

Temporal Interest Network for User Response Prediction

Haolin Zhou*
Shanghai Jiao Tong University
koziello@sjtu.edu.cn

Xihua Chen
Tencent
tinychen@tencent.com

Junwei Pan*
Tencent
jonaspan@tencent.com

Jie Jiang
Tencent
zeus@tencent.com

Xinyi Zhou
Shanghai Jiao Tong University
zhouxy1003@sjtu.edu.cn

Xiaofeng Gao(✉)
Shanghai Jiao Tong University
gao-xf@cs.sjtu.edu.cn

Guihai Chen
Shanghai Jiao Tong University
gchen@cs.sjtu.edu.cn

ABSTRACT

User behaviors are among the most critical features for user response prediction in recommendation systems. Many works have revealed that a user’s behavior reflects her interest in the candidate item, owing to the semantic or temporal correlation between them. While the literature has individually examined each of these correlations, researchers have yet to analyze them in combination, that is, the semantic-temporal correlation. We empirically measure this correlation and observe intuitive yet robust patterns. We then examine several popular user interest models and find that, surprisingly, none of them learn such correlation well.

To fill this gap, we propose a Temporal Interest Network (TIN) to capture the semantic-temporal correlation simultaneously between behaviors and the target. We achieve this by incorporating target-aware temporal encoding, in addition to semantic encoding, to represent behaviors and the target. Furthermore, we conduct explicit 4-way interaction by deploying target-aware attention and target-aware representation to capture both semantic and temporal correlation. We conduct comprehensive evaluations on two popular public datasets, and our proposed TIN outperforms the best-performing baselines by 0.43% and 0.29% on GAUC, respectively. During online A/B testing in Tencent’s advertising platform, TIN achieves 1.65% cost lift and 1.93% GMV lift over the base model. It has been successfully deployed in production since October 2023, serving the WeChat Moments traffic. We have released our code at <https://github.com/zhouxy1003/TIN>.

CCS CONCEPTS

• **Computing methodologies** → **Semantic networks**.

*Equal Contribution

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 the author(s) 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.

WWW ’24 Companion, May 13–17, 2024, Singapore, Singapore

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0172-6/24/05...\$15.00

<https://doi.org/10.1145/3589335.3648340>

KEYWORDS

User Response Prediction, Target Attention, Sequential Recommendation, CTR Prediction

ACM Reference Format:

Haolin Zhou, Junwei Pan, Xinyi Zhou, Xihua Chen, Jie Jiang, Xiaofeng Gao(✉), and Guihai Chen. 2024. Temporal Interest Network for User Response Prediction. In *Companion Proceedings of the ACM Web Conference 2024 (WWW ’24 Companion)*, May 13–17, 2024, Singapore, Singapore. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3589335.3648340>

1 INTRODUCTION

In recent decades, users have been confronted with an overwhelming amount of information while browsing websites or mobile apps. This poses significant challenges for electronic retailers, content providers, and online advertising platforms in their quest to recommend suitable items to individual users within specific contexts. Consequently, recommendation systems have gained widespread deployment to capture users’ interests and predict their preferences from an extensive pool of candidate items. For instance, in cost-per-click (CPC) advertising, advertising platforms must bid for each ad based on the estimated value of the impression, which relies on the bid value and the estimated Click-Through Rate (CTR). As a result, the accurate prediction of user response becomes a critical factor and has garnered considerable research attention.

Recommendation systems usually use categorical ID features to represent users, items, and the context, such as user demographic features, item attribute features, *etc.* Among them, users’ history behaviors are crucial since they represent a user’s preference over items in the past and, therefore, may reflect her interest in the candidate item. Many works tried to capture such user interest by recurrent neural networks [7, 44], self-attention [3, 10, 29], or target attention [3, 5, 44, 45].

Intuitively, a user’s interest is reflected by what items she has interacted with and when these interactions happened. For example, if she has clicked ads on video games many times, then it’s a good idea to recommend ads for new video games to her in the future. However, if all her clicks on video games happened 1 year ago, but she clicked on clothing ads several times recently, then it’s better to recommend clothing ads. Therefore, when capturing a user’s interest in a candidate item, it’s crucial to consider both the semantic correlation and the temporal correlation between behaviors and

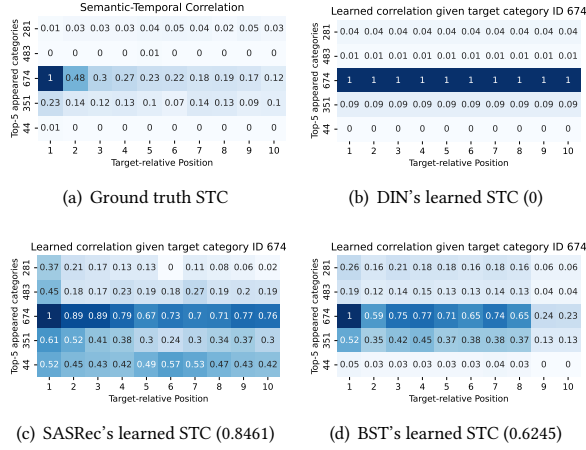


Figure 1: (a): The ground truth Semantic-Temporal Correlation (STC) between history behaviors of Top-5 categories (y-axis) at various target-relative positions from 1 to 10 (x-axis), and the target category of 674; (b,c,d): the learned STC in DIN, SASRec and BST. Numbers in parentheses denote the Pearson correlation coefficient between ground truth and learned correlation on behaviors with the same category as the target (3rd row).

the candidate. To our surprise, there is neither a study measuring such semantic-temporal correlation nor an examination of how well existing methods capture it.

Motivated by the lack of such analysis, we propose an empirical measurement of the ground truth semantic-temporal correlation by mutual information. On the Amazon dataset, which contains users' reviews on items, we calculate the *mutual information* between history reviews with Top-5 categories at various target-relative positions and the target item with a specific category. As illustrated in Fig. 1(a), our results reveal a compelling semantic-temporal pattern. We then examined the learned correlation of several popular user interest models such as DIN [45], SASRec [10] and BST [3], and show the results in Fig. 1. To our surprise, none of these models is capable of capturing the semantic-temporal pattern well.

To capture such correlation, the model needs to conduct a 4-way interaction over the quadruplet (behavior semantics, target semantics, behavior temporal, target temporal). Nevertheless, existing user interest methods are not equipped with such quadruple interaction, making them unable to learn semantic-temporal correlation. For instance, DIN [45] doesn't consider any temporal information of behaviors. SASRec [10] and Bert4Rec [29] leverage positional embedding to model behavior sequences, but their position is target-agnostic, leading to inadequate modeling of target-aware temporal correlations. BST [3] conducts self-attention over behaviors and the target, enabling only 3-order explicit interactions.

To this end, we propose a novel model named Temporal Interest Network (TIN), to capture the semantic-temporal correlation. To achieve this, we adopt Target-aware Temporal Encoding (TTE) to encode the temporal information, which applies to both behaviors and the target itself. Specifically, TTE employs the relative position or time interval of each behavior with respect to the target. Furthermore, we incorporate Target-aware Attention (TA) and

Target-aware Representation (TR), each interacting behaviors with the target. Finally, we explicitly incorporate a 4-way interaction by multiplying the output of target-aware attention with that of target-aware representation to capture the quadruple semantic-temporal correlation.

Each of the three components in TIN is critical in acquiring the semantic-temporal correlation. To empirically validate the significance of each component, we conduct an ablation study on the Amazon dataset. The ablation of each component results in a notable decrease in the GAUC by 0.99%, 0.73%, and 0.61%, respectively. It is noteworthy that all existing methods can be considered as ablated variants of TIN, as they lack one or more of these critical components. The contribution of this paper can be summarized as follows:

- We pioneer a study on quantifying the semantic-temporal correlation between behaviors and the target, and reveal the presence of strong semantic-temporal patterns. We examine several popular user interest models and observe that they fail to capture such correlation.
- We propose a simple yet effective user interest model named Temporal Interest Network (TIN) to capture semantic-temporal correlation. TIN adopts target-aware temporal encoding and incorporates a 4-way interaction by deploying both target-aware attention and representation.
- We conduct comprehensive experiments on two publicly available datasets, and the results demonstrate that TIN outperforms state-of-the-art user interest methods.

2 SEMANTIC-TEMPORAL CORRELATION

This section presents a novel measurement to quantize the ground truth semantic-temporal correlation between behaviors and the target via *mutual information*. We then examine several popular user interest models on their ability to learn such correlation.

