

Neural News Recommendation with Collaborative News Encoding and Structural User Encoding

Zhiming Mao^{1,2}, Xingshan Zeng³, Kam-Fai Wong^{1,2}

¹The Chinese University of Hong Kong, Hong Kong, China

²MoE Key Laboratory of High Confidence Software Technologies, China

³Huawei Noah’s Ark Lab, China

^{1,2}{zmmao, kfwong}@se.cuhk.edu.hk ³zeng.xingshan@huawei.com

Abstract

Automatic news recommendation has gained much attention from the academic community and industry. Recent studies reveal that the key to this task lies within the effective representation learning of both news and users. Existing works typically encode news title and content separately while neglecting their semantic interaction, which is inadequate for news text comprehension. Besides, previous models encode user browsing history without leveraging the structural correlation of user browsed news to reflect user interests explicitly. In this work, we propose a news recommendation framework consisting of collaborative news encoding (CNE) and structural user encoding (SUE) to enhance news and user representation learning. CNE equipped with bidirectional LSTMs encodes news title and content collaboratively with cross-selection and cross-attention modules to learn semantic-interactive news representations. SUE utilizes graph convolutional networks to extract cluster-structural features of user history, followed by intra-cluster and inter-cluster attention modules to learn hierarchical user interest representations. Experiment results on the MIND dataset validate the effectiveness of our model to improve the performance of news recommendation¹.

1 Introduction

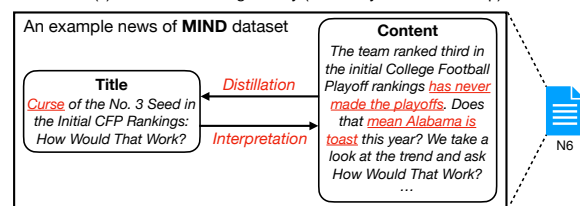
Online news applications, such as CNN News and MSN News, have become more and more people’s first choices to obtain the latest news (Das et al., 2007). With a deluge of news generated every day, an efficient news recommendation system should push relevant news to users to satisfy their diverse personalized interests (IJntema et al., 2010).

From the perspective of representation learning (Bengio et al., 2013), existing works mainly study how to effectively encode news and users into discriminative representations (Okura et al., 2017; Wu

¹Our code is released at <https://github.com/Veason-silverbullet/NNR>

News	Category	Title
N1	Travel	The Spookiest Place in Every State
N2	Travel	The Best Islands in the World: 2019 Readers’ Choice Awards
N3	Sports	Long Island high school football coach suspended for running up the score against previously undefeated foe
N4	Travel	50 states, 50 places: The top natural wonder in your state
N5	Sports	Potential Florida State football coaching targets
N6	Sports	Curse of the No. 3 Seed in the Initial CFP Rankings: How Would That Work?

(a) A user’s browsing history (sorted by click timestamp)



(b) Semantic interaction between news title and content

Figure 1: (a) An example of user browsing history. (b) An example of news title-content semantic interaction.

et al., 2019a; Wang et al., 2020). News encoders typically extract semantic representations of news from the textual spans (e.g., news title and content). User encoders are employed to learn the representation of a user from her browsing history. News recommendation models predict the matching probabilities between candidate news and users by measuring the similarity of their representations.

Existing news recommendation models typically encode news title and content separately and encode users’ browsing histories without explicit structural modeling. We argue that these encodings restrict the power of the news and user representations. To enhance news and user encoding, this work is established based on the two aspects of news and user representation learning:

