shown in Fig. 2(a), there is a statistically significant positive correlations between *CI Lift* and averaged *PI* (slope=0.0448, p-value<0.01). The Spearman correlation between *PI* and *CI Lift* is 0.336 (p-value<0.01). In other words, as the item’s *PI* increases, its *CI Lift* becomes higher; i.e., *Taocode* sharing can

³Price Index, defined as $PI = f(\text{price})$, reflecting the price level of the product, where $f(\cdot)$ is a monotonically increasing non-linear scalar mapping due to privacy concerns.

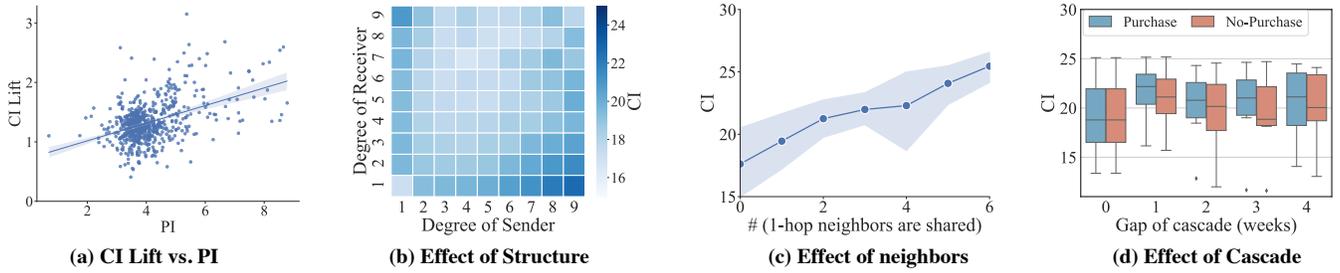


Figure 2: Observational studies on *Taocode* diffusion. (a) illustrates the positive correlations between the PI and CI Lift of items, indicating that *Taocode* has a stronger influence over expensive goods. Both (b) and (c) reflect the structural properties of *Taocode* diffusion; for example, the numbers of *Taocode* messages being sent and received by a user and her neighbors would affect her purchasing behavior. (d) finds that *Taocode* messages could have a long-term impact on users’ purchases.

further promote its purchase rate, which indicates that when buying expensive goods, users often engage in more consultation before making a final decision. As the only condition we control is that of whether the products are shared via *Taocode*, we conclude that ***Taocode* diffusion does indeed affect users’ purchasing behaviors**: users will be more willing to make a purchase if they receive the *Taocode* message from someone else.

3.3 Structural dynamics of *Taocode*

We next investigate how *Taocode* affect the purchases in the following two subsections, starting with the following question: **Are there any structural characteristics of *Taocode* diffusion patterns? (Q2)** We first examine the structural properties of *Taocode* diffusion networks. We merge all edges among each dynamic interest diffusion network \mathcal{G}_t at every time step t into one graph \mathcal{G} , and calculate the sum of the in and out degree with respect to different item categories for each user. Since more than 80% users have a degree between 1 and 9, and the degree distribution exhibits a long-tail shape, we only count *Taocodes* where the degrees of both sender and receiver lie between 1 and 9. The results are shown in Fig. 2(b): the average CI of all *Taocode* records is 19.0, reaching a maximum of 22.7 when the sender’s degree is 9 and the receiver’s is 1, and 20.9 when the degrees of the sender and receiver are 1 and 9 respectively. It is consistent with previous findings in the literature that “consumers who are central in networks are quite susceptible to others’ influences” [19]. However, CI of *Taocodes* where the degrees of both sender and receiver are 1 or 9 are only 17.2 and 17.9 (much lower than average). Contradicting the initial assumption that *Taocode* sharing between high degree users would have a higher CI, we find that when the gap of degrees between the sender and receiver is larger, CI becomes higher. This may be explained by noting that a user’s online purchases are likely to happen more frequently if she interacts socially with another user who has a very different pattern of social behaviors. For example, if a user u who often shares products with others, sends a *Taocode* message to user v that who uses *Taocode* relatively infrequently, one possible motivation for this sharing is that the item is of urgent demand for v ; thus, v is more likely to make a purchase than usual.

Next, we investigate the influence of the receiver’s neighbors in the *Taocode* diffusion. More specifically, for a *Taocode* message (u, v, p, t) , we look at v ’s 1-hop neighbors who have received and

bought the same product p , denoted as *close neighbors*. The intuition here is that, the number of *close neighbors* of the receiver can reflect the popularity of the item during the diffusion flow to some degree. We group *Taocode* records by the number of *close neighbors* of the receivers, and compare their average CI across different product categories, as illustrated in Fig. 2(c). We find that, in general, the more close neighbors the receiver has, the more likely it is that the receiver will buy the item; i.e., the average CI increases from 18.8 to 25.3 as the x -axis becomes large. We can therefore conclude that users would have a higher CI on an item if many of their neighbors are in the diffusion flow of this product, (i.e., receive the *Taocode* message of this item and buy it). Our observations are consistent with [17], who found that users with social connections may have similar shopping behaviors. In conclusion, **the network structure of interest diffusion networks may help us in modeling user purchasing behaviors.**

3.4 Temporal dynamics of *Taocode*

Finally, we tackle the temporal dynamics of *Taocode* diffusion: **Does the senders’ past receiving history do have impact on receivers’ purchasing? (Q3)**

In the literature, a diffusion cascade can be defined as a information flow, as item p flows by $u_1 \rightarrow u_2 \rightarrow u_3 \rightarrow \dots \rightarrow u_k$. To simplify the diffusion process, we only look at the 2-hop cascades that end with the user u_3 in the above example. For convenience of explanation, we refer to the target user as u_3 , and his/her 1-hop and 2-hop neighbor on an in-edge of a diffusion path as u_2 and u_1 respectively. We aim to figure out 1) Does the 1-hop neighbor(u_2)’s receiving records have impact on the receiver(u_3)’s purchasing and 2) How does the impact changes over time?