2.1 Measurement of Ground Truth Semantic-Temporal Correlation

We investigate the semantic-temporal correlation on the public Amazon dataset [44, 45], which consists of user's reviews on items. Regarding temporal correlation, we choose the relative position of each behavior regarding the target without loss of generality. We choose the category feature of behaviors and the target since its cardinality is intermediate. We quantify the ground truth semantic-temporal correlation by the following metric.

Definition 1 (Category-wise Target-aware Correlation (CTC)). CTC is defined as the mutual information [18, 26] between behaviors with category c_i whilst occurring at position p : $\mathcal{X}_{C(X)=c_i \wedge P(X)=p}$ and the user response label¹ on the target item with category c_t : $\mathcal{Y}_{C(Y)=c_t}$. Formally,

$$\text{Cor} = \text{MI} \left(\mathcal{X}_{C(X)=c_i \wedge P(X)=p}, \mathcal{Y}_{C(Y)=c_t} \right), \quad (1)$$

where $C(\cdot)$, $P(\cdot)$ denotes the category or position of the behavior or target. For example, the CTC of behaviors with Top-5 categories at various target-relative positions and target with category $c_t = 674$

¹The Amazon dataset only contains reviews, i.e., positive responses, on the target item. We generate negative responses by replacing the target item with a random item.

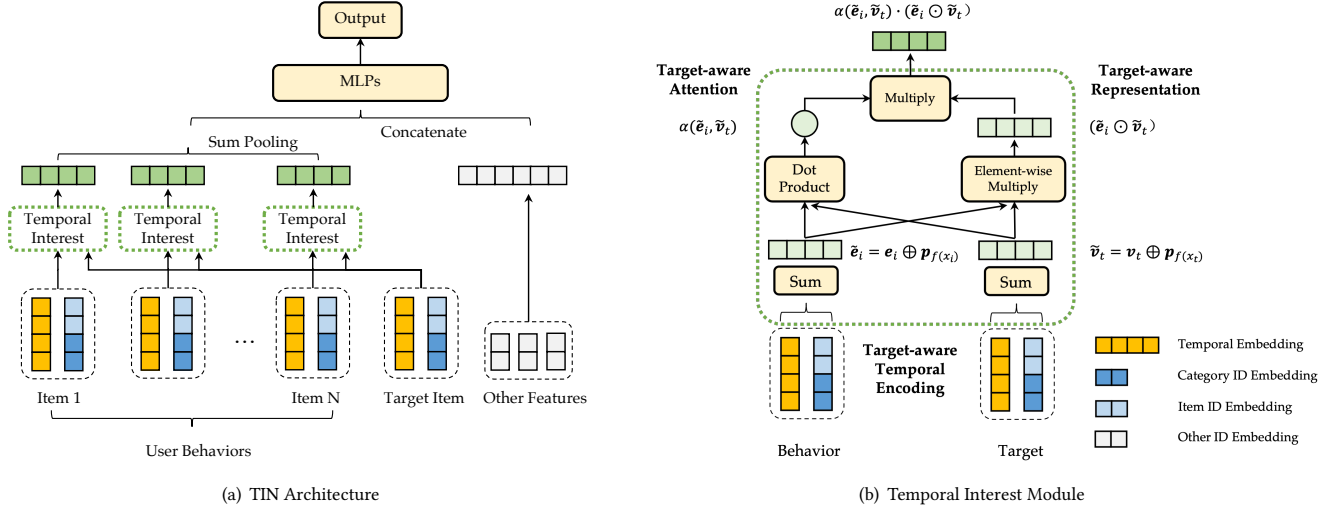


Figure 2: The left part illustrates the whole architecture of TIN; the right part illustrates the architecture of the Temporal Interest Module, which introduces target-aware temporal encoding to preserve the temporal information, and target-aware attention and representation for explicit 4-way interaction.

is illustrated in Fig. 1(a). We observe 1) *semantic pattern between matching categories*: behaviors belonging to the same category as the target (the 3rd row, category 674) exhibit a higher degree of correlation compared to other categories. 2) *temporal decaying pattern*: among the semantically correlated behaviors (i.e., the 3rd row), there is a compelling correlation decrease from the most recent behaviors to the oldest ones.

2.2 How do Existing Methods Capture Semantic-Temporal Pattern?

We are intrigued by how existing user interest methods capture the observed patterns. To investigate this, we select three popular user interest models: DIN [45], SASRec [10], and BST [3], and evaluate the learned correlation of these models on the target category 674, as depicted in Fig. 1(b), Fig. 1(c) and Fig. 1(d).

We observe notable disparities in the learned correlations of these models compared to the ground truth one. DIN fails to capture any temporal pattern, as evidenced by the absence of decay along positions. SASRec and BST demonstrate limited proficiency in learning temporal correlation, as there is a slight decay along positions. However, they fall short in capturing the semantic correlation, as they learn high correlations for many other categories besides target category 674 (3rd row). Existing methods' failure to learn semantic-temporal correlation motivates us to develop a more effective user interest model. For a more comprehensive understanding of how we measure the learned correlation of each model, please refer to Sec. 4.3.

3 TEMPORAL INTEREST NETWORK

Our model adopts the widely used Embedding & MLP paradigm. To begin with, all features from the user side, item side, and context side are transformed into embeddings. We then process user behavior

features by a Temporal Interest Module (TIM), generating a fixed-length user interest representation. Subsequently, the output of the TIM and other one-hot encoding features are concatenated and fed into Multi-Layer Perceptrons (MLPs). The final output is obtained through a sigmoid function, and the entire model is optimized using the cross-entropy loss. Refer to Fig. 2 for an overview of the model architecture.

3.1 Temporal Interest Module

In order to conduct 4-way interaction over the quadruplet (behavior semantics, target semantics, behavior temporal, target temporal) to capture the semantic-temporal correlation, we incorporate the following components. a) Target-aware Temporal Encoding (TTE): TTE preserves the temporal information of behaviors regarding the target. Combining TTE with the semantic ID embedding, we can capture both semantic and temporal information about behaviors and the target. b) Target-aware Attention (TA) and c) Target-aware Representation (TR) over the temporally encoded behaviors and target: each conducts a 2-way behavior-target interaction. We multiply the output of TA and TR, resulting in an explicit 4-way interaction between behaviors and the target, to capture their semantic and temporal correlation. Formally,

$$\mathbf{u}_{\text{TIN}} = \sum_{i \in \mathcal{H}} \alpha(\tilde{\mathbf{e}}_i, \tilde{\mathbf{v}}_t) \cdot (\tilde{\mathbf{e}}_i \odot \tilde{\mathbf{v}}_t) \quad (2)$$

where $\tilde{\mathbf{e}}_i = \mathbf{e}_i \oplus \mathbf{p}_{f(x_i)}$ denotes the representation of the i -th behavior X_i , consisting of both semantic encoding \mathbf{e}_i and *target-aware temporal encoding* $\mathbf{p}_{f(i)}$, similarly, $\tilde{\mathbf{v}}_t = \mathbf{v}_t \oplus \mathbf{p}_{f(x_t)}$ denotes the representation of target item. $\alpha(\tilde{\mathbf{e}}_i, \tilde{\mathbf{v}}_t)$ and $(\tilde{\mathbf{e}}_i \odot \tilde{\mathbf{v}}_t)$ denotes the target-aware attention and target-aware representation between the i -th behavior and the target, $\alpha(\cdot)$ denotes the softmax function, and \odot denotes element-wise multiplication. We illustrated the architecture of TIM in Fig. 2(b).

3.1.1 Target-aware Temporal Encoding (TTE). In addition to encoding the ID features of each behavior as embeddings $\mathbf{e}_i \in \mathcal{R}^d$ to capture semantic meaning, we also incorporate the temporal information of each behavior, focusing on the *target-aware* temporal aspects. To achieve this, we employ two distinct methods: TTE based on relative Position and TTE based on Time interval.

Target-aware Temporal Encoding based-on Position (TTE-P). Denote the sequence that contains both history clicked items and the target whilst sorted by time as $\{X_1, \dots, X_H, X_t\}$. TTE-P encodes the temporal information of the i -th behavior as $H - i + 1$. The resulting embedding for this behavior is denoted as $\mathbf{p}_{f_{\text{TTE-P}}(X_i)} \in \mathcal{R}^d$. Please note that the TTE-P of the target item is assigned a value of 0, representing *the origin*, and its corresponding embedding is denoted as $\mathbf{p}_{f_{\text{TTE-P}}(X_t)} = \mathbf{p}_0$.

