

Positive, Negative and Neutral: Modeling Implicit Feedback in Session-based News Recommendation

Shansan Gong
Shanghai Jiao Tong University
Shanghai, China
gongshansan@sjtu.edu.cn

Kenny Q. Zhu*
Shanghai Jiao Tong University
Shanghai, China
kzhu@cs.sjtu.edu.cn

ABSTRACT

News recommendation for anonymous readers is a useful but challenging task for many news portals, where interactions between readers and articles are limited within a temporary login session. Previous works tend to formulate session-based recommendation as a next item prediction task, while they neglect the implicit feedback from user behaviors, which indicates what users really like or dislike. Hence, we propose a comprehensive framework to model user behaviors through positive feedback (i.e., the articles they spend more time on) and negative feedback (i.e., the articles they choose to skip without clicking in). Moreover, the framework implicitly models the user using their session start time, and the article using its initial publishing time, in what we call “neutral feedback”. Empirical evaluation on three real-world news datasets shows the framework’s promising performance of more accurate, diverse and even unexpectedness recommendations than other state-of-the-art session-based recommendation approaches.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

News recommendation; session-based; implicit feedback; time aware; cold-start

ACM Reference Format:

Shansan Gong and Kenny Q. Zhu. 2022. Positive, Negative and Neutral: Modeling Implicit Feedback in Session-based News Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*, July 11–15, 2022, Madrid, Spain. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3477495.3532040>

1 INTRODUCTION

Online news portals such as BBC, CNN, and Bing News have a huge number of readers daily. Many of them are anonymous or logged

*Kenny Q. Zhu is the corresponding author and is partially supported by NSFC Grant No. 91646205, SJTU-CMBCC Joint Research Scheme, and SJTU-Meituan Joint Research Scheme.

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 '22, July 11–15, 2022, Madrid, Spain

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-8732-3/22/07...\$15.00
<https://doi.org/10.1145/3477495.3532040>

in as guests who typically do not read many stories in a single login session. Given the limited interactions users engage with the portals, it is often hard for the systems to fully understand the user behaviors, posing significant challenges to recommendation systems.

Conventional news recommendation approaches tend to formulate the recommendation task as CTR prediction task, and they mainly rely on collaborative filtering and factorization machine [2, 6, 8, 12, 40, 52], which requires the system to keep track of the user history and can not be applied to anonymous visits or guest logins. Recent neural approaches for news recommendation mostly focus on encoding the text feature of articles with attention mechanism [37, 39, 44–46, 61] when modeling the user interest while paying little attention to the click behavior or the article-to-article transition. For example, they have not taken full advantage of the temporal information associated with the reading behavior, which is important especially when the interactions with the user are sparse.

Considering the above issues, it’s natural and realistic to formulate the streaming-like news recommendation task for anonymous users as a session-based recommendation task [4, 32, 34, 56, 57]. The task is to recommend the next item that the user might be interested in given the previous sequence of behaviors within a session, where a *session* is usually a short period of time (e.g., 30 minutes) during which the user is logged on. Session-based recommendation is widely used in the e-commerce or video streaming domain [24, 55], and can successfully capture users’ short-term intention transition process [3, 33]. However, they rarely consider the implicit feedback from user behaviors.

In this paper, we are interested in exploiting user actions outside the clicks themselves. We call them “implicit feedback”¹, as illustrated in Figure 1. Typical implicit feedback can be extracted from browsing the main page, reading an article, closing an article, backtracking [30], etc. We believe that modeling such implicit feedback “explicitly” in the session-based recommendation system can help the recommender understand user intention better. In this work, we focus on answering these questions:

- If a user clicked an article, did she really *like* it?
- If a user did not click an article, did she *dislike* it?
- How do we model the temporal characteristics of the user and the articles in the system?

First, in traditional recommendation systems, “clicks” usually indicate a “like” or a vote from the user, but things are a bit different for news reading. Users may be “tricked” into clicking an article [43] and once they realize that, they will quickly back out and switch to other articles. Thus the time a user spends on reading an article

¹To avoid misconception, here the term “feedback” is interchangeable with “signals”.

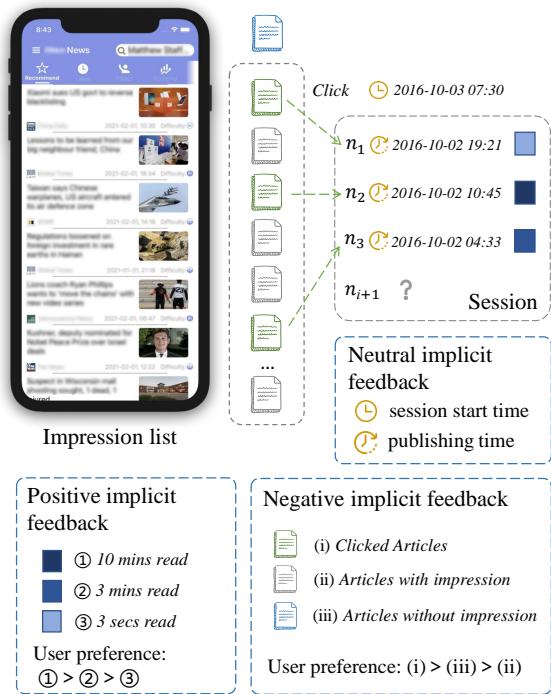


Figure 1: In session-based news reading, a user may spend different amounts of time on different clicked articles, representing different level of preference in an implicit form of positive feedback; a user’s impression of an article without an eventual click on the article represents an implicit form of negative feedback; the start time of the click and publishing time of articles can be viewed as neutral feedback.

is a better, finer-grained gauge of the user’s preference for the article [46], than just the click alone, which is only binary. We model this as the **implicit positive feedback** in this paper.

Second, just because the user did not click on an article does not necessarily mean the user does not like it; maybe she was never exposed to this article! We can infer what articles might have an impression [52] on the user during a session by assuming that articles are presented to the user roughly in the order of their publication time. Only those articles within her list of impressions but not clicked are considered “not interesting” to her. This is called **implicit negative feedback**.

Finally, while the positive and negative feedback helps us estimate the connection between the user and articles, some critical temporal information is useful to model the user and the articles individually. The session start time of a user may suggest the daily routine of that user. We can expect users who read on the same day of a week or same time of a day to have to share the same reading behavior or even background. On the other hand, the publishing time of each article can also be formed into a sequence in a session, which reflects the user’s sensitivity of the timeliness of the stories. We thus carefully design the representation of session start time and article publishing time as **implicit neutral feedback**.

In this paper, we formulate a session-based recommendation task to predict the next possible article in each session from the candidate articles pool. Our main contributions are:

- For the first time, we leverage the positive/negative/neutral implicit feedback in anonymous session-based news recommendation (Section 3);
- We design novel representations for temporal information and incorporate it with positive and negative feedback into a deep attention network;
- Our comprehensive offline evaluations on three real-world datasets show the clear advantage of our proposed method in terms of overall performance on diversity, accuracy and serendipity in both normal and cold-start scenarios (Section 4).

2 TASK DEFINITION

Assume that an anonymous user u produces a sequence of click events S_u with length T , which can be represented as

$$S_u : (n_1^u) \rightarrow (n_2^u) \rightarrow \dots \rightarrow (n_i^u) \rightarrow \dots,$$

where n_i^u denotes the id of the i -th clicked article.

In the training phase, the recommendation system aims to model the user feedback sequence as vector \mathbf{x}_s , maximizing the similarity between the vector of predicted next-article (n_{i+1}^u) and \mathbf{x}_s . While in the test phase, given interaction sequence S_u as input, we produce a ranking list R with the probability that the target user u is likely to click next from high to low. Typically a recommender needs to recommend more than one item for a user, thus the top- k items in R are recommended. Note that in our anonymous settings, each user u only appears once, which means training set and test set do not share the same users.

Table 1 summarizes some critical symbols and notations we use, and in later sections, for brevity, we ignore the superscript u when discussing the information for a particular user. To explain further, the set Imp_u denotes the articles in the vicinity of the articles being clicked, which may have left an impression on the user.