We present the observational result in Fig. 2(d), where x -axis δ donate the time gap(weeks) between the time when u_2 received the same *Taocode* and the current *Taocode* sent to u_3 , note that if u_2 does not receive the same *Taocode* before, the δ is set to 0. while y -axis donate the *conversion index*. The first of our finding is that the average CI of *Taocode* with $\delta = 0$ (no 2-hop neighbors) is 19.1 ± 4.32 , which is lower than the case when 2-hop neighbors exists($\delta > 0$). Moreover, there are no significant differences in the averaged CI between different time gaps δ : the results of remaining groups ($\delta = 2, 3, 4$ weeks) are in 19.5 ± 4.49 , 19.2 ± 4.59 , and 19.99 ± 4.16 respectively. However when it comes to the case in which the sender buys the item later (the blue bars), we can observe that the average

CI is much higher than that of the red bars, even if the previous *Taocode* sharing occurs more than four weeks ago (20.4 ± 4.07 , vs. 19.99). This suggests that **both the long- and short-term temporal dynamics of *Taocode* diffusion do indeed exist**. This is consistent with the common-sense conclusion that (1) the more frequently *Taocode* messages are spread within a short time span, the more likely it is that the item has great popularity with a higher CI; (2) alternatively, a user u will probably share some product with others long after she receives, and after spending some time making sure that the item deserves the recommendation. Thus, the temporal dynamics of *Taocode* diffusion could be an important factor in influencing user purchasing behaviors.

4 MODEL: INFNET

In this section, we introduce the proposed model *InfNet* for modeling the structure and temporal dynamics of *Taocode* diffusion in the interest diffusion networks. Given the input of a purchasing query at a specific time step that contains the target user and item, our framework first builds a sequence of “dynamic sub-networks” for each query, after which we apply two different attention-based aggregation strategies on dynamic node and edge attributes to learn users’ hidden representations. Finally, the user’s embedding, along with the item features are fed into the outer classifier, to predict whether the target user will purchase the item during the given time span. An overview of the framework is presented in Fig. 3.

4.1 Preprocessing: Sub-graph construction

For each purchasing query $q = (u, p, t_i)$ and the entire dynamic interest diffusion network $\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_{t_i}\}$, we build a sequence of sub-graphs \mathcal{G}^q of \mathcal{G} to conduct the end-to-end graph classification, i.e., the purchasing prediction of query q . The process of constructing the sub-networks can be described as follows :

- **Seed node sampling:** We define the target user u and its 1-hop neighbors $\mathcal{S}_u(p)$ as the set of seed nodes, denoted as $\{u\} \cup \mathcal{S}_u(p)$, where $\mathcal{S}_u(p)$ is the set of *sender nodes* that shares item p with u .
- **Sub-network construction:** Given the set of seed nodes, we adopt a breadth-first search (BFS) strategy to sample the nodes from each original diffusion network $\mathcal{G}_t (1 \leq t \leq t_i)$, by searching for two-hop neighbors, i.e., the search depth is set to two; and moreover, an edge exists between two nodes (u, v) in \mathcal{G}_t^q if the user u shares some product p with v at time step t .
- **Node and edge attributes:** The node features of a user u , denoted as $\text{Attr}(u)$, consists of two parts: the first part is the user’s purchasing histories which are independent from the diffusion networks, i.e., the price she pays for different products in the past year; the other part is a one-hot vector, indicating whether this node is the target user in the query or it is among the seed nodes, which can help the model better identify the nodes that it should focus on. As for the edge features, we first define the item feature $\text{Attr}(p)$ as a one-hot vector, where each bin represents a specific range of prices. Then for an edge (u, v, t) , its features e_{uv}^t is summed over features of all items that are sent from the user u to v during the period t . Note that edge features may change across queries at different time, since for a specific time span t , we only aggregate features of items that are shared during t .

4.2 Structural block

As Fig. 2 (b) shows, the local graph structure in the diffusion network is an important factor to infer whether *Taocode* sharing can prompt purchases. Therefore, we use a *Structural block* to learn the users’ embeddings to capture such structural characteristics. *Structural block* can be implemented by any GNNs with node aggregation strategy, which can embed nodes’ structural information [40]. It can be denoted by a uniform Message-Passing framework: for a specific user u , the first layer is a non-linear transformation on user attributes, and output the initial hidden representation x^u :

$$x_u^k = \text{Combine}(x_u^{k-1}, \text{Aggr}(x_v^{k-1} | v \in \mathcal{N}_u)) \quad (1)$$

where k is the layer index, $\text{Combine}(\cdot, \cdot)$ combines hidden states of u with her neighbors, and $\text{Aggr}(\cdot)$ denotes the aggregation function (sum or mean operation). By this step, the representation vector x_u^k is expected to contain the information of both u and her neighbors. After stacking K_1 layers of Eq. (1), the output is the final structural embedding of the user u , denoted as $s_u = x_u^{K_1}$.