Target-aware Temporal Encoding based-on Time interval (TTE-T). The TTE-T method encodes the temporal information of the i -th behavior by its time interval from the target. Let us denote the timestamp of the i -th behavior and the target as TS_i and TS_t . The time interval between them can be represented as $\tau_i = TS_t - TS_i$. To facilitate the learning process, we discretize the time interval by a binning function $b(\cdot)$, such as equal frequency binning. Consequently, the TTE-T embedding for the i -th behavior is defined as $\mathbf{p}_{f_{\text{TTE-T}}(X_i)} = \mathbf{p}_{b(\tau_i)}$. Similar to TTE-P, the TTE-T of the target item also represents the origin.

Once we obtain the temporal encoding $\mathbf{p}_{f(X_i)}$ for a behavior X_i , its final embedding is obtained through element-wise summation with its semantic embedding: $\tilde{\mathbf{e}}_i = \mathbf{e}_i \oplus \mathbf{p}_{f(X_i)}$. Similarly, the target item X_t is represented as $\tilde{\mathbf{v}}_t = \mathbf{v}_t \oplus \mathbf{p}_{f(X_t)}$. By combining the semantic and temporal embeddings, we get enriched representations that incorporate both the inherent semantic meaning and the temporal context of the behaviors and the target item.

TTE v.s. Chronological Order Encoding. Another widely used temporal encoding is Chronological Order Encoding (COE) [3, 5, 10, 29], which encodes behaviors based on their chronological order from earliest to oldest, *i.e.*, $f_{\text{COE}}(X_i) = i$. While COE is widely adopted, it has limitations as it is target agnostic. It may assign the same weights to behaviors regardless of their relative positions to the target. For instance, consider two behavior sequences with lengths of 10 and 100, respectively. According to COE, the first behaviors of both sequences have the same position, *i.e.*, $f_{\text{COE}}(X_1) = 1$. However, their relative positions to the target item differ significantly, *i.e.*, 10 versus 100, indicating varying levels of importance. Unfortunately, COE assigns the same position and encoding embedding to both behaviors, disregarding their distinct importance.

There are also other temporal encoding schemes such as Timestamp Encoding (TE) [37], which applies an embedding layer to encode a timestamp t containing hour h , weekday w , month m and etc., into several time embeddings of different granularities, *i.e.*, $f_{\text{TE}}(X_t) = \text{Emb}(h, w, m, \dots)$. While TE can be considered target-aware as long as the target item is incorporated in the attention or representation component, our experiments reveal that it performs worse than TTE. We will present a performance comparison of temporal encoding methods in Section 4.4.

3.1.2 Target-aware Attention and Target-aware Representation. Building upon target attention methods [5, 44, 45], we aim to capture the importance of each behavior X_i through an attention function. In this regard, we adopt the Scaled Dot-Product [32, 43] as our activation function: $\alpha(\tilde{\mathbf{e}}_i, \tilde{\mathbf{v}}_t) = \sigma(\frac{\langle \tilde{\mathbf{e}}_i, \tilde{\mathbf{v}}_t \rangle}{\sqrt{d}})$, where $\sigma(z)$ denotes the softmax function over all behaviors of a given user.

Existing methods typically multiply the attention weight with the representation of each behavior and then concatenate it with the embeddings of the target and other features to be fed into Multi-Layer Perceptrons (MLPs): $g_{\text{MLP}}([\sum_i \sigma(\mathbf{e}_i, \mathbf{v}_t) \mathbf{e}_i, \mathbf{v}_t])$. However, it has been proved that MLPs struggle to effectively learn dot product or explicit interactions [25, 42], limiting these methods to conducting only 3-way explicit interactions through the self-attention, *i.e.*, $\sigma(\mathbf{e}_i, \mathbf{v}_t) \cdot \mathbf{e}_i$.

In TIN, each behavior is represented as a second-order representation with respect to the target item, resulting in a target-aware representation: $\tilde{\mathbf{e}}_i \odot \tilde{\mathbf{v}}_t$. Such explicit second-order representation or interaction is widely adopted in existing explicit high-order interaction click-through rate (CTR) models [9, 13, 18, 24, 30, 33]. By multiplying the target-aware attention (TA) with the target-aware representation (TR), TIN explicitly conducts 4-way interactions over the quadruplet (behavior semantics, target semantics, behavior temporal, target temporal). Note that the absence of either TA or TR would cause performance deterioration, which will be discussed in Sec. 4.2.

3.2 Connection to Existing Methods

The Temporal Information (TI), Target-aware Attention (TA), and Target-aware Representation (TR) have been widely used in existing user interest models for user response prediction [3, 5, 7, 10, 29, 43–45]. In order to analyze the presence of these components, we assign a 3-bit code to represent each model, where each bit indicates the existence (✓) or absence (✗) of TI, TA, and TR. It is important to note that TI refers to any form of temporal information, including Temporal/Position Encoding or Recurrent Neural Networks (RNN), and is not limited to Target-aware Temporal Encoding (TTE).

For instance, the code for TIN is ✓✓✓, indicating the presence of all three components. On the other hand, DIN's code is ✗✓✓, as it lacks TI. Models such as GRU4Rec, SASRec, and BERT4Rec share the same code, ✓✗✓, as they incorporate TI and TR but lack TA. Conversely, other target attention methods like DIEN, DSIN, and BST have the code ✓✓✗, due to the failure of Concat & MLP to learn dot product [25]. Please refer to Table 5 for a comprehensive overview of each method's code.

3.3 Connections to Transformer

There are several commonalities between the Temporal Interest Module (TIM) and Scaled Dot-Product Attention proposed in Transformer [32]: 1. they both employ Scaled Dot-Product as the attention function; 2. they consider both semantic embedding as well as temporal embedding to represent each token or item. Using the Transformer terminology, $\tilde{\mathbf{v}}_t$ and $\tilde{\mathbf{e}}_i$ in target-aware attention corresponds to *query* (target query) and *key* (behavior key) in Dot-Product Attention, while $\tilde{\mathbf{e}}_i$ and $\tilde{\mathbf{v}}_t$ in target-aware representation corresponds to *values* (behavior value and target value).

Table 1: Evaluation results on the Alibaba and Amazon dataset. Models with the same background color have the same component code.

Model	Code	Amazon				Alibaba			
		Logloss	$\Delta\%$	GAUC	$\Delta\%$	Logloss	$\Delta\%$	GAUC	$\Delta\%$
Avg Pooling & Concat	XXX	0.4908 (1E-3)	-	0.8445 (2E-3)	-	0.1969 (1E-3)	-	0.6074 (5E-4)	-
Avg Pooling & Product	XX✓	0.4824 (5E-4)	-1.71	0.8523 (8E-4)	+0.92	0.1960 (1E-3)	-0.46	0.6106 (9E-4)	+0.53
DIN'	X✓X	0.4803 (9E-4)	-2.14	0.8536 (4E-4)	+1.08	0.1962 (1E-3)	-0.36	0.6096 (6E-4)	+0.36
DIN	X✓✓	0.4703 (2E-3)	-4.18	0.8590 (1E-3)	+1.72	0.1963 (1E-3)	-0.30	0.6113 (4E-4)	+0.64
GRU4Rec	✓X✓	0.4766 (2E-3)	-2.89	0.8574 (2E-3)	+1.53	0.1972 (1E-3)	+0.15	0.6091 (4E-4)	+0.28
SASRec	✓X✓	0.4837 (7E-3)	-1.45	0.8497 (4E-3)	+0.62	0.1959 (8E-4)	-0.51	0.6091 (4E-4)	+0.28
BERT4Rec	✓X✓	0.4833 (5E-3)	-1.53	0.8501 (2E-3)	+0.66	0.1961 (1E-3)	-0.41	0.6096 (9E-4)	+0.36
DIEN	✓✓X	0.4807 (8E-3)	-2.06	0.8590 (1E-3)	+1.72	0.1973 (9E-4)	+0.20	0.6108 (6E-4)	+0.56
DSIN	✓✓X	0.4726 (2E-3)	-3.71	0.8592 (1E-3)	+1.74	0.1964 (2E-3)	-0.25	0.6106 (9E-4)	+0.53
BST	✓✓X	0.4850 (5E-4)	-1.18	0.8500 (9E-4)	+0.65	0.1959 (2E-3)	-0.51	0.6096 (6E-4)	+0.36
TIN	✓✓✓	0.4636 (3E-3)	-5.54	0.8629 (9E-4)	+2.18	0.1954 (2E-3)	-0.76	0.6144 (4E-4)	+1.15
TIN w/o TTE	X✓✓	0.4752 (2E-3)	-3.18	0.8544 (8E-4)	+1.17	0.1963 (1E-3)	-0.30	0.6135 (7E-4)	+1.00
TIN w/o TA	✓X✓	0.4758 (3E-3)	-3.06	0.8566 (1E-3)	+1.43	0.1960 (7E-4)	-0.46	0.6094 (8E-4)	+0.33
TIN w/o TR	✓✓X	0.4743 (2E-3)	-3.36	0.8576 (9E-4)	+1.55	0.1965 (2E-3)	-0.20	0.6127 (1E-3)	+0.87

They differ from each other regarding: 1. Scaled Dot-Product Attention is applied between each pair of tokens within the sequence, while TIM is only applied between the target and each item in the sequence, that is, Scaled Dot-Product Attention is a Self-Attention upon a sequence, while TIM is a Target-Attention upon a sequence and a target; 2. The representation in Scaled Dot-Product Attention is 1st-order only, that is, V , while it's 2nd-order in TIM, that is, $(\tilde{e}_i \odot \tilde{v}_t)$.