3 APPROACH

We first lay down some foundations for session-based recommendation, then present our base model which is a content-aware session-based model. After that, we introduce the key ideas of neutral, positive, and negative feedback, which are additional mechanisms that strengthen the base model. The high-level overview of our framework is illustrated in Figure 2.

3.1 Session-based Recommendation Basics

In a typical session-based setting, given the prefix sequence of the session, denoted as $S_u = (n_1, n_2, \dots, n_T)$, our goal is to predict n_{T+1} article that the target user u is most likely to click next. Following [20], they use an $N \times d_n$ item embedding matrix, where d_n is the embedding dimensionality, to provide article n_i ’s embedding vector as \mathbf{x}_i . Then methods like RNN [10], GNN [24, 50], or attention-based approaches [15, 20] can be used to encode the session information into vector \mathbf{x}_s from the sequence $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$, which represents the user’s history preferences. Meanwhile, the same item embedding matrix can be also regarded as N encoded

Table 1: Notations and descriptions

Symbol	Description
S_u	A session produced by an anonymous user u .
T	The number of articles in S_u .
i	$i \in [1, T]$ indexes articles in S_u .
n_i	The id of i -th article clicked by u .
c_i	The Content representation of the article n_i in S_u .
t_i	The active time that u stays in the article n_i .
\mathbf{ta}_i	The encoded active duration that u stays in n_i .
\mathbf{ts}_i	The encoded click time that u click for article n_i .
\mathbf{tp}_i	The encoded publishing time of article n_i .
Imp_u	A set of articles that appear in u 's impression list.
\mathbf{x}_i	The item embedding vector of n_i .
\mathbf{xc}_i	The embedding vector of n_i combined \mathbf{x}_i and c_i .
\mathbf{xc}_s	The contextual session vector of S_u .
\mathbf{xt}_s	The temporal session vector of S_u .
\mathbf{x}_s	The final session vector of S_u .
N	The total number of articles.
j	$j \in [1, N]$ indexes total N articles and is article id.
\hat{y}_j^u	The score of article j clicked by u in the next step.
R	Ranking list generated according to \hat{y}_j^u .
Ne_u	The set of negative samples for each S_u .

candidates $[\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$. For u , the cosine similarity score \hat{z}_j^u between the session representation and the article j is calculated by the inner product of the session vector \mathbf{x}_s and the candidate news article's vector:

$$\hat{z}_j^u = \mathbf{x}_j^T \mathbf{x}_s, j \in [1, N], \quad (1)$$

$$\hat{y}^u = \text{softmax}(\hat{\mathbf{z}}^u), \quad (2)$$

\hat{y}_j^u is normalized by softmax function in Eq. (2) to be the probability of the article j being clicked next in the session. The cross-entropy is usually used to compute loss:

$$\mathcal{L}_1 = -\frac{1}{|S|} \sum_{S_u \in S} \sum_{j=1}^N (y_j^u \log(\hat{y}_j^u) + (1 - y_j^u) \log(1 - \hat{y}_j^u)), \quad (3)$$

where S is the whole training sessions, $y_j^u = 1$ if the article j is indeed the next-clicked articles n_{T+1} in S_u and $y_j^u = 0$ otherwise.

3.2 Our Base Model: Content-aware Recommendation (CAR)

In order to recommend new and emerging articles, our starting point is a basic content-aware recommendation model. To encode articles' content information, Some use pre-trained article content embeddings [4], which is supervised trained based on the Word2Vec word embeddings in news titles and the metadata attributes of articles, such as categories. Specifically, we get the d_c -dimensional vectors from Word2Vec to represent the topic-oriented content of articles. Once we get the content vector c_i of article n_i , we concatenate c_i and \mathbf{x}_i to represent the article as \mathbf{xc}_i . To model the varying user preference to the articles in the same session, mainly following [20], we adopt a simple attention network, using the weighted sum of the input sequence of vectors to encode the whole session.

We define α_i , the attention weight of i -th articles n_i in session S_u as:

$$\alpha_i = W_0 \times \sigma(W_1 \times \mathbf{xc}_i + b_0), \quad (4)$$

$$\alpha'_i = \frac{\exp(\alpha_i)}{\sum_{i=1}^T \exp(\alpha_i)}, \quad (5)$$

where $W_0 \in \mathbb{R}^{1 \times d_n}$, $W_1 \in \mathbb{R}^{d_n \times (d_n + d_c)}$ are weighting parameters, \mathbf{xc}_i is the vector representing the article, and $b_0 \in \mathbb{R}^{d_n}$ is a bias. α'_i is normalized by softmax function.

Finally, the contextual session vector \mathbf{xc}_s of session S_u is defined as the weighted sum:

$$\mathbf{xc}_s = \sum_{i=1}^T \alpha'_i \mathbf{xc}_i. \quad (6)$$

Noted that in order to obtain the sequential information of the input sequence, the attention mechanism extract \mathbf{xc}_s utilizing the positional information by means of the positional encoding, which is the same as it is in Transformer architecture [35]. In the end, we can optimize the loss function according to Eq. (1)~Eq. (3); when computing \mathcal{L}_1 , the \mathbf{x}_j in Eq. (1) should be replaced by \mathbf{xc}_j and \mathbf{x}_s is equal to \mathbf{xc}_s in this base model.

3.3 Modeling Time as Neutral Feedback

Active time that one user stays on one particular article is the duration time. For the different and continuous active time t_i , we design a duration encoder, which bucketizes t_i into a discrete variable [46] by:

$$t'_i = \lfloor \log_2 t_i \rfloor, \quad (7)$$

and we map discrete t'_i into m distinct categories, where each category shares the same embedding vector \mathbf{ta}_i .

Different from the duration time which reveals the positive implicit feedback from users, the date-time when a certain event happens carries physical meaning and also conveys the regular pattern behind the temporal information, hence we design a temporal encoder for this kind of temporal representation. We encode the publishing time of an article and the start time of a session. In order to extract periodic temporal change which is also known as seasonality, we feed the month, day of the month, day of the week, hour, minute ($s \in \mathbb{R}^{12}$, $d \in \mathbb{R}^{31}$, $w \in \mathbb{R}^7$, $h \in \mathbb{R}^{24}$, $m \in \mathbb{R}^{60}$) into a time embedding matrix $\mathbf{E}_t \in \mathbb{R}^{134 \times d_t}$ with d_t dimension, and concatenate five parts as $\mathbf{t} \in \mathbb{R}^{5d_t}$ to build the whole temporal vector and to represent the date time. We next introduce how to utilize start time and publishing time, then we focus on modeling positive and negative implicit feedback which both derive from temporal information.

3.3.1 Start time. Users who start reading at a similar time are more likely to share the same reading behavior, which means that user interests are influenced by the start time. For example, some people tend to read financial news in the morning but instead read entertainment in the evening. We denote the click time from each click behavior of a session as $\mathbf{ts}_i \in \mathbb{R}^{2d_t}$ using the week and the hour (w, h), which is enough to capture different user's daily routine. To model the different informativeness of the articles in S_u for users' reading at different start reading time, we apply this information to compute personalized attention. We first transform the start time