(1) *Encoding the semantic interaction between news title and content*: Title and content play different roles in news, but they are complementary. News title distills the information of content, while content interprets the details of title, as shown in Figure 1(b). Previous works treat news title and content as two separate textual features, leading to a “semantic encoding dilemma”. This dilemma is dyadic as: (i) Although a news title is much shorter than content, the performance of title-encoding is

empirically better than content-encoding (Wu et al., 2020). This can be attributed to the crucial information that the human-summarized title naturally represents; (ii) News titles are always subjective and rhetorical to attract potential readers. This leads to a severe textual *data sparsity* problem. News titles with *unseen terminology*, *metaphor* and *ambiguity* make it difficult to comprehend news with limited title wording (Shree et al., 2019). For example in Figure 1(b), the word “curse” is a *metaphor*, which cannot be resolved by the training corpus or the title itself, due to its unique semantic occurrence. News encoders must turn to the content to interpret the semantics of the word “curse” (i.e., “has never made the playoffs” and “mean Alabama is toast”). However, news encoders proposed by previous works either extract features solely from the title, or encode title and content separately, then perform concatenation or attention fusion on them (Zhu et al., 2019; Wu et al., 2019a). Such separate encodings of title and content without leveraging their semantic interaction are inadequate for news text comprehension.

(2) *Encoding the user-interest-news correlation with hierarchical cluster-structure*: While a user usually has diverse interests in news topics, her browsed news with the same topic is often linked by some logical correlation. For example in Figure 1(a), the news N_3 , N_5 and N_6 are labeled as sports news and logically correlated with the topic “football”, forming a virtual user interest cluster. None of the single news can precisely represent the overall user interest in “football”. However, refined user interest in “football” can be encoded by aggregating the news N_3 , N_5 and N_6 . With aspects of user interests encoded within specific clusters, overall user representations can be aggregated by leveraging the correlation among interest clusters. Previous works typically formulate user history as an ordered linear sequence of news. Based on this sequential formulation, recurrent neural networks (Okura et al., 2017; An et al., 2019) and attention networks (Zhu et al., 2019; Wu et al., 2019a,b,c) are proposed to encode user history. These encoding methods viewing user history as a sequence of news cannot explicitly model the hierarchical *user-interest-news* correlation. Compared to linear sequences, hierarchical clusters are more suitable to represent a user’s diverse interests. User history can be structurally formulated into certain interest clusters, as correlated news shares information in a specific interest cluster. Encoding user history with

hierarchical cluster-structure is more precise to represent the correlation of news and user interests.

To address the above issues, in this work, we propose collaborative news encoding (CNE) and structural user encoding (SUE) to learn semantic-interactive news representations and hierarchical user representations. We conduct experiments on the MIND dataset (Wu et al., 2020), showing the encoding effectiveness of our proposed model. Experiments and further analyses validate that (i) CNE can enhance news encoding by exploiting the word-level semantic interaction between news title and content with cross-selective and cross-attentive mechanisms; (ii) SUE utilizes hierarchical cluster graphs to model the correlation of a user’s browsed news, which can extract more precise user interest representations; (iii) our model significantly outperforms existing state-of-the-art news recommendation models on the real-world MIND dataset.

2 Related Work

News recommendation is not only an important research task in NLP (Wu et al., 2020) but also a core component of industrial personalized news service (Okura et al., 2017). Conventional collaborative filtering (CF) approaches (Wang and Blei, 2011) exploit the interaction relationship between news and users. Since news only lasts for a short period, CF-based methods suffer from severe *cold-start* problem. To tackle this, content-based methods used handcrafted features to encode news and users (Li et al., 2010; Son et al., 2013; Bansal et al., 2015). In recent years, deep neural models have achieved superior performance in news recommendation. Many studies pinpointed that this improvement came from the fine-grained news and user representations, which were extracted by deep neural networks (Wu et al., 2019a,c; Wang et al., 2020).

For news representation learning, existing works used convolutional neural networks (CNN) (An et al., 2019), knowledge-aware CNN (Wang et al., 2018), personalized attention networks (Wu et al., 2019b), and multi-head self-attention networks (Wu et al., 2019c) to extract features from news title text as news representations. Zhu et al. (2019) employed parallel CNNs to encode news title and content respectively and then concatenated them into a unified representation. Wu et al. (2019a) encoded news title and content separately and incorporated them with multi-view attention.

For user representation learning, Okura et al.