4 EXPERIMENTS

In this section, we aim to address the following research questions (RQ) through experiments on real-world datasets:

RQ1: How does the performance of our proposed Temporal Interest Network (TIN) compare to state-of-the-art methods?

RQ2: Is each component of TIN essential for capturing semantic-temporal correlations? What would be the impact if any of these components were removed?

RQ3: How can we effectively measure the learned semantic-temporal correlation of TIN and other methods? How does TIN capture such correlations?

RQ4: How do various temporal encoding methods perform in comparison to each other?

4.1 RQ1: Overall Performance Evaluation

We compare TIN with the following user interest methods as baselines: Avg Pooling & Concat [45], Avg Pooling & Product [45], DIN¹, DIN², GRU4Rec [7], SASRec, BERT4Rec [29], DIEN [44], DSIN [5] and BST. Refer to the appendix for the experimental setting.

¹DIN' denotes the architecture described in the paper, consisting of only target-aware attention

²DIN denotes the architecture implemented in the released code, which multiplies the attention-weighted pooling of behaviors with the target outside the attention, therefore consists of both target-aware attention and target-aware representation

The performance evaluation results are presented in Table 1. It is evident that Avg Pooling & Concat exhibits the poorest performance on both datasets. This can be attributed to its lack of all three essential components, rendering it incapable of capturing either temporal correlation or semantic correlation (feature interaction). In contrast, Avg Pooling & Product demonstrates a significant performance improvement compared to Sum Pooling & Concat, with gains of $9.2e-3$ and $5.3e-3$ on each dataset, respectively. This improvement can be attributed to the inclusion of target-aware representation through the product operation, as discussed in [25]. While DIN' performs worse than Avg Pooling & Product on three metrics, except for GAUC on the Alibaba dataset, this can be attributed to the limited expressiveness of Temporal Attention (TA) compared to TR. DIN outperforms DIN' due to its incorporation of target-aware representation.

All the remaining methods, namely GRU4Rec, SASRec, BERT4Rec, DIEN, DSIN, and BST, take into account temporal information. It is worth noting that many of these methods outperform DIN on the Amazon dataset due to the presence of a strong target-aware temporal correlation, as shown in Fig. 1(a). However, on the Alibaba dataset, some of these methods fail to surpass DIN. This discrepancy can potentially be attributed to the dense temporal nature of behaviors in the Alibaba dataset, as discussed in detail in Section 4.4. The high density of behaviors in relation to time renders the position information less influential or useful in this context.

TIN achieves a GAUC of 0.8629 on the Amazon dataset, surpassing the performance of the best-performing baseline (DSIN) by 0.43%. Similarly, on the Alibaba dataset, TIN achieves a GAUC of 0.6144, outperforming the best-performing baseline (DIN) by 0.51%. Both improvements are statistically significant, demonstrating the superior performance of TIN.

4.2 RQ2: Ablation Study

4.2.1 Disabling Target-aware Temporal Encoding. Disabling TTE in TIN renders it incapable of learning the temporal behavior-target

Table 2: Measurement of learned semantic-temporal correlation.

Model	Overall	Attention	Representation
TIN	$e^{\langle \tilde{\mathbf{e}}_i, \tilde{\mathbf{v}}_t \rangle} \cdot \ \tilde{\mathbf{e}}_i \odot \tilde{\mathbf{v}}_t\ _2$	$e^{\langle \tilde{\mathbf{e}}_i, \tilde{\mathbf{v}}_t \rangle}$	$\ \tilde{\mathbf{e}}_i \odot \tilde{\mathbf{v}}_t\ _2$
TIN w/o TTE	$e^{\langle \mathbf{e}_i, \mathbf{v}_t \rangle} \cdot \ \mathbf{e}_i \odot \mathbf{v}_t\ _2$	$e^{\langle \mathbf{e}_i, \mathbf{v}_t \rangle}$	$\ \mathbf{e}_i \odot \mathbf{v}_t\ _2$
TIN w/o TA	$\ (\mathbf{e}_i \oplus \mathbf{p}_{f_{\text{TTE-P}}(i)}) \odot (\mathbf{v}_t \oplus \mathbf{p}_{f_{\text{TTE-P}}(t)})\ _2$	–	$\ \tilde{\mathbf{e}}_i \odot \tilde{\mathbf{v}}_t\ _2$
TIN w/o TR	$e^{\langle \mathbf{e}_i \oplus \mathbf{p}_{f_{\text{TTE-P}}(i)}, \mathbf{v}_t \oplus \mathbf{p}_{f_{\text{TTE-P}}(t)} \rangle} \cdot \ \mathbf{e}_i \oplus \mathbf{p}_{f_{\text{TTE-P}}(i)}\ _2$	$e^{\langle \tilde{\mathbf{e}}_i, \tilde{\mathbf{v}}_t \rangle}$	$\ \mathbf{e}_i \oplus \mathbf{p}_{f_{\text{TTE-P}}(i)}\ _2$
DIN	$e^{\langle \mathbf{e}_i, \mathbf{v}_t \rangle} \cdot \ \mathbf{e}_i \odot \mathbf{v}_t\ _2$	$e^{\langle \mathbf{e}_i, \mathbf{v}_t \rangle}$	$\ \mathbf{e}_i \odot \mathbf{v}_t\ _2$
SASRec	$\ (\mathbf{e}_i \oplus \mathbf{p}_{f_{\text{COE}}(i)}) \odot \mathbf{v}_t\ _2$	–	$\ (\mathbf{e}_i \oplus \mathbf{p}_{f_{\text{COE}}(i)}) \odot \mathbf{v}_t\ _2$
BST	$e^{\langle \mathbf{e}_i \oplus \mathbf{p}_{f_{\text{COE}}(i)}, \mathbf{v}_t \oplus \mathbf{p}_{f_{\text{COE}}(t)} \rangle} \cdot \ \mathbf{e}_i \oplus \mathbf{p}_{f_{\text{COE}}(i)}\ _2$	$e^{\langle \mathbf{e}_i \oplus \mathbf{p}_{f_{\text{COE}}(i)}, \mathbf{v}_t \oplus \mathbf{p}_{f_{\text{COE}}(t)} \rangle}$	$\ \mathbf{e}_i \oplus \mathbf{p}_{f_{\text{COE}}(i)}\ _2$

correlation, consequently failing to capture the semantic-temporal correlation. This deficiency results in a significant performance drop of 9.9e-3 on the Amazon dataset, where a strong temporal behavior-target correlation has been established, as discussed in Section 2. Similarly, on the Alibaba dataset, the performance experiences a drop of 1.5e-3. This relatively smaller performance drop can be attributed to the narrower time range of behaviors in the Alibaba dataset, which diminishes the strength of the temporal pattern compared to that observed in the Amazon dataset.