4.3 Diffusion block

While *Structural block* only aggregates on the node attributes, edge features encode much more information of *Taocode* diffusion: as Fig. 2 (c) and (d) suggest, if we want to better predict the user u ’s purchasing decision on the item p , we may need to know how many close neighbors of u , as well as whether there exist some users who have shared p to u having purchased this item. Moreover, the temporal information of *Taocode* diffusion can also help us capture both short- and long-term influence of such social interactions on users’ purchasing behavior. However, existing methods rarely consider the dynamic item diffusion among users: most of them apply frameworks similar to Eq. (1), which is unsuitable for a *Taocode* diffusion network. Consider a situation, in which user v has received many *Taocodes* of computers and clothes, and sends a clothing *Taocode* to u . The diffusion from v to u pertains only to clothes; therefore, node information of v related to computers may evidently introduces noise that hampers the prediction of u ’s purchase behavior. So we further design *Diffusion block* to fully model the *Taocode* diffusion via edge-level aggregations and learn users’ final representations.

Inspired by real-world situations and our observations (Sec. 3), we design a special multi-layer aggregation strategy. The key idea is that, rather than directly aggregating neighbors’ node features, we first learn the weight of each edge, then aggregate the edge features between the user and her neighbors to get an updated representation. Through the attention mechanism, neighbors’ node representations will affect the nodes’ final representation in an implicit way, and it can avoid noisy aggregations.

Formally, let e_{uv}^t denotes the edge feature of user v sent to user u at time t . To capture the dynamics of *Taocode*, we independently aggregate the attributes of edges connected with the user to get the hidden state of each time step, denoted as $h_u^{k,t}$, where k is the layer number. Each hidden state $h_u^{k,t}$ only contains all interplay on social networks involving u . We then utilize a sequence encoder to combine users’ hidden states in different time steps to get the dynamic representations of users. In detail, the same user at different time step has the same initial hidden state $h_u^{0,t}$, which is obtained by a non-linear transformation of user attributes $\text{Attr}(u)$; then for

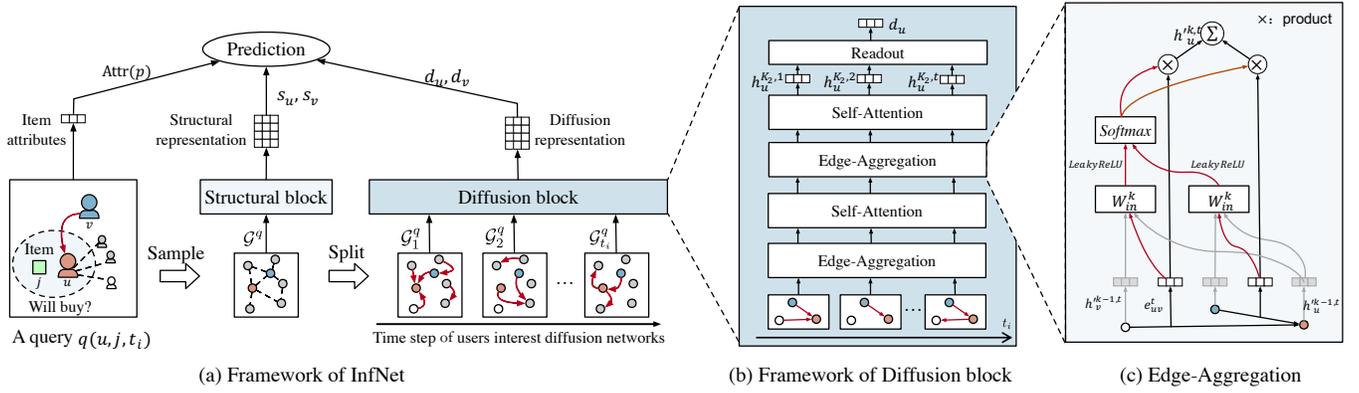


Figure 3: Framework of *InfNet*. (a) is the overview of *InfNet*, (b) is the framework of diffusion encoder in *InfNet*. (c) is the detail of edge-level aggregation. Note, in (c), user’s hidden states only used to count attention weight.

user u in k^{th} layer at time step t , we conduct separate aggregations of users’ in- and out-edges and obtain two representations $In_u^{k,t}$ and $Out_u^{k,t}$. They represent the information users receive and send. We combine them to get the hidden state $h_u^{k,t}$ at time step t :

$$In_u^{k,t} = \sum_{v \in S_u^t} \text{ATT}_{in}^k(h_u^{k-1,t}, h_v^{k-1,t}, e_{uv}^t) e_{vu}^t \quad (2)$$

$$Out_u^{k,t} = \sum_{v \in \mathcal{R}_u^t} \text{ATT}_{out}^k(h_u^{k-1,t}, h_v^{k-1,t}, e_{vu}^t) e_{vu}^t \quad (3)$$