(2017) used GRU to encode user history by ordered timestamp. An et al. (2019) utilized RNN to learn short-term user representations from the browsing history, combined with long-term user embeddings. Various attention networks are also widely used to attend to important news in user history (Wu et al., 2019a,b; Zhu et al., 2019). Wu et al. (2019c) employed multi-head self-attention (Vaswani et al., 2017) to capture deep interaction of user browsed news. These works formulated user history as an ordered linear sequence of news, to which recurrent or attention models were applied without modeling the structural correlation of user browsed news. Hu et al. (2020) formulated news and user jointly with a bipartite graph and disentangled user preferences with routing mechanism, which however implicitly relied on the manually-set latent preference factor.

3 Methodology

Our model is composed of the Collaborative News Encoding (CNE) module presented in Section 3.1 and Structural User Encoding (SUE) module presented in Section 3.2. CNE and SUE extract representations of candidate news and users respectively. The overall model architecture is illustrated in Figure 2. Finally, Section 3.3 will describe the click predictor and details of model training.

3.1 Collaborative News Encoding

3.1.1 Cross-selective Encoding

The news encoder is employed to learn semantic news representations from news title and content. Given the title word sequence $x^t = [x_1^t, x_2^t, \dots, x_N^t]$ and content word sequence $x^c = [x_1^c, x_2^c, \dots, x_M^c]$, they are mapped to word embeddings $W^t = [w_1^t, w_2^t, \dots, w_N^t]$ and $W^c = [w_1^c, w_2^c, \dots, w_M^c]$, where N and M are the word sequence lengths. For simplicity, we only formulate the title encoding part of our model (denoted by superscript t). The content encoding formula is symmetric as a counterpart (denoted by superscript c) and omitted.

First, two parallel bidirectional LSTMs are employed to extract the sequential features from the title and content word embeddings respectively.

$$\vec{h}_i^t = \overrightarrow{LSTM}^t(w_i^t, \vec{h}_{i-1}^t, \vec{c}_{i-1}^t) \quad (1)$$

$$\overleftarrow{h}_i^t = \overleftarrow{LSTM}^t(w_i^t, \overleftarrow{h}_{i+1}^t, \overleftarrow{c}_{i+1}^t) \quad (2)$$

where $\{h^t\}$ and $\{c^t\}$ are LSTM hidden states and cell states. The i -th title sequential feature is fused

as $h_i^t = [\vec{h}_i^t; \overleftarrow{h}_i^t]$, where $[\cdot; \cdot]$ denotes vector concatenation. We consider the global semantic information of title (content) preserved in its LSTM cell states and concatenate the last forward \vec{c}_N^t and backward \overleftarrow{c}_1^t as the semantic memory vector m^t .

$$m^t = [\vec{c}_N^t; \overleftarrow{c}_1^t] \quad (3)$$

To facilitate semantic interaction between title and content, we design a gated cross-selective network, inspired by Geng et al. (2020). Concretely, we utilize the semantic memory vector $m^{c(t)}$ to perform feature recalibration (Hu et al., 2018) on the sequential features $\{h^{t(c)}\}$ by a sigmoid gate function. The motivation behind this gate function is to utilize the memory vector of content (title) $m^{c(t)}$ to cross-select important semantic information from the i -th title (content) sequential feature $h_i^{t(c)}$.

$$Gate_i^t = \sigma(W_g^h h_i^t + W_g^m m^c + b_g) \quad (4)$$

$$\tilde{h}_i^t = Gate_i^t \odot h_i^t \quad (5)$$

where σ is sigmoid activation, \odot denotes element-wise multiplication. \tilde{h}_i^t is the cross-selective feature of the i -th title sequential feature h_i^t interacting with the content memory vector m^c . It is the first stage of collaborative title-content semantic interaction.