4.2.2 Disabling Target-aware Attention. In the absence of target-aware attention in TIN, the model regresses to 2-way interaction solely based on the target-aware representation. Furthermore, excluding the attention mechanism results in a target-aware representation that merely pools the semantic and temporal embeddings from all behaviors, thereby losing the temporal information associated with each behavior. To illustrate this, consider a user with only two behaviors, denoted as i and j . The resulting representation can be expressed as:

$$\begin{aligned} \tilde{\mathbf{e}}_i \odot \tilde{\mathbf{v}}_t &= (\mathbf{v}_i \oplus \mathbf{p}_{f(X_i)}) \odot (\mathbf{v}_t \oplus \mathbf{p}_{f(X_t)}) \\ &\quad + (\mathbf{v}_j \oplus \mathbf{p}_{f(j)}) \odot (\mathbf{v}_t \oplus \mathbf{p}_{f(X_t)}) \\ &= (\mathbf{v}_i \oplus \mathbf{p}_{f(X_i)} \oplus \mathbf{v}_j \oplus \mathbf{p}_{f(j)}) \odot (\mathbf{v}_t \oplus \mathbf{p}_{f(X_t)}). \end{aligned} \quad (3)$$

It is evident that this formulation disrupts the correspondence between each behavior and its respective position. Consequently, this leads to a performance drop of 7.3e-3 and 8.1e-3 on the Amazon and Alibaba datasets, respectively.

4.2.3 Disabling Target-aware Representation. Similar to target-aware attention, disabling target-aware representation makes the model degenerate to the 2nd or 3rd order, disabling it to capture semantic-temporal correlation well. However, with the presence of position encoding, TIN w/o TR is still able to capture the target-aware temporal correlation to some extent within the target-aware attention. Therefore, even though there is a 6.1e-3 and 2.7e-3 performance drop on each dataset, TIN w/o TR beats TIN w/o TA on both datasets and beats TIN w/o TTE on the Amazon dataset.

4.3 RQ3: Measurement of Learned Correlation

In this section, we outline how to measure the learned semantic-temporal correlation of various models, including TIN, its three ablated variants, and several representative user interest models.

The learned semantic-temporal correlation between behavior X_i at position $f(X_i)$ and target X_t is defined as:

$$\overline{\text{Cor}}(f(X_i), \mathbf{e}_i, \mathbf{v}_t) = e^z \cdot \|\mathbf{r}\|_2. \quad (4)$$

where z represents the attention logits and \mathbf{r} denotes the representation embedding. We chose this form since there are connections between the mutual information and the parameterized (intermediate) output of the interaction of two variables [17]. In our model, the intermediate parameterized output of the interaction before the MLPs is proportional to $e^z \cdot \mathbf{r}$. For instance, the learned semantic-temporal correlation of TIN can be defined as:

$$\begin{aligned} \overline{\text{Cor}}(f(X_i), \mathbf{e}_i, \mathbf{v}_t)_{\text{TIN}} &= \exp^{z_{\text{TIN}}} \cdot \|\mathbf{r}_{\text{TIN}}\|_2 \\ &= \exp^{\frac{\langle \tilde{\mathbf{e}}_i, \tilde{\mathbf{v}}_t \rangle}{\sqrt{d}}} \cdot \|\tilde{\mathbf{e}}_i \odot \tilde{\mathbf{v}}_t\|_2 \end{aligned} \quad (5)$$

For the three ablated variants of TIN, the corresponding ablated component is disabled in the measurement. In the case of DIN, it possesses both TA and TR but lacks temporal encoding. For SASRec, it lacks TA as it solely applies self-attention over behaviors. The learned semantic-temporal correlation is measured using the representation, which is simplified as the element-wise product between a COE-encoded embedding of the behavior and the target. As for BST, it incorporates TA with COE to capture temporal correlation. However, its representation is target-agnostic since there is no explicit interaction between the behavior and the target. For a summary of how to measure the learned correlation of all models, please refer to Table 2.

As depicted in Figure 3, the learned semantic-temporal correlation of TIN exhibits a remarkable proximity to the ground truth 3(a). However, when excluding the Target-aware Temporal Encoding (TTE) component, TIN w/o TTE is only capable of learning the semantic correlation between behavior and target categories (as observed in the darkened third row) while failing to capture the underlying temporal pattern. Similarly, when excluding the TA component, TIN w/o TA not only fails to capture the temporal pattern but also struggles to learn the semantic correlation. On the other hand, TIN w/o TR demonstrates decent learning of the temporal pattern, but it falls short in capturing the semantic correlation, as evidenced by the non-zero weights assigned to almost all behaviors from other categories.




Figure 3: Learned semantic-temporal correlation on the Amazon dataset. The numbers in parentheses represent the Pearson correlation coefficient between the ground truth (as shown in Fig. 1(a)) and the learned correlation on behaviors with the same category as the target (3rd row).

Table 3: Performance evaluation of temporal encoding methods on the Amazon dataset with various behavior sequence length ranges.

Range of sequence lengths	[1, 5)	[5, 10)	[10, 15)	[15, 20)	[20, 430]
TTE-P	0.8571	0.8630	0.8742	0.8752	0.8728
COE	0.8564	0.8615	0.8718	0.8731	0.8675
$\Delta\%$ (GAUC)	-0.0009	-0.0017	-0.0027	-0.0025	-0.0061

Table 4: Performance evaluation of various encoding methods.

	Amazon				Alibaba			
	Logloss	$\Delta\%$	GAUC	$\Delta\%$	Logloss	$\Delta\%$	GAUC	$\Delta\%$
TTE-P	0.4636	-	0.8629	-	0.1954	-	0.6144	-
TTE-T	0.4484	-3.28	0.8703	+0.86	0.1956	+0.10	0.6139	-0.08
Timestamp Encoding, PH	0.4647	+0.24	0.8621	-0.09	0.1960	+0.31	0.6132	-0.20

4.4 RQ4: Temporal Encoding Methods

In TIN, we employ TTE-P to encode the target-relative position of historical behaviors. This approach differs from the commonly used self-attention position encoding method, Chronological Order Encoding (COE), which encodes behaviors in chronological order from earliest to oldest. TTE-P distinguishes itself from COE in that behaviors with the same TTE-P may receive different COE values due to variations in sequence length. Consequently, as the sequence length increases, the disparity between TTE-P and COE becomes more pronounced.

To evaluate these two temporal encoding methods, we partitioned the Amazon dataset based on the sequence length and conducted a comparative analysis. As shown in Table 3, TTE-P consistently outperforms COE across all sequence length ranges, with the performance gap widening for longer sequences. Unfortunately, we are unable to perform a similar comparison on the Alibaba dataset since most sequences are padded to a fixed length of 50.

Next, we present the comparison results of TTE-P, TTE-T, and Timestamp Encoding on both datasets in Table 4. Specifically, TTE-T

split the time interval into 10 buckets with equal frequency binning function. Upon replacing TTE-P with TTE-T, we observe an increase in the GAUC by $8.6e-3$ on the Amazon dataset and a slight drop of $0.8e-3$ on the Alibaba dataset. This discrepancy can be attributed to the distinct time ranges of behaviors in the two datasets. The Amazon dataset encompasses review data from 1996 to 2014, whereas the Alibaba dataset comprises user behaviors recorded over a continuous period of 8 days. The larger time range of the Amazon dataset renders TTE-T more informative than TTE-P. Conversely, on the Alibaba dataset with a relatively short time range, TTE-P proves to be more effective in capturing temporal patterns.




Figure 4: Time interval distribution and binning on Amazon and Alibaba dataset.

4.5 Existing models regarding the three components

Table 5 presents a comprehensive summary of existing models, including TIN, with respect to Temporal Information (TI), Target-aware Attention (TA), and Target-aware Representation (TR). DIN consists solely of TA, while DIN' incorporates both TA and TR. In DIEN [44], the Interest Extractor Layer (ExtLayer) and Interest Evolution Layer (EvoLayer) capture TI using the Gated Recurrent Unit (GRU). However, ExtLayer lacks both TA and TR, while EvoLayer lacks TR. DSIN [5] employs the Session Interest Extractor (SIE), which applies self-attention over historical behaviors. The Session Interest Activating (SIA) layer computes an attention weight for each behavior based on the behavior itself and the target item. DSIN then feeds the output of SIA to a Concatenation and Multi-Layer Perceptron (MLP) layer. Consequently, DSIN incorporates TI through self-attention and TA via SIA, but it lacks TR due to the usage of Concatenation and MLP. BST [3] combines both TI (achieved through position encoding) and TA using the Transformer model, which can be viewed as an extension of the DIN model. However, BST simply concatenates the output of the Transformer, which represents an attentive pooling of behavior embeddings, with the target item embedding. This concatenated representation then undergoes several MLP layers to make predictions. As mentioned earlier, the Concatenation and MLP are hard to learn the dot product (interaction) effectively [25].