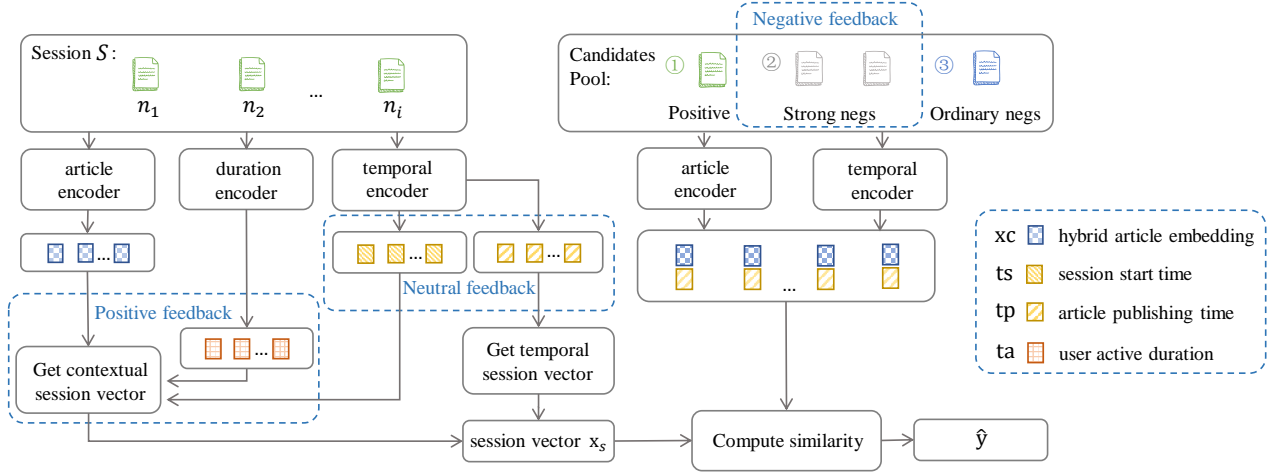


Figure 2: The architecture of our model. Squares in the figure represent vectors, and their colors refer to the different encoders that produce them.

embedding vector \mathbf{ts}_i into a preference query \mathbf{q}_i , which is similar to the “query” part in Transformer architecture:

$$\mathbf{q}_i = \text{ReLU}(W_t \times \mathbf{ts}_i + b_t), \quad (8)$$

where $W_t \in \mathbb{R}^{d_n \times 2d_t}$ is the projection parameter, $b_t \in \mathbb{R}^{d_n}$ is a bias.

Then we evaluate the importance of the interactions between preference query \mathbf{q}_i and article representation \mathbf{c}_i as attention α^t :

$$\alpha_i^t = \mathbf{c}_i \times \tanh(W'_t \times \mathbf{q}_i + b'_t), \quad (9)$$

$$\alpha_i^{t'} = \frac{\exp(\alpha_i^t)}{\sum_{j=1}^{|S_u|} \exp(\alpha_j^t)}, \quad (10)$$

where $W'_t \in \mathbb{R}^{(d_c+d_n) \times d_n}$ and $b'_t \in \mathbb{R}^{d_c+d_n}$ are weighting parameters. The contextual vector representation in Eq. (6) is now modified to:

$$\mathbf{xc}_s = \sum_{i=1}^{|S_u|} (\alpha_i^t + \alpha_i^{t'}) \mathbf{xc}_i. \quad (11)$$

3.3.2 Publish time. Users’ reading habits are reflected in the sequence of publishing time tp_1, \dots, tp_i in S_u . We can make inferences whether the user tends to browse new articles or older ones from this. The publishing time of clicked articles is a relatively independent sequence thus we model it separately. Due to the high density of article publishing time, we construct publishing time embedding vector $\mathbf{tp}_i \in \mathbb{R}^{5d_t}$ using (s, d, w, h, m) . We obtain the session temporal representation vector \mathbf{xt}_s by applying a similar attention mechanism in Section 3.2. We add the content vector of each article to capture the attention relation between the article content and its publishing time. The attention weight with click-level content information involved is formulated as:

$$\alpha_i^{tp} = W'_0 \times \sigma(W'_1 \times \mathbf{tp}_i + W'_2 \times \mathbf{c}_i + b'_0), \quad (12)$$

$$\alpha_i^{tp'} = \frac{\exp(\alpha_i^{tp})}{\sum_{j=1}^{|S_u|} \exp(\alpha_j^{tp})}, \quad (13)$$

where $W'_0 \in \mathbb{R}^{1 \times d_n}$, $W'_1 \in \mathbb{R}^{(5d_t) \times d_n}$, $W'_2 \in \mathbb{R}^{d_c \times d_n}$ and $b'_0 \in \mathbb{R}^{1 \times d_n}$. The final temporal session representation is:

$$\mathbf{xt}_s = \sum_{i=1}^{|S_u|} \alpha_i^{tp'} \mathbf{tp}_i. \quad (14)$$

In the end, we concatenate \mathbf{xt}_s and \mathbf{xc}_s as the aggregated representation $\mathbf{x}_s \in \mathbb{R}^{d_n+d_c+5d_t}$ for the whole session. As for computing \mathcal{L}_1 , the \mathbf{x}_j in Eq. (1) should be replaced by $\mathbf{xc}_j \oplus \mathbf{tp}_j$ (\oplus stands for the concatenation operation).

3.4 Modeling Positive Feedback

Our implicit positive feedback takes the form of the *active time* interval that a user spent on each article after clicking on it. If the user spends a short time in an article, it’s probably because the user is fooled by the title but actually does not like the article [21]. Note that if the active time is not explicitly available, it can be estimated by the time interval between the user’s two consecutive clicks.

As illustrated in Section 3.3, each degree of active time shares the same embedding vector \mathbf{ta}_i , representing to what extent the positive feedback is. We feed this vector into the attention computation as extra click-level feedback. Now, α_i in Eq. (4) is modified to:

$$\alpha_i = W_0 \times \sigma(W_1 \times \mathbf{xc}_i + W_2 \times \mathbf{ta}_i + b_0), \quad (15)$$

where $W_2 \in \mathbb{R}^{d_n \times d_t}$ is the projection parameter that map the active time embedding with d_t dimension into another dimension space. The contextual session vector \mathbf{xc}_s still follows Eq. (6) and final session vector \mathbf{x}_s is combined with \mathbf{xt}_s and \mathbf{xc}_s .

3.5 Modeling Negative Feedback

The most straight-forward and widely adopted negative sampling strategy is the random sampling from a non-clicked set of items, or from a global buffer with the last N clicks [4]. The major problem of randomly sampled items is that these items might be completely unrelated to the user, posing too little challenge for the model to learn. On the contrary, an informative item should be able to

confuse the model whether it has discovered more complex hidden meaning of user interaction or not.

While reading news, a user scrolls up and down the news stream, and the articles that are exposed to the user collectively form an impression list Imp_u . We take unclicked articles in Imp_u as more informative negative signals than other candidates [52] and thus we should treat them differently when counting loss, which means we should penalize the similarity between \mathbf{x}_s and those strong negative samples more strictly. This idea is similar to utilizing grayscale data [19] and contrastive learning [27], where we both consider the different degrees of information carried from different items.

As we discussed before, since the impression list is not always explicitly available, we assume an article is more likely to be in Imp_u if it was published nearby an article that has been clicked by u . Specifically, we sort the candidate articles according to their publishing time, and keep the nearby articles with the window size 300 and sample items from this window. We aim to minimize the cosine score between \mathbf{x}_s and the vector \mathbf{x}_j of negative sample j when $j \in Ne_u$, where $Ne_u \subseteq Imp_u$ is the set of negative samples for session S_u , thus we add this constraint into the final loss:

$$\mathcal{L}_2 = -\frac{1}{|S|} \sum_{S_u \in S} \sum_{j=1}^N (y_j^u \log(\hat{y}_j^u) + (1 - y_j^u) \log(1 - \hat{y}_j^u)) + \lambda \mathbb{1}(j \in Ne_u) \log(\sigma(1 - \mathbf{x}_j^T \mathbf{x}_s)), \quad (16)$$

where $\mathbb{1}(\cdot)$ returns 1 if the expression is true, λ is the weighting parameter of loss from negative articles. We jointly optimize these two losses with Adam optimizer.

4 EXPERIMENTS

In this section, we conduct experiments on three real-world news datasets: Adressa [7], Globo [4, 23] and MIND [47].

4.1 Experimental Setup

4.1.1 Data Preprocessing. In dataset preprocessing, we treat a series of click events from one anonymous user within 30 minutes as a session. We discard the sessions no longer than 1 click. To augment limited training data, for a session of n clicks, we create $n - 1$ mini-sessions, each starting from the first click of the original session and ending at the 2nd, 3rd through the last click of the original session. The article clicked at the end of every mini-session is the label to be predicted. The dataset statistics are in Table 2, where each session is quite short. The public Globo dataset covers 16 days, and only provides the extracted vectors of articles. We only choose a subset of days (20 days) in Adressa (the whole dataset lasts for 3 months) following [4] for simplicity. As for MIND, the data from the first several weeks is users’ historical logs without the session information, so we choose the data from the last week (7 days). The active time is missing in MIND, but it explicitly contains the impression list of each session.