3.1.2 Cross-attentive Encoding

Based on the cross-selective sequential feature $\{\tilde{h}\}$, a two-phase attention module is designed to learn cross-attentive representations of title and content. First, we employ self-attention layers to learn the self-attentive representation of the sequential $\{\tilde{h}\}$.

$$\alpha_{self}^t = softmax(v^T tanh(W\tilde{h}^t + b))$$

$$r_{self}^t = \sum_{i=1}^N \alpha_{self,i}^t \tilde{h}_i^t \quad (6)$$

Then we employ the self-attentive representation r_{self} as a query, and the $\{\tilde{h}\}$ as key-value pairs to build cross-attention layers². It is the second stage of collaborative title-content semantic interaction.

$$\alpha_{cross}^t = Attention(r_{self}^c, \{\tilde{h}^t\})$$

$$r_{cross}^t = \sum_{i=1}^N \alpha_{cross,i}^t \tilde{h}_i^t \quad (7)$$

²Practically, we employ the scaled dot-product attention proposed by Vaswani et al. (2017). $Attention(Q, K) = softmax(\frac{\tilde{Q}\tilde{K}^T}{\sqrt{d}})$, where $\tilde{Q} = QW^Q$, $\tilde{K} = KW^K$. The same attention functions are also applied in Eq. (10) and (12).

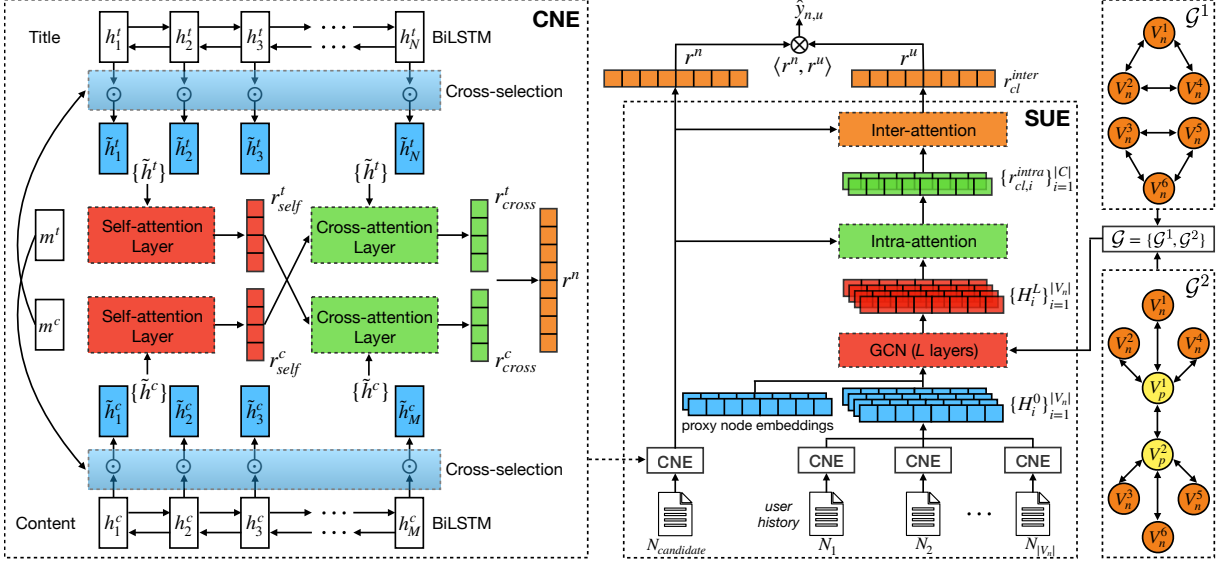


Figure 2: The overall architecture of our model. The graph construction is based on the user history in Figure 1(a).

We compute element-wise summation (denoted as \oplus) of the self-attentive representation r_{self} and cross-attentive representation r_{cross} to derive the title and content semantic-interactive representations, i.e., r^t and r^c . Finally, we concatenate r^t and r^c as the collaborative news representation r^n .

$$r^n = [r^t; r^c] = [r_{self}^t \oplus r_{cross}^t; r_{self}^c \oplus r_{cross}^c] \quad (8)$$

3.2 Structural User Encoding

3.2.1 Cluster-based Encoding of User History

The user encoder is employed to learn user interest representations from their browsing histories. To formulate the cluster-structure of user interests, we construct a hierarchical cluster graph in two steps:

Intra-cluster Subgraph \mathcal{G}^1 . We construct an original cluster graph³ with the topic category label of news (e.g., “Sports” and “Travel” in Figure 1(a)). We build the subgraph $\mathcal{G}^1 = (V_n, E_n)$ by treating the browsed news as nodes $\{V_n\}$ and adding bidirectional edges $\{E_n\}$ to those nodes, which share the same category labels. Each news node of $\{V_n\}$ is associated with its embedding r^n in Eq. (8). Each cluster contains multiple browsed news with a specific topic, reflecting an aspect of user interests.