4.6 Online A/B Testing

We apply TIN to the user's 2-year clicked ads category sequence feature in Tencent's WeChat Moments, and use both TTE-P and TTE-T to capture the temporal correlation. The backbone model employs a Heterogeneous Experts with Multi-Embedding architecture [6, 19, 28]. Specifically, we learn multiple different feature interaction experts, e.g., GwPFM [19] (a variant of FFM [9] and

Table 5: A summary of all models w.r.t. the three components: Temporal Information (TI), Target-aware Attention (TA), and Target-aware Representation (TR). COE denotes chronological order encoding, TTE denotes target-aware temporal encoding, $g_{MLP}(\cdot)$ denotes a Multi-Layer Perceptron and $[\cdot]$ denotes concatenation. Q , K , and V represent query, key, and value.

Model	Temporal Information?	How	TA	Attention	TR	Representation	Notes on Representation
Avg Pooling & Concat	✗	-	✗	-	✗	$g_{MLP}([e_i, v_i])$	-
Avg Pooling & Product	✗	-	✗	-	✓	$g_{MLP}([e_i, v_i, e_i \odot v_i])$	-
DIN' [45]	✗	-	✓	$\sigma(\langle e_i, v_i \rangle)$	✗	$g_{MLP}([e_i, v_i])$	Concat & MLP's hard to learn dot product [25].
DIN [45]	✗	-	✓	$\sigma(\langle e_i, v_i \rangle)$	✓	$g_{MLP}([e_i, v_i, e_i \odot v_i])$	According to code released by the authors.
SIM [21]	✓	TTE-T	✓	$\sigma(\langle e_i, v_i \rangle)$	✗	$g_{MLP}([e_i, v_i])$	Concat & MLP's hard to learn dot product [25].
GRU4Rec [7]	✓	GRU	✗	-	✓	$\langle h_i, v_i \rangle$	According to open source code.
SASRec [10]	✓	COE	✗	-	✓	$\langle \tilde{e}_i^V, v_i \rangle$	According to Eqn. $r_{t,i} = F_t^{(b)} N_i^T$ of [10].
BERT4Rec [29]	✓	COE	✗	-	✓	$\langle \tilde{e}_i^V, v_i \rangle$	According to Eqn.(7) in [29].
ExtLayer [44]	✓	AUGRU	✗	-	✗	-	-
EvoLayer [44]	✓	AUGRU	✓	$\sigma(\langle h_i, v_i \rangle)$	✗	$g_{MLP}([h_i, v_i])$	-
DSIN [5]	✓	Bi-LSTM & COE	✓	$\sigma(\langle \tilde{e}_i^K, v_i^Q \rangle)$	✗	$g_{MLP}([\tilde{e}_i^V, v_i])$	Concat & MLP's hard to learn dot product [25].
BST [3]	✓	COE	✓	$\sigma(\langle \tilde{e}_i^K, \tilde{v}_i^Q \rangle)$	✗	$g_{MLP}([\tilde{e}_i^V, v_i])$	Concat & MLP's hard to learn dot product [25].
TIN	✓	TTE	✓	$\sigma(\langle \tilde{e}_i, \tilde{v}_i \rangle)$	✓	$g_{MLP}([\tilde{e}_i, \tilde{v}_T, \tilde{e}_i \odot \tilde{v}_i])$	-
TIN w/o TTE	✗	-	✓	$\sigma(\langle \tilde{e}_i, \tilde{v}_i \rangle)$	✓	$g_{MLP}([\tilde{e}_i, \tilde{v}_i, \tilde{e}_i \odot \tilde{v}_i])$	-
TIN w/o TA	✓	TTE	✗	-	✓	$g_{MLP}([\tilde{e}_i, \tilde{v}_i, \tilde{e}_i \odot \tilde{v}_i])$	-
TIN w/o TR	✓	TTE	✓	$\sigma(\langle \tilde{e}_i, \tilde{v}_i \rangle)$	✗	$g_{MLP}([\tilde{e}_i, \tilde{v}_i])$	-

FwFM [18]), IPNN [22], DCN V2 [33], or FlatDNN. There are hundreds of features in total, mostly user behavior features, along with ads and context side features. Multiple embedding tables are learned for all features, each corresponding to one or several experts. TIN functions as an additional expert alongside the existing three experts, sharing embeddings with GwPFM and FlatDNN.

We conduct online A/B testing from September 2023 to October 2023. In the baseline, the user interest is modeled by feeding both user behavior features as well as target ads features to GwPFM or IPNN. We previously attempted to use DIN, but it did not surpass the baseline's performance. This is possibly due to that explicit interaction models such as GwPFM or IPNN are already able to learn the semantic correlation between behaviors and the target. During the two-week 20% A/B testing, TIN demonstrated promising results, achieving a 1.65% increase in cost and a 1.93% increase in GMV (Gross Merchandise Value) compared to the baseline. These improvements were statistically significant according to t-tests. Based on our estimation, these performance gains could potentially lead to a revenue increase of hundreds of millions of dollars per year. TIN has been successfully deployed on production in several scenarios.

5 RELATED WORK

Deep Interest Network (DIN) [45] introduced the concept of target attention, enabling the learning of attentive weights for each user behavior with respect to a target item. Subsequently, several works have emerged following the DIN framework. For instance, DIEN [44], DSIN [5], and BST [3] focus on modeling interest evolution. MIMN [20], SIM [21], HPMN [23], LimaRec [36], and ETA [2] address the challenge of modeling long sequence interests. Additionally, MIND [12], ComiRec [1], and LimaRec [36] tackle the task of modeling multiple interests.

Another line of research is sequential recommenders [8, 10, 14, 15, 27, 31, 38, 39, 41, 46], which aim to capture temporal correlations using either Recurrent Neural Networks (RNN) [7, 34, 40] or Self-Attention [10, 16, 29, 35]. However, these models lack target-aware attention, limiting them to only being able to capture 2nd or 3rd-order interactions and, therefore, unable to learn quadruple semantic-temporal correlation.

6 CONCLUSION

This paper investigates the crucial semantic-temporal correlation between user behaviors and the target. Examination of existing methods reveals that they fail to capture such correlation. We propose the Temporal Interest Network (TIN) to capture such correlation, which incorporates target-aware temporal encoding, target-aware attention, and target-aware representation. Comprehensive experiments on two public datasets demonstrate the superiority of TIN over the best-performing baselines. TIN has been successfully deployed to WeChat Moments on Tencent's advertising platform.

7 ACKNOWLEDGEMENT

We gratefully acknowledge the contributions of the following: Xueming Qiu, Lei Mu, Xian Hu, Yaqian Zhang, Ming Yue, Wenbo Liu, Xiaobo Li, Zezhen Xu, Jiayu Sun, Xun Liu, Yi Li, Ximei Wang, Junwen Cheng, Yan Tan, Yuxiong Li, Zhaohua Li, Kuo Zhang and Yufei Zheng.

This work was supported by the National Natural Science Foundation of China [U23A20309, 62272302, 62172276, 62372296], Shanghai Municipal Science and Technology Major Project [2021SHZDZX-0102], and the Tencent Rhinoceros Project. Haolin Zhou, Xinyi Zhou, Xiaofeng Gao and Guihai Chen are in the MoE Key Lab of Artificial Intelligence, Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China. Xiaofeng Gao is the corresponding author.