4.1.2 Train/test set split. In order to simulate the real-time recommendation scenario, we choose a standard framework [13] which trains the model in the first few days and test the trained model in the remaining days. Each dataset can be split into several folds, we

Table 2: Dataset statistics (after preprocessing)

Dataset	# sessions	# articles	# topics	clicks/ session	clicks/ article
Globo	1M	45k	461	2.69	64
Adressa	0.5M	12k	23	2.79	117
MIND	0.2M	7k	16	2.38	59

will average the results over these folds. For Globo dataset, we split every 4 days into one fold, with 3 days for training and 1 day for testing, and 4 folds in total. For Adressa dataset, we split every 10 days into one fold due to its fewer session data in one day, and we need to extend the training days to keep the similar size of training data with Globo. We average the metrics performance of each fold in the end. For MIND dataset, we leave the last day as the test set to make one fold. After data preprocessing, the ratio between training data size and test data size is around 6:1 for Globo, and 10:1 for the other two datasets.

4.1.3 Metrics. During the test, given the first few click events in a session, the model generates a top- k recommendation list R with descending probabilities. We use widely-used metrics HR@ k , NDCG@ k to measure the model’s prediction accuracy.

Intra List Diversity (ILD@ k) [33] evaluates the topical/semantic diversity in R , and reflects the model’s ability to recommend different items to the same user.

$$ILD@k = \frac{1}{|R|(|R| - 1)} \sum_{a \in R} \sum_{b \in R} d(a, b), \quad (17)$$

where $d(a, b)$ is a distance measure between item a and b , and $d(a, b) = 1$ if item a, b belong to different topics (categories), 0 otherwise.

Besides, we expect the system to recommend unseen items to surprise users. The content-based unexpectedness metric (unEXP) [14] can be used to measure this kind of unexpectedness, which is calculated as follows:

$$unEXP_u@k = \frac{1}{|R|} \sum_{a \in R} \frac{1}{|S_u|} \sum_{b \in S_u} d(a, b). \quad (18)$$

4.1.4 Baseline Algorithms. Strong baselines are listed as follows. Their detail explanations are in Section 5.2.

Simple session-based recommenders: Despite simplicity of some of those methods, they still have competitive accuracy. CBCF [32] is a news recommender combines Content-Based similarity with session-based Collaborative Filtering similarity. STAN [5] is an extended version of SKNN (Session-KNN) with three controllable temporal decay factors.

Session-based neural recommenders: GRU4Rec [10, 11] is a Gated Recurrent Unit for recommendation, building with gated recurrent neural networks, which is similar to LSTM in [4]. SASRec [15] is a self-attention based Sequential model, adopting Transformer architecture to model user’s action. STAMP [20] is a Short-Term Attention/Memory Priority Model introducing the attention mechanism to model the relationship of each historical click and the last click. SRGNN [50] is a Session-based Recommender using Graph Neural

Table 3: Main and ablation results ($k = 20$ by default in our all tables). All results are averaged over all folds. The best baseline result on each metric is marked with * and overall best results are bolded. The “Ours” is our whole model and (-) means to ablate the corresponding module, where “neut”, “pos” and “neg” respectively refer to our neutral, positive and negative feedback modules. The last column is to replace our negative sampling strategy with random sampling. ↓ indicates performance drop over the whole model.

Datasets	Metrics	CBCF	STAN	GRU4Rec	SASRec	SRGNN	SGNNHN	STAMP	Ours	(-)neut	(-)pos	(-)neg	random
Adressa	HR	0.0957	0.1130	0.1120	0.1205	0.1152	0.1285	0.1287*	0.1658	0.1344↓	0.1619↓	0.1658	0.1646↓
	NDCG	0.0341	0.0500	0.0511	0.0509	0.0536	0.0562	0.0575*	0.0730	0.0613↓	0.0690↓	0.0720↓	0.0693↓
	ILD	0.2337	0.2409	0.8170	0.7856	0.8611*	0.8403	0.8445	0.8085	0.8204	0.8249	0.8237	0.8234
	unEXP	0.2509	0.2407	0.6949	0.8010	0.4754	0.8059*	0.5728	0.8279	0.8243↓	0.8333	0.8267↓	0.8346
Globo	HR	0.1185	0.1273	0.1280	0.1409	0.1280	0.1414	0.1435*	0.1852	0.1460↓	0.1817↓	0.1821↓	0.1847↓
	NDCG	0.0474	0.0647	0.0599	0.0620	0.0627	0.0611	0.0698*	0.0936	0.0727↓	0.0907↓	0.0940	0.0933↓
	ILD	0.3874	0.3087	0.9377	0.9864*	0.9248	0.9415	0.7980	0.8702	0.8362↓	0.8685↓	0.8927	0.8739
	unEXP	0.3730	0.2921	0.9771*	0.9690	0.6383	0.9467	0.8437	0.8358	0.8142↓	0.8252↓	0.8489	0.8317↓
MIND	HR	0.0315	0.0312	0.0338	0.0355	0.0334	0.0366	0.0371*	0.0495	0.0445↓	-	0.0471↓	0.0457↓
	NDCG	0.0110	0.0142	0.0132	0.0139	0.0144	0.0122	0.0151*	0.0211	0.0198↓	-	0.0180↓	0.0204↓
	ILD	0.7166	0.3193	0.8662	0.8562	0.8706	0.8775*	0.8452	0.8813	0.8779↓	-	0.8808↓	0.8858
	unEXP	0.6039	0.1064	0.8578	0.8654*	0.4508	0.4514	0.7544	0.8617	0.8415↓	-	0.8623	0.8680

¹ We conduct t significance test between the best score (if ours) and the second-best score for the main results, and the improvement is strongly significant as $p < 0.001$. Between the results of the whole model and the ablation model, the decline is significant as $p < 0.01$.

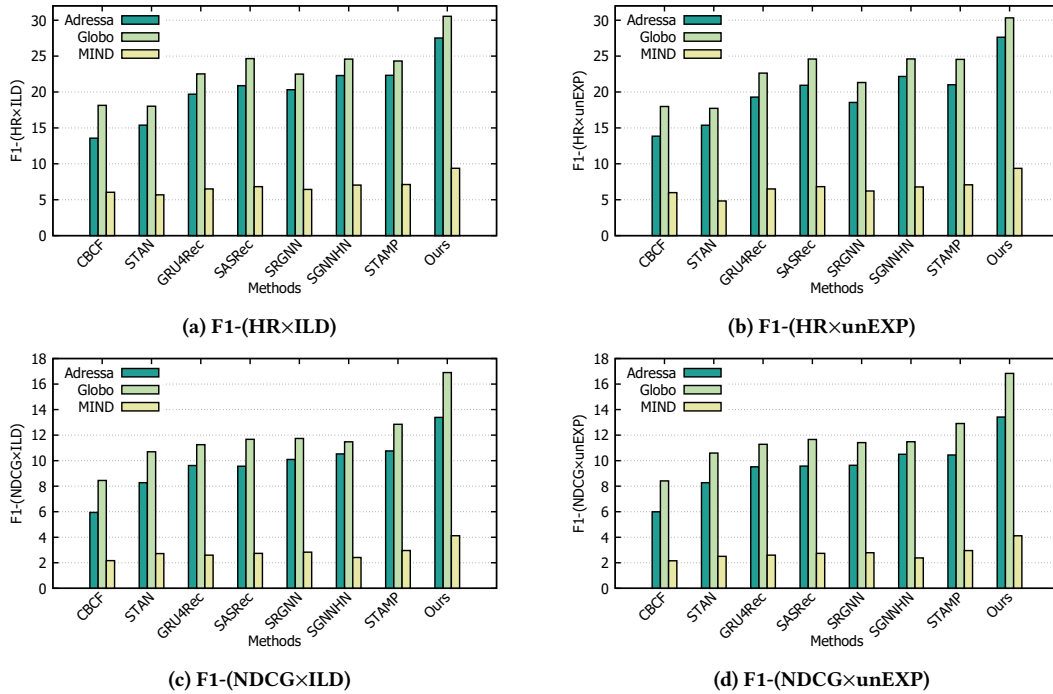


Figure 3: The graphical comparison of all methods on 4 different metrics and 3 different datasets. “Ours” is our approach.