$$h_u^{k,t} = [In_u^{k,t}, Out_u^{k,t}] \quad (4)$$

Here $\text{ATT}_*^k(*)$ is an attention function that takes the user’s hidden states in the last layer, as well as edge attributes as inputs, and output the attention weights of the edge attributes. It fuses these three inputs, then employs a non-linear transformation to obtain the scalar values. The attention mechanism can be clarified as:

$$w_{uv}^{k,t} = \text{LeakyReLU}(W_{in}^k[h_u^{k-1,t}, h_v^{k-1,t}, e_{uv}^t]) \quad (5)$$

$$\text{ATT}_{in}^k(h_u^{k-1,t}, h_v^{k-1,t}, e_{uv}^t) = \frac{\exp(w_{uv}^{k,t})}{\sum_{c \in S_u^t} \exp(w_{uc}^{k,t})} \quad (6)$$

where $W_*^k \in \mathbb{R}^{1 \times 3c}$, and we adopt a softmax function [31] on edges to make the summed weights of all u ’s neighbors equal to 1. This operation can alleviate over-fitting and facilitate fair comparison of various edges. The out-edge attention is the same as that of in-edge, except that in Eq. (6), it is summed over \mathcal{R}_u^t rather than S_u^t .

After obtaining the hidden state at each time step t , i.e., a sequence of representations of different time steps for each user u , we use a sequence encoder to capture the temporal dynamics:

$$h_u^{k,t} = \text{Enc}(h_u^{k,i} | i \leq t) \quad (7)$$

The $\text{Enc}(*)$ can be a simple operator such as Mean-pooling or a more complex operator such as GRU [3], Masked self-attention [30], etc. However, as this aspect is not our main focus, here we do not expand on it in detail, and will analyze the performance of different encoding methods in Sec. 5.

After K_2 layers of *Diffusion block*, we obtain the final dynamic representation $d_u = \text{Readout}(h_u^{K_2,t} |_{t=1}^{t=t_i})$, where $\text{Readout}(*)$ denotes readout the overall representation of sequence. For example,

we can directly readout the representation in last time step in GRU [3], or readout by sum with weights in Self-Attention [5].

4.4 Model learning

After we obtain the users’ structural and *Taocode* diffusion representations, we combine them to predict the users’ purchasing behavior: whether user u will purchase item p , namely query $q = (u, p)$. We focus not only on u ’s representation and item p ’s attributes, but also consider the users who send item p to u :

$$gq = \sum_{v \in S_u(p)} [s_v, d_v] \quad (8)$$

Here $S_u(p)$ is the set of seed users (Sec. 4.1), in which all users are u ’s neighbors who send item p to u . We then combine these three parts and make the predictions through a simple output layer:

$$y' = \sigma([d_u, s_u, gq, \text{Attr}(p)]) \quad (9)$$

where $\text{Attr}(p)$ denotes the features of item p .

In general, the downstream task can be regarded as a binary classification problem, thus we use *cross-entropy* as the loss function:

$$\mathcal{L}(\Theta) = \sum_{q_i \in Q} -y_i \log y'_i + (1 - y_i) \log(1 - y'_i) \quad (10)$$

Here, Θ denotes all learnable parameters in *InfNet*, Q is the set of all queries, and y_i is the ground truth of the query q_i .

5 EXPERIMENTS

5.1 Experimental setup

Prediction Datasets. We select *Taocode* sharing messages from six categories for our experiments. In more detail, in order to formulate the purchasing prediction problem with clear input and avoid data selection biases, we collect 30,000 *Taocode* records for each category on each day during the three days following the observational time window in Sec. 3, with a total of 540,000 samples, referred to as “the prediction datasets”.

We regard the receiver and the item of each *Taocode* record as a purchasing query q , which asks whether the user will buy this item within a specific time span. We then inquire into the purchasing logs during the time period to label these queries (i.e., whether or not the user will buy the item). For each query sample, the

Table 1: Statistics of the prediction datasets. Due to data security requirement, we rename the categories into ranked ids.

Category	# Users	# Items	# Edge	# Purchase
Cate#1	1,261,486	33,184	2,730,610	21,633
Cate#2	1,289,775	38,776	2,832,443	29,538
Cate#3	664,732	5,196	1,396,300	80,278
Cate#4	997,957	19,632	2,160,293	11,657
Cate#5	891,134	74,487	1,911,424	9,181
Cate#6	1,267,316	47,920	2,806,235	41,008
All	6,372,400	21,9195	13,837,305	122,739

input for our model is the dynamic diffusion sub-network defined in Sec. 4.1, within the given time period. More specifically, we divide the whole time span into four equal segments, meaning that there are a total of four diffusion sub-graphs constructed for each query with one week as a time step. To keep the offline experiment consistent with the online situation, where we could not obtain any information ahead of the querying time when making predictions, we randomly split the datasets on the first two days into training and validation sets with a ratio of 7:3, with all queries on the third day are used for testing. Overall statistics of the datasets for the item-level purchasing prediction task are presented in Table 1.