REFERENCES

- [1] Yukuo Cen, Jianwei Zhang, Xu Zou, Chang Zhou, Hongxia Yang, and Jie Tang. 2020. Controllable multi-interest framework for recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2942–2951.
- [2] Qiwei Chen, Changhua Pei, Shanshan Lv, Chao Li, Junfeng Ge, and Wenwu Ou. 2021. End-to-End User Behavior Retrieval in Click-Through Rate Prediction Model. *arXiv preprint arXiv:2108.04468* (2021).
- [3] Qiwei Chen, Huan Zhao, Wei Li, Pipei Huang, and Wenwu Ou. 2019. Behavior sequence transformer for e-commerce recommendation in alibaba. In *Proceedings of the 1st International Workshop on Deep Learning Practice for High-Dimensional Sparse Data*. 1–4.
- [4] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research* 12, 7 (2011).
- [5] Yufei Feng, Fuyu Lv, Weichen Shen, Menghan Wang, Fei Sun, Yu Zhu, and Keping Yang. 2019. Deep session interest network for click-through rate prediction. In *IJCAI*.
- [6] Xingzhuo Guo, Junwei Pan, Ximei Wang, Baixu Chen, Jie Jiang, and Mingsheng Long. 2023. On the Embedding Collapse when Scaling up Recommendation Models. *arXiv preprint arXiv:2310.04400* (2023).
- [7] Balázs Hidasi, Alexandros Karatzoglou, Lina Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [8] Yidan Hu, Yong Liu, Chunyan Miao, and Yuan Miao. 2022. Memory bank augmented long-tail sequential recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 791–801.
- [9] Yuchin Juan, Damien Lefortier, and Olivier Chapelle. 2017. Field-aware factorization machines in a real-world online advertising system. In *International Conference on World Wide Web (WWW)*. 680–688.
- [10] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.
- [11] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [12] Chao Li, Zhiyuan Liu, Mengmeng Wu, Yuchi Xu, Huan Zhao, Pipei Huang, Guoliang Kang, Qiwei Chen, Wei Li, and Dik Lun Lee. 2019. Multi-interest network with dynamic routing for recommendation at Tmall. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2615–2623.
- [13] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. xDeepFM: Combining explicit and implicit feature interactions for recommender systems. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*. 1754–1763.
- [14] Xiaolin Lin, Jinwei Luo, Junwei Pan, Weiwei Pan, Zhong Ming, Xun Liu, Shudong Huang, and Jie Jiang. 2024. Multi-Sequence Attentive User Representation Learning for Side-information Integrated Sequential Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 414–423.
- [15] Chang Liu, Xiaoguang Li, Guohao Cai, Zhenhua Dong, Hong Zhu, and Lifeng Shang. 2021. Noninvasive self-attention for side information fusion in sequential recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4249–4256.
- [16] Zhiwei Liu, Ziwei Fan, Yu Wang, and Philip S Yu. 2021. Augmenting sequential recommendation with pseudo-prior items via reversely pre-training transformer. In *Proceedings of the 44th international ACM SIGIR conference on Research and development in information retrieval*. 1608–1612.
- [17] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).
- [18] Junwei Pan, Jian Xu, Alfonso Lobos Ruiz, Wenliang Zhao, Shengjun Pan, Yu Sun, and Quan Lu. 2018. Field-weighted factorization machines for click-through rate prediction in display advertising. In *Proceedings of the 2018 World Wide Web Conference*. 1349–1357.
- [19] Junwei Pan, Wei Xue, Ximei Wang, Haibin Yu, Xun Liu, Shijie Quan, Xueming Qiu, Dapeng Liu, Lei Xiao, and Jie Jiang. 2024. Ad Recommendation in a Collapsed and Entangled World. *arXiv preprint arXiv:2403.00793* (2024).
- [20] Qi Pi, Weijie Bian, Guorui Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Practice on long sequential user behavior modeling for click-through rate prediction. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*. 2671–2679.
- [21] Qi Pi, Guorui Zhou, Yujing Zhang, Zhe Wang, Lejian Ren, Ying Fan, Xiaoqiang Zhu, and Kun Gai. 2020. Search-based user interest modeling with lifelong sequential behavior data for click-through rate prediction. In *ACM International Conference on Information & Knowledge Management (CIKM)*. 2685–2692.
- [22] Yanru Qu, Han Cai, Kan Ren, Weinan Zhang, Yong Yu, Ying Wen, and Jun Wang. 2016. Product-based neural networks for user response prediction. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 1149–1154.
- [23] Kan Ren, Jiarui Qin, Yuchen Fang, Weinan Zhang, Lei Zheng, Weijie Bian, Guorui Zhou, Jian Xu, Yong Yu, Xiaoqiang Zhu, et al. 2019. Lifelong sequential modeling with personalized memorization for user response prediction. In *International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 565–574.
- [24] Steffen Rendle. 2010. Factorization machines. In *2010 IEEE International conference on data mining*. IEEE, 995–1000.
- [25] Steffen Rendle, Walid Krichene, Li Zhang, and John Anderson. 2020. Neural collaborative filtering vs. matrix factorization revisited. In *RecSys*. 240–248.
- [26] Ravid Shwartz-Ziv and Naftali Tishby. 2017. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810* (2017).
- [27] Jiajie Su, Chaochao Chen, Zibin Lin, Xi Li, Weiming Liu, and Xiaolin Zheng. 2023. Personalized Behavior-Aware Transformer for Multi-Behavior Sequential Recommendation. In *Proceedings of the 31st ACM International Conference on Multimedia*. 6321–6331.
- [28] Liangcai Su, Junwei Pan, Ximei Wang, Xi Xiao, Shijie Quan, Xihua Chen, and Jie Jiang. 2023. STEM: Unleashing the Power of Embeddings for Multi-task Recommendation. *arXiv preprint arXiv:2308.13537* (2023).
- [29] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1441–1450.
- [30] Yang Sun, Junwei Pan, Alex Zhang, and Aaron Flores. 2021. FM2: field-matrixed factorization machines for recommender systems. In *Proceedings of the Web Conference 2021*. 2828–2837.
- [31] Jiaxi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the eleventh ACM international conference on web search and data mining*. 565–573.
- [32] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*. 5998–6008.
- [33] Ruoxi Wang, Rakesh Shivanna, Derek Cheng, Sagar Jain, Dong Lin, Lichan Hong, and Ed Chi. 2021. Dcn v2: Improved deep & cross network and practical lessons for web-scale learning to rank systems. In *Proceedings of the Web Conference 2021*. 1785–1797.
- [34] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J Smola, and How Jing. 2017. Recurrent recommender networks. In *Proceedings of the tenth ACM international conference on web search and data mining*. 495–503.
- [35] Liwei Wu, Shuqing Li, Cho-Jui Hsieh, and James Sharpnack. 2020. SSE-PT: Sequential recommendation via personalized transformer. In *Fourteenth ACM Conference on Recommender Systems*. 328–337.
- [36] Yongji Wu, Lu Yin, Defu Lian, Mingyang Yin, Neil Zhenqiang Gong, Jingren Zhou, and Hongxia Yang. 2021. Rethinking Lifelong Sequential Recommendation with Incremental Multi-Interest Attention. *arXiv preprint arXiv:2105.14060* (2021).
- [37] Sicong Xie, Qunwei Li, Weidi Xu, Kaiming Shen, Shaohu Chen, and Wenliang Zhong. 2022. Denoising Time Cycle Modeling for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1950–1955.
- [38] Xu Xie, Fei Sun, Zhaoyang Liu, Shiwen Wu, Jinyang Gao, Jiandong Zhang, Bolin Ding, and Bin Cui. 2022. Contrastive learning for sequential recommendation. In *2022 IEEE 38th international conference on data engineering (ICDE)*. IEEE, 1259–1273.
- [39] Yueqi Xie, Peilin Zhou, and Sunghun Kim. 2022. Decoupled Side Information Fusion for Sequential Recommendation. *arXiv preprint arXiv:2204.11046* (2022).
- [40] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. A dynamic recurrent model for next basket recommendation. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 729–732.
- [41] Fajie Yuan, Alexandros Karatzoglou, Ioannis Arapakis, Joemon M Jose, and Xiangan He. 2019. A simple convolutional generative network for next item recommendation. In *Proceedings of the twelfth ACM international conference on web search and data mining*. 582–590.
- [42] Jiaqi Zhai, Zhaojie Gong, Yueming Wang, Xiao Sun, Zheng Yan, Fu Li, and Xing Liu. 2023. Revisiting Neural Retrieval on Accelerators. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 5520–5531.
- [43] Zuowu Zheng, Xiaofeng Gao, Junwei Pan, Qi Luo, Guihai Chen, Dapeng Liu, and Jie Jiang. 2022. AutoAttention: Automatic Field Pair Selection for Attention in User Behavior Modeling. *arXiv preprint arXiv:2210.15154* (2022).
- [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. In *AAAI*. 5941–5948.
- [45] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *SIGKDD*. 1059–1068.
- [46] Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020. S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *Proceedings of the 29th ACM international conference on information & knowledge management*.

A APPENDIX

A.1 Experimental Setting

We assess the effectiveness of TIN by performance evaluation on two popular datasets for user response prediction.

The **Amazon**² dataset is derived from real product reviews and ratings on the Amazon e-commerce website. We conduct experiments on the Electronics subset, which contains 192,403 users, 63,001 goods, 801 categories and 1,689,188 samples. Each user in the dataset has a minimum of 5 reviews with various items. We adopt the widely employed leave-one-out strategy [45], that is, we predict the user’s feedback for the n -th item based on the user’s first $n - 1$ interaction records. The negative samples are generated by replacing the target item in positive samples with a random item from the entire item collection.