Networks to capture complex transitions of items. *SGNNHN* [24] is improved SR-GNN using Star Graph Neural Network.

For session-based neural recommenders, we initialize the item embeddings with items’ content vector for fair comparison.

Neural news recommendation approaches: CPRS [46] is a typical news recommendation approach that utilizes the textual feature of articles to model user’s interests. It also uses the dwell time (i.e. active time) to measure user’s satisfaction. We make this approach adapt to the session-based scenario. Since only Adressa dataset

contains complete information for CPRS, we only compare it with our method in Adressa dataset and this is discussed in Table 4.

4.1.5 Implementation Details. For fair comparisons, we apply all baselines and our method to the same augmented data and train models on one GTX1080Ti GPU². We use the same latent dimension $d_n = d_c = 250, d_t = 64$, choose different learning rate in $\{0.002, 0.001, 0.0005\}$, batch size in $\{512, 1024, 2048\}$ and other hyper-parameters to select the best model using early stopping strategy based on the HR@20 score on the validation set, which is the last 10% of samples in the training set sorted by time. All embedding tables are normally initialized with 0 mean, 0.002 standard deviation, and for weighting parameters 0 mean, 0.05 standard deviation.

4.2 Main Results

In Table 3 and Figure 3, we compare the performance of all baselines and our approach, and we can make the following observations.

Non-neural methods CBCF and STAN are either considering the content information or the recency of the current session, and their results are somehow comparable to deep learning methods in three datasets. However, they generate recommendation lists with low diversity/novelty, mainly because their simple algorithms cannot capture enough personalized information. For session-based approaches, generally speaking, STAMP and SGNNHN yield better performance on HR and NDCG, but not always good at ILD/unEXP, while SASRec and SRGNN recommend more diverse but less accurate, showing the trade-off between diversity and accuracy. From the user’s aspect, though, when ILD/unEXP is over a threshold (like around 0.8^3), it’s hard for them to distinguish the difference, thus the ILD/unEXP score of our model is bearable.

As for our whole model, when compared with STAMP, it performs better or close on both accuracy and diversity. This result shows that **our model mitigates the dilemma between accuracy and diversity** to a certain extent. In the MIND dataset, the improvement is comparatively small and the possible reasons are: on the one hand, MIND did not provide active interval, nor did they give click time of each article (just start time of a session), we cannot get positive feedback from the data; on the other hand, from the results of CBCF, we assume the article transition information is too sparse and thus it is hard to recommend. Note that this dataset is not designed for the session-based recommendation, hence some information may be inaccurate (e.g., one session may last for days, longer than 30 minutes).

4.3 Effectiveness of Components

From Table 3 we can verify the effects of modeling user positive negative and neutral implicit feedback in our model. Compared with the whole model, there is a huge drop after removing neutral information and this is the most consistent over all metrics, which reveals the importance of neutral information (temporal information), and we will discuss the modeling of it in detail (Section 4.6). We cannot get positive feedback from MIND so this column is empty, and the reasons why the improvement is limited for MIND are analyzed previously.

²The implementation of our approach and baselines is released at <https://github.com/summeer/session-based-news-recommendation>

³We recruit two volunteers to measure the diversity of 50 samples to get this consensus

Adressa provides **the most complete information**, which not only releases the original text of articles instead of the extracted vectors in Globo, but also gives the accurate active time of the user in each article, while we can only estimate the active time by the interval between two successive clicks for Globo, which may not be accurate. After removing the positive implicit feedback module, in Adressa dataset, the HR and NDCG drop by 2.4% and 5.5% respectively, while in Globo dataset, they drop by 1.9% and 3.1%. The positive information performs similarly in both Adressa and Globo datasets, implying that our approximate estimation is reasonable. Further, the positive implicit feedback is more favorable on the Adressa dataset due to the more precise information.

We observe that negative information is less effective than positive information, especially by diversity/novelty metrics. One explanation is that the negative samples from the impression list are reconstructed based on their publishing time, so the information is not totally reliable. Negative sampling module lowers diversity, possibly because in the dataset the negative samples and the positive article usually belong to different categories, thus adding this module forces the model to recommend similar articles to the positive one. Negative feedback is better modeled in MIND due to its complete impression data. To verify the effect of the negative sampling strategy more accurately, we set the control group with random sampling, and we find that even though the random sampling would decrease the performance slightly, our negative feedback shows superior performance over it. The possible reason for the worse performance of using random sampling is that randomly sampled negative items have the possibility to be liked by this user, and this module imports some noise instead because this sample strategy does not consider **what the user really likes**.

4.4 Effects of Positive Feedback

CPRS considers the active time to represent the user’s satisfaction using personalized reading speed, which is quite similar but more complex than our positive feedback modeling. We firstly modify this method to meet the session-based scenario, and secondly plug the personalized reading speed into our model. The experiment is conducted in Adressa due to its complete information (Globo does not provide the text-level content and the active time is missing for MIND). In Table 4, the poor score of CPRS shows that when it is adapted to the session-based scenario, limited interactions are the bottleneck. When we adopt their personalized reading speed instead of the reading time, there is no significant improvement, and we hypothesize that for this dataset, when news reading **the reading speed is quite similar** for different users.

Table 4: Results of CPRS and CPRS module plugged into our model in Adressa dataset.

Methods	HR	NDCG	ILD	unEXP
Our whole model	0.1658	0.0730	0.8085	0.8279
CPRS	0.0812	0.0371	0.8191	0.8109
Ours using speed	0.1641	0.0674	0.8457	0.8245
Ours using 1-d t_i	0.1603	0.0705	0.8293	0.8220

To verify the effectiveness of the duration encoder, we compare it with the continuous active time vector t_i regarded as a one-dimensional vector instead of using distinct categories. As the result shows, the accuracy score of 1-d t_i is inferior to our whole model, indicating that our duration encoder catch more personalized information by bucketizing it.

4.5 Effects of Negative Feedback

In this section, we first validate our assumption for the negative user feedback, which is that articles whose publishing time is close to the clicked articles are likely presented to the user, or within their impressions. We do that on the MIND dataset, in which the real impressions and the publishing time of articles are all available for the sessions. For each session S_u , we sample negative items Ne_u using our strategy, and compute Jaccard similarity between Ne_u and real Imp_u , the overall score is 0.0062 when $|Ne_u| = 100$, compared with 0.0044 when random sampling. This shows that our assumption is reasonable and our strategy can **better reconstruct the impression list**.

We further conduct a parameter analysis on loss weights λ and the number of negative samples $|Ne|$ in Eq. (16). We report results on Adressa as an example, and results from other datasets are similar. The number of negative samples does not matter much, as shown in Figure 4a, so for simplicity we choose $|Ne| = 20$. The performance gets worse if λ is set too low or too high in Figure 4b, we conclude that the negative loss is useful but too many weights on it will **harm** the learning of the user’s positive feedback.

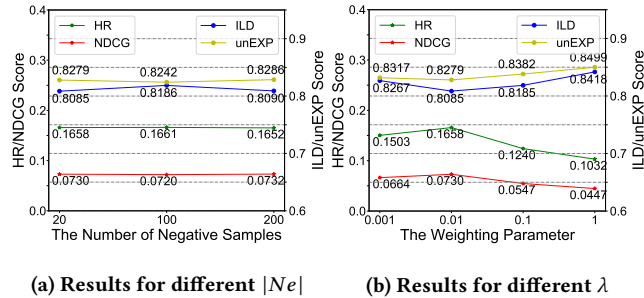


Figure 4: Hyper parameter discussion.