Baselines. We compare our proposed model, *InfNet*, with several groups of state-of-the-art baselines:

- Traditional methods. We first compare our model with a classical machine learning method: Logistic Regression (LR) based on feature engineering. We also choose one of the most popular frameworks in the item predictions, Bayesian Personalized Ranking (BPR) [23], which proposes a maximum posterior estimator as the generic optimization criterion for personalized ranking.
- Session-based methods. Another line of online purchasing predictions works are the session-based models. Since *Taocode* message sharing sequence can be naturally regarded as user actions within a session, we choose three popular session-based baselines: SR-GNN [38], MGNN-SPred [33] and GCE-GNN [35].
- Social recommendation methods. To verify the effectiveness of *InfNet* in capturing social dynamics across the interest diffusion networks, we compare three popular social recommendation frameworks: EATNN [2], GraphRec [7] and DiffNet [36]. EATNN introduces attention mechanisms to model users’ preferences and assigning a personalized transfer scheme for each user. GraphRec incorporates user-item and user-user graphs, and uses attentions to model the importance of different social relationships. DiffNet considers the social diffusion by applying GCN [16] on the social networks and SVD++ to the item recommendations.

Evaluation metrics and implementation details. We evaluate the predictive performance of *InfNet* and baseline methods in terms of AUC for the Precision-Recall (PR) and ROC curves, as these metrics are widely applied in recommendation systems [4].

In order to fairly compare the performance of different models, and to adjust to our problem settings, we make some modifications to the deployment of baselines:

- LR based on feature engineering. We carefully extract three categories of features: (1) cost-related features, which contain the

price of the product in the query and the total spending of the user in the past 30 days; (2) graph-related features, which contains the in- and out-degree of users in the diffusion network; (3) features of users’ historical *Taocode* sharing records, i.e., the information about products the user has shared and received.

- Session-based methods. We regard users’ *Taocode* sharing records during the given time span as a session (these are used to extract node and edge features in *InfNet*), then add the price-related features when encoding the item information as a one-hot vector.
- Social-based recommendations. For all GNN-based models, the input is the static interest diffusion network by aggregating all dynamic sub-networks in *InfNet*, which are constructed from users’ previous sharing records. Moreover, since above-mentioned methods are based on static graphs and features, to fairly compare their performance with *InfNet*, we design the *InfNet-S*, where all user and item features are completely consistent with the social-based recommendation baselines.

Besides, we use the Adam [15] optimizer with a learning rate of 0.001, and the batch size is set as 512. In *InfNet*, we search the hidden sizes c in the range [16, 32, 64, 128], and $c = 64$ reaches the best performance. For the GNNs in Eq. (1), we have tried GIN [40], GCN [16] and GAT [31]. We find that there are little differences on the performance of different GNNs models, and choose GAT whose average performance is the best among all candidates. For the sequence encoder in Eq. (7), we select Masked Self-Attention (Sec. 5.3), where the Readout function is similar to [31].

5.2 Experimental results

We present the performance of all methods on item-level purchasing prediction task on the *Taocode* dataset in Table 2. Overall, *InfNet* achieves the best performances on all six categories: *InfNet* is expected to be able to capture the spread of product interests along with both structural and temporal dynamics of *Taocode* diffusions, and brings in an averaged relative performance increase of 4.6% and 6.3% in terms of AUC_PR and AUC_ROC respectively compared with the best baseline method, demonstrating the effectiveness of our approach for purchasing predictions:

- In detail, although BPR only utilizes the information of user-item interactions while LR infers purchasing behaviors using only the user and item features, their performances are competitive with many baselines, which indicates that handcrafted features are crucial in item-level purchasing predictions.
- Most of the session-based methods focus primarily on the correlations between items across different time steps. However, due to the huge number of possible action sequences (*Taocode* sharing) within a session, this group of methods could be easily over-fitted on the training set, and they also ignore some key properties of users’ purchasing behavior. And when the scale of the dataset is relatively large (Cate#1, #2, #5 and #6), session-based models achieve only a limited improvement over LR and BPR.
- Social recommendation methods are the most appropriate to our experimental setting, and most similar to *InfNet*, as they not only consider the user-item interactions but also model the social dynamics among users. For instance, EATNN uses a whole data-based optimization strategy for neural models, GraphRec constructs both item-item and user-user graphs based on users’

Table 2: Experimental results of item-level purchasing prediction tasks. We use * to denote the best result among all models, while an underline indicates the best performance among the baseline methods.

Methods \ Datasets	Cate#1		Cate#2		Cate#3		Cate#4		Cate#5		Cate#6	
	ROC	PR										
LR	0.607	0.310	0.565	0.371	0.523	0.048	0.641	0.200	0.636	0.152	0.567	0.479
BPR	0.599	0.317	0.550	0.367	0.598	0.057	0.614	0.164	0.562	0.134	0.549	0.492
SR-GNN	0.600	0.318	0.569	0.382	0.599	0.059	0.603	0.187	0.558	0.126	0.558	0.495
MGNN-SPred	0.604	0.319	0.554	0.370	0.573	0.054	0.637	0.190	0.559	0.126	0.567	0.514
GCE-GNN	0.615	0.330	0.573	0.389	<u>0.613</u>	0.064	0.655	0.206	0.561	0.127	0.569	0.515
EATNN	0.508	0.246	0.512	0.340	0.508	0.042	0.507	0.120	0.512	0.107	0.505	0.450
GraphRec	0.547	0.277	0.562	0.375	0.542	0.048	0.623	0.171	0.573	0.132	0.515	0.474
DiffNet	<u>0.639</u>	<u>0.357</u>	<u>0.582</u>	<u>0.392</u>	0.612	<u>0.071</u>	<u>0.687</u>	<u>0.233</u>	<u>0.654</u>	<u>0.187</u>	<u>0.612</u>	<u>0.598</u>
<i>InfNet-S</i>	0.650	0.368	0.620	0.428	0.587	0.065	0.669	0.240	0.704	0.244	0.635	0.578
<i>InfNet</i>	0.671*	0.409*	0.644*	0.466*	0.643*	0.100*	0.726*	0.319*	0.742*	0.323*	0.657*	0.606*

Table 3: In-deep analysis on *InfNet*. “-” means we remove the corresponding part in *InfNet*. The reported metric is AUC_PR.