The **Alibaba**³ dataset encompasses a comprehensive collection of data obtained from a random selection of 1.14 million users on Taobao’s website, resulting in an 8-day advertising display /click log comprising 26 million records. We partition the dataset as follows: the samples from the first six days (2017-05-06 to 2017-05-11) are allocated for training purposes, the samples occurring on 2017-05-12 are utilized for validation, and the samples from the final day (2017-05-13) are reserved for testing. We follow the data preprocessing procedure of DSIN [5].

For the Amazon dataset, we leverage the `cate_id` and `item_id` fields to construct the embeddings for both the historical behavior and the target item. Specifically, we compose the embedding \mathbf{e}_i as follows: $\mathbf{e}_i = [\mathbf{c}(X_i), \mathbf{it}(X_i)]$, where $\mathbf{c}(X_i)$ denotes the category embedding of X_i , $\mathbf{it}(X_i)$ represents the item embedding of X_i , and $[\cdot]$ denotes the concatenation operator. On the other hand, for the Alibaba dataset, we utilize the `cate_id` and `brand_id` fields. In this case, the embedding \mathbf{e}_i is constructed as follows: $\mathbf{e}_i = [\mathbf{c}(X_i), \mathbf{b}(X_i)]$, where $\mathbf{b}(X_i)$ denotes the brand embedding of X_i .

For both datasets, we set the embedding dimension d to 64 for all feature embeddings. The hidden layers of the two-layer MLP have dimensions of 80 and 40 for the Amazon dataset, and 200 and 80 for the Alibaba dataset. As the optimizer, we employ Adam [11] with a learning rate of 0.001 for the Amazon dataset, and Adagrad [4] with a learning rate of 0.01 for the Alibaba dataset. To ensure consistency and comparability, we maintain the remaining settings identical to those used in DIN on the Amazon dataset, and DSIN⁴ on the Alibaba dataset. We evaluate the performance using Logloss and GAUC. To account for variability, each experiment is repeated five times, and we report the mean and standard deviation of the results.

A.2 Expressiveness Study of Target-aware Attention and Target-aware Representation

Both Target-aware Attention (TA) and Target-aware Representation (TR) are able to learn semantic correlation between behaviors and target. However, TR should be more expressive because there exists a solution *by construction* to make it the same as TA by simply

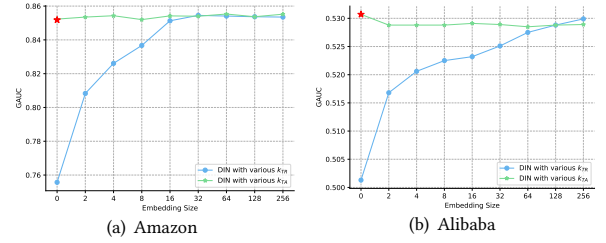


Figure 5: Expressiveness of target-aware attention and target-aware representation in DIN. The green line denotes the performance of DIN with various attention embedding sizes d_{TA} , while the blue line denotes the performance of DIN with various representation embedding sizes d_{TR} . The red star corresponds to totally disabling the target-aware attention.

summing up all elements of the representation vector of TR. We conduct a simple experiment to investigate their expressiveness on DIN which consists of both TA and TR. Specifically, we separate the embedding space of TA and TR in DIN, *i.e.*, separate $\{\mathbf{v}^{TA}, \mathbf{e}^{TA}\} \in \mathbb{R}^{d_{TA}}$ from $\{\mathbf{v}^{TR}, \mathbf{e}^{TR}\} \in \mathbb{R}^{d_{TR}}$. Formally,

$$\begin{aligned} \hat{y}_{DIN} &= \text{sigmoid}(g_{MLP}([\mathbf{u}_{int}, \mathbf{v}_i^{TR}, \mathbf{u}_{int} \odot \mathbf{v}_i^{TR}])) \\ \mathbf{u}_{int} &= \sum_{i \in \mathcal{H}} \alpha_{DIN}(\mathbf{e}_i^{TA}, \mathbf{v}_i^{TA}) \cdot \mathbf{e}_i^{TR} \\ \alpha_{DIN}(\cdot) &= \sigma\left(\frac{\langle \mathbf{v}_i^{TA}, \mathbf{e}_i^{TA} \rangle}{\sqrt{d_{TA}}}\right) \end{aligned} \quad (6)$$

We then tune the embedding size of each space, *i.e.*, d_{TA} , d_{TR} from 0 to 256 in one component while retaining it as 128 for the other one. As shown in Fig. 5, the performance drops dramatically with lower embedding size in TR, from 0.854 with $d_{TR} = 256$ to 0.808 with $d_{TR} = 2$, shown as the blue line. On the other hand, its performance keeps around 0.853 with various embedding sizes (d_{TA}) from 256 to 2, when retaining d_{TR} as 128, shown as the green line. This verifies that TR is more expressive than TA in learning the semantic correlation between behaviors and target.

Analysis of semantic-temporal correlation on more categories. We have also computed the learned semantic-temporal correlation for all the models on the top 200 frequent target categories. To evaluate the correlation between the ground truth and learned Category-wise Target-aware Correlation (CTC), we select the history behaviors with the same category as the target and calculate the Pearson coefficient. The coefficients are then categorized as follows: poor (range of [-1, 0.2]), medium (0.2, 0.8], and strong (0.8, 1.0]. We plot the distribution of these coefficients across the three bins for TIN and several other models. As depicted in Figure 6, a significant majority (90.0%) of the categories demonstrate a strong ability to learn the CTC in TIN. Similarly, TIN w/o TR also exhibits a commendable performance in learning the CTC. However, the remaining models do not achieve comparable results.

²<http://jmcauley.ucsd.edu/data/amazon/>

³<https://tianchi.aliyun.com/dataset/56>

⁴<https://github.com/shenweichen/DSIN/tree/master/code>

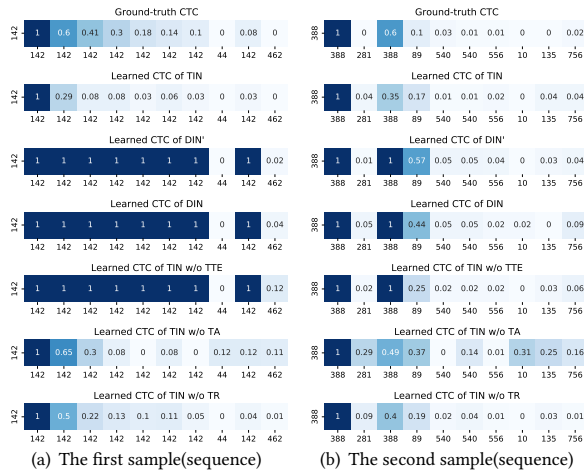


Figure 7: Visualization of ground-truth CTC and learned CTC between behavior sequence and the target item of 2 random samples from the Amazon dataset.

A.3 Visualization on User Behavior Sequences

We have randomly selected two samples (user behavior sequences) and visualized the ground truth and learned Category-wise Target-aware Correlation (CTC) for the latest 10 behaviors of each sequence in Figure 7. In the first sequence, 8 out of 10 behaviors share the same category as the target (142). The ground truth CTC exhibits a strong temporally decaying pattern over these 8 behaviors. Notably, both TIN and TIN without Target-aware Representation (TR) successfully capture this correlation, while other models fail to do so. Specifically, DIN⁺, DIN, and TIN without Target-aware Temporal Encoding (TTE) learn the same correlation for all 8 behaviors. In the second sequence, only 2 behaviors have the same category with the target (388), displaying an evident temporal decaying pattern. Conversely, the remaining 8 behaviors belong to diverse categories and has weaker correlation with the target according to the mutual information. TIN and TIN without TR effectively capture both the temporal and semantic correlations. However, DIN⁺, DIN, and TIN without TTE fail to learn the temporal correlation for the 1st and 3rd behaviors, while TIN without Target-aware Attention (TA) fails to capture the weak semantic correlation in the other behaviors. These visualizations demonstrate the superior performance of TIN in capturing both temporal and semantic correlations in specific user behavior sequences.




Figure 6: Distribution of Pearson coefficient regarding ground-truth and learned temporal correlation on the top-200 target categories. The range for poor, medium, and strong coefficients are $[-1, 0.2]$, $(0.2, 0.8]$, and $(0.8, 1.0]$, respectively. The results of DIN are omitted as it's the same with TIN w/o TTE.