4.6 Temporal Representations

Since both the start time and the publishing time use the temporal encoder, we wonder if it would be better to have them share the embedding space. Table 5 shows the findings.

We can see that sharing the time embedding between publishing time and session start time has clear advantages in most of the metrics. This is because publishing time is associated with every article and there are a lot of such data for training, whereas the session start time suffers from lack of data and is less trained. Training the two jointly implicitly helps each other. It also makes physical sense because a Monday is a Monday regardless a story breaks on that day or a reader pops in to read that story on that day.

The ablation tests of using only publishing time or only start time in Table 5 also clearly indicates that temporal modeling both

Table 5: Different ways of utilizing temporal information in Adressa, where “p” stands for the publishing time and “s” stands for the start time.

Methods	HR	NDCG	ILD	unEXP
whole (shared)	0.1658	0.0730	0.8085	0.8279
whole (no share)	0.1620	0.0727	0.8310	0.8215
whole-p-s	0.1353	0.0612	0.8280	0.8276
whole-p	0.1344	0.0613	0.8204	0.8243
whole-s	0.1620	0.0726	0.8325	0.8415

from the item point of view and the user point of view are useful in garner latent information between the two. In other words, Table 5 shows that the neutral feedback works.

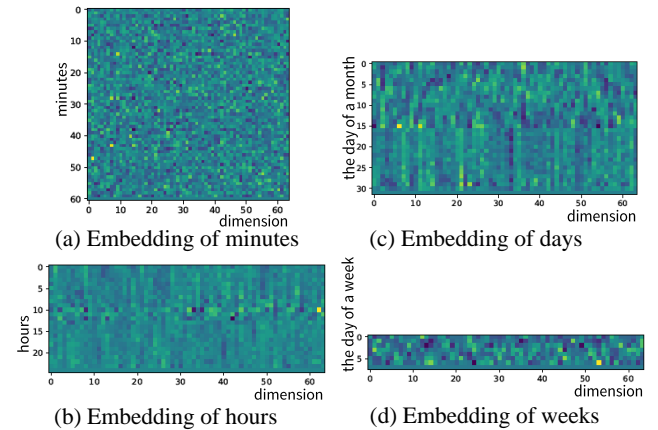


Figure 5: The visualization of time embedding tables for the day of a month, the day of a week, hour and minute, trained on Globo dataset.

To give some concrete evidence that the time embedding that we train carries some physical meanings, we visualize the embedding tables for *minutes*, *hours*, *weekdays* and *days of a month*, in Figure 5, which is trained on a subset of the Globo dataset. Some interesting patterns can be observed. For example, the representation of the minutes is rather uniform and random, because news publishing and reading can happen any minute of an hour. But there are certainly more activities at certain hours during a day. There are also some irregular patterns for weekends as shown in (c). Finally, because we only have the first 15 days of training data in this dataset, values for dates 16-31 in (d) are not fully trained, which only has the chance to update when the publishing date of articles falls in the range of 16-31, but most of articles are published nearby the click time according to the dataset statistics.

4.7 Discussion

4.7.1 Article cold-start. For news recommendation, all methods suffer from article cold-start problem due to the continuously published news, the analyses of the article cold-start scenario can help

us figure out where our improvement comes from. Another concern about the article cold-start problem is that if fresh articles can not get exposure reasonably, they will suffer from the Matthew effect and will not be clicked anymore. According to [4], instead of using user-oriented metric ILD/unEXP, we thus consider the system-oriented Item Coverage (COV@ k) as an additional metric. COV@ k is also called “aggregate diversity”, and reflects the model’s ability to recommend different items to different users, which forces the model to make a larger fraction of the item catalog visible to the users. We compute COV@ k as the number of distinct articles that appeared in any R divided by the number of distinct last clicked articles in the test set.

Table 6: Cold-start performance on Globo in the first fold. HR score is listed as percentage due to its relatively small value, and all scores are reserved to the second decimal.

Methods	Cold(80.3%)		non-Cold(19.7%)		Total	
	HR(%)	COV	HR(%)	COV	HR(%)	COV
CBCF	3.69	5.06	24.88	5.54	7.87	3.95
STAN	-	-	26.52	1.63	5.22	0.93
GRU4Rec	1.51	0.03	20.93	0.88	5.33	0.50
SASRec	0.80	0.01	23.35	1.28	5.25	0.73
SR-GNN	1.00	0.01	23.65	0.99	5.46	0.57
STAMP	1.72	0.01	21.84	1.04	5.68	0.59
SGNNHN	0.89	0.01	24.86	0.05	5.61	0.04
Ours	4.96	0.74	25.27	1.87	8.96	1.20

Table 6 lists the results of one fold in Globo dataset, and we choose it because it suffers from the most severe cold-start problem. For cold situation, where the test articles are completely disjoint from the training data, STAN does worse because it can not handle unseen items. Deep learning methods tend to predict the same articles for different users. Even though methods like SASRec yields not bad results, the models tend to overfit to popular articles. Our model, on the other hand, not only performs well on HR@20 but also gets the comparable COV@20 score, and the difference with the other deep learning methods is remarkable. In non-cold situation, the performance of all methods is close. The overall recommendation results largely depend on how a method does for cold-start scenarios.

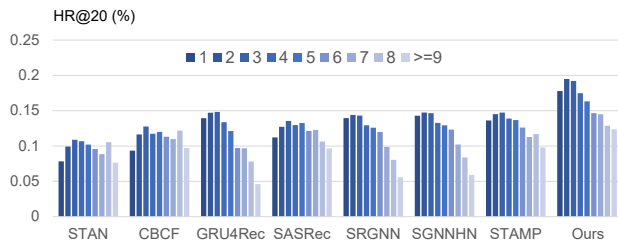


Figure 6: Accuracy with different session lengths.

4.7.2 User cold-start. Since anonymous news sessions are short with the average length of less than 3, the user cold-start problem is severe. We show the results for different input lengths in Figure 6. We report results on Globo, and other datasets perform similarly. Interestingly, our model reaches its peak accuracy from length 1 to 2. In contrast, other methods all reach the peak at 3. This shows our model is capable of **capturing user interests earlier** in the session by leveraging the user’s implicit feedback. For longer input length, the difference between our model and others narrows, indicating the similar ability to predict a user’s interests given a longer history. We can observe that the performance drops with the longer length input, and this may contribute to the noise that is imported from the longer reading history, which means it is harder to recommend when considering the longer and more diverse interests from users.

5 RELATED WORK

In this section, we discuss works in the area of news recommendation and session-based recommendation in news and other domains like e-commerce or music, and we also compare them with our proposed approach.

5.1 News Recommendation

First, the news recommendation task can be formulated as conventional recommendation tasks, the account of a user is reserved and articles are recommended based on users’ long-term click history. Some works use a well-designed graph network to represent the target user and the clicked article [6, 12]. In this situations, the relation of items and users are well exploited. Unfortunately, in real-time scenarios, new articles and anonymous users emerge, causing a severe cold-start problem. Then if we want to capture users’ preferences within sessions and recommend articles with their several interactions as input, this kind of approach with the static user-item matrix is not suitable. Some propose incremental matrix factorization algorithm based on classic MF algorithm by adding a new user or a new article to the matrices with a random initialization [1], and others apply meta-learning which aims to train a model that can rapidly adapt to a new task with a few examples [16], but do not solve the problem fundamentally.

Second, some news recommendation systems use clicked articles to represent a user, which can be adaptive to anonymous users. Some of them encode the article text with fine-grained attention mechanism [37, 44, 45, 61], some consider the relation between the dwell time of the user and satisfaction of the user [46], and others use the knowledge graph of entities in the news title as affiliated information [38, 39]. They mainly focus on the textual feature of articles in order to aggregate users’ preference while paying less attention to the click behavior. Although they can be applied for anonymous users by replacing long-term history clicks with articles within the session when fetching user representations, challenges are that they cannot take full advantage of the textual information due to the limited interactions and the overload of training cannot be avoided. Besides, they evaluate their methods by classifying the real clicked article and several constructed distractors from the impression list, and this is not consistent with the

real recommendation scenario, where the recommender retrieves top-K recommendation lists from all candidates.