Analytical Category	Variant	Cate#1	Cate#2	Cate#3	Cate#4	Cate#5	Cate#6
(a) Sequence Encoder Selection	None	0.393	0.460	0.091	0.312	0.292	0.597
	Mean-Pooling	0.402	0.464	0.086	0.308	0.313	0.603
	GRU	0.403	0.466	0.094	0.316	0.322	0.606
(b) Effect of <i>Taocode</i> .	-User feature	0.407	0.464	0.086	0.297	0.305	0.599
	-Item feature	0.399	0.456	0.092	0.289	0.290	0.578
	-Taocode feature	0.355	0.406	0.079	0.282	0.236	0.577
(c) Attention mechanism.	-Attention	0.399	0.456	0.092	0.289	0.290	0.578
(d) Structural encoder	-Structural block	0.403	0.464	0.085	0.313	0.318	0.587
<i>InfNet</i>		0.409	0.466	0.100	0.319	0.323	0.605

social interactions, and DiffNet uses a fuse layer to combine users’ and items’ features, then apply GCN to model the interest diffusion in the social network. However, EATNN and GraphRec only consider the structural properties of *Taocode* diffusion networks while ignore the attributes on nodes and edges, which may lose too much information of *Taocode* messages. Furthermore, EATNN aims to rank all items for each user without negative sampling, bringing in performance drops on purchasing predictions of each single item. Although DiffNet achieves the best performance among all baselines, it does not consider the different influence of a user on her different neighbors, as well as the various impacts of *Taocode* records with different items. Therefore, those social recommendation methods cannot make accurate predictions on item-level purchases based on *Taocode* diffusion.

5.3 In-depth analysis of *InfNet*

InfNet consists of a structural encoder based on a node-aware GNN, and a dynamic *Taocode* diffusion encoder with attention-based edge aggregations and a sequence encoder. To investigate whether these encoders or attention mechanisms actually work and how they influence the performance of *InfNet*, we conduct ablation studies by removing each part of key component, and evaluate the performances of these variants of *InfNet* in Table 3:

- **Sequence Encoder Selection.** We adopt Masked Self-attention as the sequence encoder in Eq. (7), where elements in each time step are encoded only with reference to the previous time steps. Besides, there are many other deep architectures designed to encode sequences, such as RNN-based methods. In order to verify the influence of self-attention, we compare with the model replacing with GRU. In addition, to validate the necessity of sequence encoding, we further compare our methods with simple Mean-Pooling or even None (use only the output of last time step without any encoder). According to the results shown in Table 3(a), it is important to capture the temporal dynamics at different time steps by a sequence encoder, and simply using mean-pooling is unable to capture long- and short-term dependence, thus it cannot achieve a good performance on some time-sensitive categories (Cate#3 and #4). We choose Self-Attention because it is better at capturing long-term dependencies and achieves superior performance than GRU.
- **Effect of *Taocode*.** In order to intuitively demonstrate the importance of *Taocode*, we remove different part of input features. As shown in Table 3(b), masking *Taocode* features results in the sharpest drop in performance. This indicates that *Taocode* is a key factor to predict users’ purchasing behavior, so *InfNet* can model *Taocode* diffusion better for purchase prediction.

- **Attention mechanism.** In Eq. (2) and Eq. (3), *InfNet* assigns dynamic weights to each edge via graph attentions. In order to clarify the importance of the attention mechanism, we draw a comparison between *InfNet* and its variant without attention. As Table 3(c) shows, the performances of *InfNet*-Attention are worse on all six categories. with a maximum drop of 10.2% on Cate #3. Therefore, we conclude that the attention mechanism is essential for modeling *Taocode* interest diffusion.
- **Structural encoder.** *Structural block* is designed to encode user’s structural information. To evaluate whether this module can help us infer users’ purchase behavior, we compare with *InfNet* without this component. Again, as in Table 3(d), removing *Structural block* leads to a 2.3% decline in performance.

To summarize, according to the above discussions: (1) The sequence encoder is a necessary component, and self-attention is the best choice for *InfNet*. (2) *Taocode* is the most informative among the three input features of *InfNet*. (3) Attention mechanism can help *InfNet* better model *Taocode* interest diffusion. (4) Structural information also benefits in improving the proposed model.

Table 4: Comparison on cold-start problem. C# is short for Cate# and using bold to highlight the better improvement.

Group	Method	C#1	C#2	C#3	C#4	C#5
<i>Cold.</i>	DiffNet	0.331	0.332	0.061	0.159	0.143
	Ours	0.392	0.381	0.113	0.230	0.276
	Impv.	18.4%	14.8%	85.2%	44.7%	93.0%
<i>Warm.</i>	DiffNet	0.358	0.396	0.074	0.244	0.188
	Ours	0.410	0.472	0.086	0.329	0.329
	Impv.	14.5%	19.2%	16.2%	34.8%	75.0%

5.4 Cold-start challenge

As we observed in Sec. 5.3, *Taocode* features are important for purchase predictions. To further verify the power of *InfNet*, we conduct an analysis on the challenge which traditional recommendation system suffered, namely the **cold-start problem** [25]. There are two distinct categories of cold start: product cold-start and user cold-start, while here we focus on user cold-start problem.

In our setting, we denote *Cold* for the group of users who have no purchase records in the past and *Warm* on the contrary. We then compare *InfNet* with DiffNet (the best baseline), shown in Table 4. Although both models suffer in cold-start dilemma (perform worse under the *Cold* setting compared with *Warm*), *InfNet* shows a better improvement than DiffNet with cold-start. More specifically, except for Cate#2, *InfNet*’s improvements over DiffNet on *Cold* are significantly higher than that of *Warm*. It shows that *InfNet* is a potentially better approach to tackle the cold-start problem.

6 RELATED WORK

Purchase prediction in the recommendation system. Online shopping has been investigated since the early stages of the Web. Purchase behavior prediction in e-commerce is a classic but difficult problem. Many previous works studies online shopping behaviors based on features of users and products [17, 20, 26], such as item

price and user credit records, etc. Besides, traditional collaborative filtering [24] and deep learning techniques are used to predict the probabilities of which item will be purchased by the user. Some of them try to extract the implicit correlations between users and items [10, 11, 18, 34]. Another kind of methods focus on the click and purchase records of users, which treat the click/purchase records as a sequence and encoding them with a sequence encoder such as RNNs [12], and GNNs [33, 35, 38]. The goal of these methods is digging out the regular pattern of item purchases.

Social recommendations. Due to the potential value of social relations and user interactions in the recommendation systems, social recommendation has attracted increasing attention [29]. As theories from the domain of social science show the homophily and influence [1] between social relations, various approaches are proposed to build social recommendation systems such as trust ensemble [21], trust SVD [9], and social regularization [22]. More recently, deep learning methods as well as transfer learning [2] and GNNs [7, 36, 37, 39] are also adopted to solve the problem.

Deep learning on graphs. Recently, GNNs have achieved state-of-the-art performance on many tasks [16, 31, 32, 40]. Based on the Message-Passing framework, many works have expanded the capabilities to tackle different types of graph, such as heterogeneous graphs [14] and dynamic graphs [13, 45]. These improvements make GNNs better compatible for purchase prediction in real-world applications [6, 27].

7 CONCLUSION

This paper investigates the influence of information diffusion on user purchasing behaviors in an online shopping platform. Taking *Taocode* as a case study, we collect a large-scale real-world dataset that includes over 100M *Taocode* sharing records. Based on the dataset, we conduct empirical observations to explore the intrinsic properties of *product interest diffusion* on *Taobao*, finding that *Taocode* diffusion exerts a strong influence on user purchasing behaviors, while both structural and temporal dynamics of diffusion networks play a key role in such correlations. Inspired by these observational insights, we design an end-to-end GNN-based framework, dubbed *InfNet*, to model the product interest diffusion via *Taocode*. More specifically, we apply both graph- and sequence-level attention mechanisms to capture the dynamics of user interests on different products at different time steps. On the item-level purchasing prediction task on the real-world *Taocode* diffusion datasets from six different product categories, *InfNet* achieves significantly better performance compared with state-of-the-art baselines; moreover, ablation studies demonstrate the additional predictive power introduced by the careful design of our structural and temporal components. We hope this work will bring insights for user purchase behavior modeling, especially from the perspective of user interest diffusion and social interactions.

Acknowledgments. This work is supported by the National Key Research and Development Project of China (No. 2018AAA0101900) and a research funding from Alibaba Group. Yang Yang’s work is supported by Tongdun Technology.

REFERENCES

- [1] Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. The role of social networks in information diffusion. *WWW* (2012), 519–528.