For the rest of the work, one uses the information of how frequent user returns to help improve recommendation [59], another work jointly models click, unclick and dislike as explicit/implicit feedback [52], and others excavate the quality of news articles [21] or the backtracking behavior as the user feedback [30].

5.2 Session-based News Recommendation

Many online recommender systems are proposed to deal with the session-based scenarios [3, 60], where the user interaction information is limited and items are increasingly generated. Usually session-based news recommendation approaches integrate content-based similarity [32], and many of them introduce external knowledge to recommend top-K similar articles [28, 29, 34]. Some recommenders consider the past sessions of the same user [56, 57], which is not consistent with our anonymous settings, and that is why we do not compare experiment results with them.

Many other session-based recommenders are in the e-commerce domain, which can also be converted to deal with news articles. Here RNN, LSTM and GNN possess properties that make them attractive for sequence modeling of user sessions [4, 9, 11, 42, 50]. Further, a hybrid encoder with an attention mechanism is introduced to model the sequential behavior of users [18, 20, 31, 55, 58]. Besides, many sequential recommendation systems [25, 53] on music listening, games playing construct assorted RNN-related architectures (e.g. RCNN [54], GRU [10], HGN [22, 51]), showing RNN's high capacity to modeling user shift preference.

Although above works naturally take the content information and preference shifting into account, the implicit user feedback are neglected. When sampling negative articles, an adaptive negative sampling method based on GAN is proposed [41]. Beyond that, few works pay attention to the implicit meaning of negative samples. Randomly sampling from such continuously increasing and high-volume news articles might be fast but will not be effective enough.

5.3 The Use of Temporal Information

Sequence and Time-Aware Neighborhood (STAN) [5] takes vanilla SKNN as its special case. They build static time decay functions for three factors: the position of an item in the current session, recency of a past session w.r.t. to the current session, and the position of a recommendable item in a neighboring session. This approach can be regarded as rule-based SKNN, with exponential decay function, and the experiment result on e-commerce websites even outperforms some deep-learning-based approaches. In the deep learning model, some works design different temporal kernel functions or decay functions for different consumption scenarios [36, 48, 57]. However, these functions of news articles is fixed, which may undermine the ability to model user's short-term preferences towards different articles. Dwell time is considered in [46] as the user satisfaction, but the difference of users' reading speed is hard to capture in our session-based scenario. A time-interval-based GRU is proposed to model user session-level representations [17], and some work [26, 49, 55] treat the time feature of interactions as a temporal context, while they fail to consider the publishing/click/active time in the different dimension.

6 CONCLUSION

Session-based recommendations are indispensable under the streaming like or real-time scenario when users' historical records are unavailable. By leveraging the positive and negative implicit feedback from users, as well as properly modeling the times in the problem, our proposed method is simple but effective to improve the trade-off between accuracy, diversity and serendipity, as shown in experimental results. For further investigation, our positive/negative modules can be plugged into other sophisticated session-based recommendation approaches; the automatic diversity metric may not always accord with the user experience, and attention can be paid towards the real user satisfaction; the temporal encoder can encode the physical meaning of the date-time, so maybe the pre-trained temporal embedding can improve the model's understanding of tasks which contain temporal information.

REFERENCES

- [1] Marie Al-Ghossein, Pierre-Alexandre Murena, Talel Abdesslem, Anthony Barré, and Antoine Cornuéjols. 2018. Adaptive collaborative topic modeling for online recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 338–346.
- [2] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishu Aradhya, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*. 7–10.
- [3] Elena Viorica Epure, Benjamin Kille, Jon Espen Ingvaldsen, Rebecca Deneckere, Camille Salinesi, and Sahin Albayrak. 2017. Recommending personalized news in short user sessions. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*. 121–129.
- [4] P Moreira Gabriel De Souza, Dietmar Jannach, and Adilson Marques Da Cunha. 2019. Contextual hybrid session-based news recommendation with recurrent neural networks. *IEEE Access* 7 (2019), 169185–169203.
- [5] Diksha Garg, Priyanka Gupta, Pankaj Malhotra, Lovekesh Vig, and Gautam Shroff. 2019. Sequence and time aware neighborhood for session-based recommendations: Stan. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1069–1072.
- [6] Suyu Ge, Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020. Graph enhanced representation learning for news recommendation. In *Proceedings of The Web Conference 2020*. 2863–2869.
- [7] Jon Atle Gulla, Lemei Zhang, Peng Liu, Özlem Özgöbek, and Xiaomeng Su. 2017. The adressa dataset for news recommendation. In *Proceedings of the international conference on web intelligence*. 1042–1048.
- [8] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: a factorization-machine based neural network for CTR prediction. *arXiv preprint arXiv:1703.04247* (2017).
- [9] Lei Guo, Hongzhi Yin, Qinyong Wang, Tong Chen, Alexander Zhou, and Nguyen Quoc Viet Hung. 2019. Streaming session-based recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1569–1577.
- [10] Balázs Hidasi and Alexandros Karatzoglou. 2018. Recurrent neural networks with top-k gains for session-based recommendations. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 843–852.
- [11] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [12] Linmei Hu, Siyong Xu, Chen Li, Cheng Yang, Chuan Shi, Nan Duan, Xing Xie, and Ming Zhou. 2020. Graph neural news recommendation with unsupervised preference disentanglement. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 4255–4264.
- [13] Michael Jugovac, Dietmar Jannach, and Mozghan Karimi. 2018. Streamingrec: a framework for benchmarking stream-based news recommenders. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 269–273.
- [14] Marius Kaminskis and Derek Bridge. 2014. Measuring surprise in recommender systems. In *Proceedings of the workshop on recommender systems evaluation: dimensions and design (Workshop programme of the 8th ACM conference on recommender systems)*. Citeseer.
- [15] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.
- [16] Hoyeop Lee, Jinbae Im, Seongwon Jang, Hyunsouk Cho, and Sehee Chung. 2019. Melu: Meta-learned user preference estimator for cold-start recommendation.