- [2] Chong Chen, Min Zhang, Chenyang Wang, Weizhi Ma, Minming Li, Yiqun Liu, and Shaoping Ma. 2019. An efficient adaptive transfer neural network for social-aware recommendation. *SIGIR* (2019), 225–234.
- [3] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014).
- [4] Jesse Davis and Mark Goadrich. 2006. The relationship between Precision-Recall and ROC curves. *ICML* (2006), 233–240.
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [6] Shaohua Fan, Junxiong Zhu, Xiaotian Han, Chuan Shi, Linmei Hu, Biyu Ma, and Yongliang Li. 2019. Metapath-guided heterogeneous graph neural network for intent recommendation. *SIGKDD* (2019), 2478–2486.
- [7] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. *WWW* (2019), 417–426.
- [8] Mihajlo Grbovic, Vladan Radosavljevic, Nemanja Djuric, Narayan Bhamidipati, Jaikrit Savla, Varun Bhagwan, and Doug Sharp. 2015. E-commerce in your inbox: product recommendations at scale. *SIGKDD* (2015), 1809–1818.
- [9] Guibing Guo, Jie Zhang, and Neil Yorke-Smith. 2015. Trustsvd: collaborative filtering with both the explicit and implicit influence of user trust and of item ratings. *AAAI* (2015), 123–129.
- [10] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: simplifying and powering graph convolution network for recommendation. *SIGIR* (2020), 639–648.
- [11] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. *WWW* (2017), 173–182.
- [12] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [13] Wenjie Hu, Yang Yang, Ziqiang Cheng, Carl Yang, and Xiang Ren. 2021. Time-Series Event Prediction with Evolutionary State Graph. *WSDM* (2021), 580–588.
- [14] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous graph transformer. *WWW* (2020), 2704–2710.
- [15] Diederik P Kingma and Jimmy Ba. 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [16] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [17] Farshad Kooti, Kristina Lerman, Luca Maria Aiello, Mihajlo Grbovic, Nemanja Djuric, and Vladan Radosavljevic. 2016. Portrait of an online shopper: understanding and predicting consumer behavior. *WSDM* (2016), 205–214.
- [18] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. *SIGKDD* (2008), 426–434.
- [19] Seung Hwan (Mark) Lee, June Cotte, and Theodore J. Noseworthy. 2010. The role of network centrality in the flow of consumer influence. *JCP* 20, 1 (2010), 66–77.
- [20] Jure Leskovec, Lada A Adamic, and Bernardo A Huberman. 2006. The dynamics of viral marketing. *EC* (2006), 228–237.
- [21] Hao Ma, Irwin King, and Michael R Lyu. 2009. Learning to recommend with social trust ensemble. *SIGIR* (2009), 203–210.
- [22] Hao Ma, Dengyong Zhou, Chao Liu, Michael R Lyu, and Irwin King. 2011. Recommender systems with social regularization. *WWW* (2011), 287–296.
- [23] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [24] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. *WWW* (2001), 285–295.
- [25] Andrew I Schein, Alexandrin Popescul, Lyle H Ungar, and David M Pennock. 2002. Methods and metrics for cold-start recommendations. *SIGIR* (2002), 253–260.
- [26] Catarina Sismeiro and Randolph E Bucklin. 2004. Modeling purchase behavior at an e-commerce web site: a task-completion approach. *JMR* 41, 3 (2004), 306–323.
- [27] Weiping Song, Zhiping Xiao, Yifan Wang, Laurent Charlin, Ming Zhang, and Jian Tang. 2019. Session-based social recommendation via dynamic graph attention networks. *WSDM* (2019), 555–563.
- [28] Xiaodan Song, Belle L Tseng, Ching-Yung Lin, and Ming-Ting Sun. 2006. Personalized recommendation driven by information flow. *SIGIR* (2006), 509–516.
- [29] Jiliang Tang, Xia Hu, and Huan Liu. 2013. Social recommendation: a review. *SNAM* 3, 4 (2013), 1113–1133.
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762* (2017).
- [31] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).
- [32] Chunyang Wang, Yanmin Zhu, Tianzi Zang, Haobing Liu, and Jiadi Yu. 2021. Modeling inter-station relationships with attentive temporal graph convolutional network for air quality prediction. *WSDM* (2021), 616–634.
- [33] Wen Wang, Wei Zhang, Shukai Liu, Qi Liu, Bo Zhang, Leyu Lin, and Hongyuan Zha. 2020. Beyond clicks: modeling multi-relational item graph for session-based target behavior prediction. *WWW* (2020), 3056–3062.
- [34] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. *SIGIR* (2019), 165–174.
- [35] Ziyang Wang, Wei Wei, Gao Cong, Xiao-Li Li, Xian-Ling Mao, and Minghui Qiu. 2020. Global context enhanced graph neural networks for session-based recommendation. *SIGIR* (2020), 169–178.
- [36] Le Wu, Junwei Li, Peijie Sun, Richang Hong, Yong Ge, and Meng Wang. 2020. DiffNet++: a neural influence and interest diffusion network for social recommendation. *TKDE* (2020), 1–1.
- [37] Le Wu, Peijie Sun, Yanjie Fu, Richang Hong, Xiting Wang, and Meng Wang. 2019. A neural influence diffusion model for social recommendation. *SIGIR* (2019), 235–244.
- [38] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based recommendation with graph neural networks. *AAAI* (2019), 346–353.
- [39] Fengli Xu, Jianxun Lian, Zhenyu Han, Yong Li, Yujian Xu, and Xing Xie. 2019. Relation-aware graph convolutional networks for agent-initiated social e-commerce recommendation. *CIKM* (2019), 529–538.
- [40] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018).
- [41] Yang Yang, Jia Jia, Boya Wu, and Jie Tang. 2016. Social role-aware emotion contagion in image social networks. *AAAI* (2016), 65–71.
- [42] Yang Yang, Jie Tang, Cane Leung, Yizhou Sun, Qicong Chen, Juanzi Li, and Qiang Yang. 2015. Rain: Social role-aware information diffusion. *AAAI* (2015), 367–373.
- [43] Mao Ye, Xingjie Liu, and Wang-Chien Lee. 2012. Exploring social influence for recommendation: a generative model approach. *SIGIR* (2012), 671–680.
- [44] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. *AAAI* (2019), 5941–5948.
- [45] Lekui Zhou, Yang Yang, Xiang Ren, Fei Wu, and Yueting Zhuang. 2018. Dynamic network embedding by modeling triadic closure process. *AAAI* (2018), 571–578.