- In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1073–1082.
- [17] Chenyi Lei, Shouling Ji, and Zhao Li. 2019. Tissa: A time slice self-attention approach for modeling sequential user behaviors. In *The World Wide Web Conference*. 2964–2970.
 - [18] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural attentive session-based recommendation. In *Proceedings of the 2017 ACM Conference on Information and Knowledge Management*. 1419–1428.
 - [19] Zibo Lin, Deng Cai, Yan Wang, Xiaojiang Liu, Haitao Zheng, and Shuming Shi. 2020. The World Is Not Binary: Learning to Rank with Grayscale Data for Dialogue Response Selection. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 9220–9229.
 - [20] Qiao Liu, Yifu Zeng, Refuoc Mokhosi, and Haibin Zhang. 2018. STAMP: short-term attention/memory priority model for session-based recommendation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1831–1839.
 - [21] Hongyu Lu, Min Zhang, Weizhi Ma, Yunqiu Shao, Yiqun Liu, and Shaoping Ma. 2019. Quality effects on user preferences and behaviors in mobile news streaming. In *The World Wide Web Conference*. 1187–1197.
 - [22] Chen Ma, Peng Kang, and Xue Liu. 2019. Hierarchical gating networks for sequential recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 825–833.
 - [23] Gabriel de Souza P. Moreira, Felipe Ferreira, and Adilson Marques da Cunha. 2018. News Session-Based Recommendations using Deep Neural Networks. *Proceedings of the 3rd Workshop on Deep Learning for Recommender Systems - DLRS 2018* (2018), 15–23. arXiv: 1808.00076.
 - [24] Zhiqiang Pan, Fei Cai, Wanyu Chen, Honghui Chen, and Maarten de Rijke. 2020. Star Graph Neural Networks for Session-based Recommendation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 1195–1204.
 - [25] Bruno L Pereira, Alberto Ueda, Gustavo Penha, Rodrygo LT Santos, and Nivio Ziviani. 2019. Online learning to rank for sequential music recommendation. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 237–245.
 - [26] Lakshmanan Rakkappan and Vaibhav Rajan. 2019. Context-Aware Sequential Recommendations with Stacked Recurrent Neural Networks. In *The World Wide Web Conference*. 3172–3178.
 - [27] Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. 2019. A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*. PMLR, 5628–5637.
 - [28] Heng-Shiou Sheu, Zhixuan Chu, Daiqing Qi, and Sheng Li. 2021. Knowledge-guided article embedding refinement for session-based news recommendation. *IEEE Transactions on Neural Networks and Learning Systems* (2021).
 - [29] Heng-Shiou Sheu and Sheng Li. 2020. Context-aware graph embedding for session-based news recommendation. In *Fourteenth ACM conference on recommender systems*. 657–662.
 - [30] Uzi Smadja, Max Grusky, Yoav Artzi, and Mor Naaman. 2019. Understanding reader backtracking behavior in online news articles. In *The World Wide Web Conference*. 3237–3243.
 - [31] Jing Song, Hong Shen, Zijing Ou, Junyi Zhang, Teng Xiao, and Shangsong Liang. 2019. ISLF: Interest Shift and Latent Factors Combination Model for Session-based Recommendation. In *IJCAI*. 5765–5771.
 - [32] Gabriele Sottocornola, Panagiotis Symeonidis, and Markus Zanker. 2018. Session-based news recommendations. In *Companion Proceedings of the The Web Conference 2018*. 1395–1399.
 - [33] Panagiotis Symeonidis, Lidija Kirjackaja, and Markus Zanker. 2020. Session-aware news recommendations using random walks on time-evolving heterogeneous information networks. *User Modeling and User-Adapted Interaction* (2020), 1–29.
 - [34] Panagiotis Symeonidis, Lidija Kirjackaja, and Markus Zanker. 2021. Session-based news recommendations using SimRank on multi-modal graphs. *Expert Systems with Applications* 180 (2021), 115028.
 - [35] 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*. 5998–6008.
 - [36] Chenyang Wang, Min Zhang, Weizhi Ma, Yiqun Liu, and Shaoping Ma. 2020. Make it a chorus: knowledge-and time-aware item modeling for sequential recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 109–118.
 - [37] Heyuan Wang, Fangzhao Wu, Zheng Liu, and Xing Xie. 2020. Fine-grained Interest Matching for Neural News Recommendation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 836–845.
 - [38] Hongwei Wang, Fuzheng Zhang, Jialin Wang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. 2018. RippletNet: Propagating user preferences on the knowledge graph for recommender systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 417–426.
 - [39] Hongwei Wang, Fuzheng Zhang, Xing Xie, and Minyi Guo. 2018. DKN: Deep knowledge-aware network for news recommendation. In *Proceedings of the 2018 world wide web conference*. 1835–1844.
 - [40] Menghan Wang, Mingming Gong, Xiaolin Zheng, and Kun Zhang. 2018. Modeling dynamic missingness of implicit feedback for recommendation. *Advances in neural information processing systems* 31 (2018), 6669.
 - [41] Qinyong Wang, Hongzhi Yin, Zhiting Hu, Defu Lian, Hao Wang, and Zi Huang. 2018. Neural memory streaming recommender networks with adversarial training. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2467–2475.
 - [42] Shoujin Wang, Liang Hu, Yang Wang, Quan Z Sheng, Mehmet Orgun, and Longbing Cao. 2019. Modeling multi-purpose sessions for nextitem recommendations via mixture-channel purpose routing networks. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. AAAI Press, 1–7.
 - [43] Wenjie Wang, Fuli Feng, Xiangnan He, Hanwang Zhang, and Tat-Seng Chua. 2020. “Click” Is Not Equal to “Like”: Counterfactual Recommendation for Mitigating Clickbait Issue. *arXiv preprint arXiv:2009.09945* (2020).
 - [44] Chuhan Wu, Fangzhao Wu, Mingxiao An, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019. Npa: Neural news recommendation with personalized attention. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2576–2584.
 - [45] Chuhan Wu, Fangzhao Wu, Mingxiao An, Tao Qi, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019. Neural news recommendation with heterogeneous user behavior. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 4874–4883.
 - [46] Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2020. User Modeling with Click Preference and Reading Satisfaction for News Recommendation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*. 3023–3029.
 - [47] Fangzhao Wu, Ying Qiao, Jun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, et al. 2020. Mind: A large-scale dataset for news recommendation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 3597–3606.
 - [48] Jibang Wu, Renqin Cai, and Hongning Wang. 2020. Déjà vu: A Contextualized Temporal Attention Mechanism for Sequential Recommendation. In *Proceedings of The Web Conference 2020*. 2199–2209.
 - [49] Meng Wu, Ying Zhu, Qilian Yu, Bhargav Rajendra, Yunqi Zhao, Navid Aghdaie, and Kazi A Zaman. 2019. A recommender system for heterogeneous and time sensitive environment. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 210–218.
 - [50] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based recommendation with graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 346–353.
 - [51] Teng Xiao, Shangsong Liang, and Zaiqiao Meng. 2019. Hierarchical neural variational model for personalized sequential recommendation. In *The World Wide Web Conference*. 3377–3383.
 - [52] Ruobing Xie, Cheng Ling, Yalong Wang, Rui Wang, Feng Xia, and Leyu Lin. 2020. Deep Feedback Network for Recommendation. *Proceedings of IJCAI-PRICAI* (2020).
 - [53] Chengfeng Xu, Pengpeng Zhao, Yanchi Liu, Victor S Sheng, Jiajie Xu, Fuzhen Zhuang, Junhua Fang, and Xiaofang Zhou. 2019. Graph contextualized self-attention network for session-based recommendation. In *Proc. 28th Int. Joint Conf. Artif. Intell. (IJCAI)*. 3940–3946.
 - [54] Chengfeng Xu, Pengpeng Zhao, Yanchi Liu, Jiajie Xu, Victor S Sheng S. Sheng, Zhiming Cui, Xiaofang Zhou, and Hui Xiong. 2019. Recurrent convolutional neural network for sequential recommendation. In *The world wide web conference*. 3398–3404.
 - [55] Qidi Xu, Haocheng Xu, Weilong Chen, Chaojun Han, Haoyang Li, Wenxin Tan, Fumin Shen, and Heng Tao Shen. 2019. Time-aware Session Embedding for Click-Through-Rate Prediction. In *Proceedings of the 27th ACM International Conference on Multimedia*. 2617–2621.
 - [56] Lemei Zhang, Peng Liu, and Jon Atle Gulla. 2018. A deep joint network for session-based news recommendations with contextual augmentation. In *Proceedings of the 29th on Hypertext and Social Media*. 201–209.
 - [57] Lemei Zhang, Peng Liu, and Jon Atle Gulla. 2019. Dynamic attention-integrated neural network for session-based news recommendation. *Machine Learning* 108, 10 (2019), 1851–1875.
 - [58] Tingting Zhang, Pengpeng Zhao, Yanchi Liu, Victor S Sheng, Jiajie Xu, Deqing Wang, Guanfeng Liu, and Xiaofang Zhou. 2019. Feature-level Deeper Self-Attention Network for Sequential Recommendation. In *IJCAI*. 4320–4326.
 - [59] Guanjie Zheng, Fuzheng Zhang, Zihan Zheng, Yang Xiang, Nicholas Jing Yuan, Xing Xie, and Zhenhui Li. 2018. DRN: A deep reinforcement learning framework for news recommendation. In *Proceedings of the 2018 World Wide Web Conference*. 167–176.
 - [60] Fan Zhou, Zijing Wen, Kumpeng Zhang, Goce Trajcevski, and Ting Zhong. 2019. Variational session-based recommendation using normalizing flows. In *The World Wide Web Conference*. 3476–3475.
 - [61] Qiannan Zhu, Xiaofei Zhou, Zeliang Song, Jianlong Tan, and Li Guo. 2019. Dan: Deep attention neural network for news recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 5973–5980.