Inter-cluster Subgraph \mathcal{G}^2 . Besides intra-cluster refinement of user interests, modeling inter-cluster correlation is also essential to leverage the overall information of user history. For each cluster C_i in \mathcal{G}^1 , we add a new cluster proxy node V_p^i . We build the subgraph $\mathcal{G}^2 = (\{V_n, V_p\}, \{E_p^1, E_p^2\})$ by adding bidirectional edges $\{E_p^1\}$ to those news

nodes $\{V_n\}$ and proxy nodes $\{V_p\}$ within the same clusters and fully connecting $\{V_p\}$ by bidirectional edges $\{E_p^2\}$. The node embedding of $\{V_p\}$ is initiated as zero-embedding r^p . News node information among clusters aggregates via cluster proxy nodes.

The hierarchical cluster graph \mathcal{G} consists of intra- and inter-cluster subgraphs: $\mathcal{G} = \{\mathcal{G}^1, \mathcal{G}^2\}$. With d -dimensional node embedding vectors $\{r_i^n\}_{i=1}^{|V_n|}$ and $\{r_i^p\}_{i=1}^{|V_p|}$, we define the history feature matrix as $H^0 = [r^n; r^p] \in \mathbb{R}^{(|V_n|+|V_p|) \times d}$. For graph \mathcal{G} , we denote its normalized adjacency matrix as \tilde{A} and degree matrix as \tilde{D} . We use graph convolutional networks (GCN) (Kipf and Welling, 2017) to extract structural representations of user history. To mitigate the over-smoothing issue of deep GCN (Chen et al., 2020), we add residual connections to adjacent GCN layers, following Li et al. (2019).

$$H^{l+1} = ReLU(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^l W^l) + H^l \quad (9)$$

where W^l is a trainable matrix. The GCN extracts structural features on graph \mathcal{G} , refining specific user interest representations within clusters and aggregating overall user history information among clusters. We train GCN of L layers and derive the structural user history representation from the news node embeddings as $r^h = \{H_i^L\}_{i=1}^{|V_n|} \in \mathbb{R}^{|V_n| \times d}$.

3.2.2 Intra-cluster Attention

Given the $|C|$ interest clusters implied by $|V_n|$ user’s browsed news, there are $|C_i|$ news in cluster C_i . The structural user history representation r^h can be viewed as $r^h = \{r_i^h\}_{i=1}^{|C|} = \left\{ \{r_{i,j}^h\}_{j=1}^{|C_i|} \right\}_{i=1}^{|C|}$.

³https://en.wikipedia.org/wiki/Cluster_graph

To derive intra-cluster features associated with candidate news, we design an intra-cluster attention layer, regarding the candidate news representation r_{can}^n as a query, and the j -th intra-cluster feature $r_{i,j}^h$ of cluster C_i as a key-value pair.

$$\alpha_i^{intra} = \text{Attention}(r_{can}^n, \{r_i^h\})$$

$$r_{cl,i}^{intra} = \sum_{j=1}^{|C_i|} \alpha_{i,j}^{intra} r_{i,j}^h \quad (10)$$

The intra-cluster feature $r_{cl,i}^{intra}$ attends to the node-level features $\{r_{i,j}^h\}_{j=1}^{|C_i|}$ of cluster C_i , associated with the candidate news representation r_{can}^n . The $r_{cl,i}^{intra}$ refines the i -th user interest representation within the cluster C_i of graph \mathcal{G} .

3.2.3 Inter-cluster Attention

Before inter-cluster modeling, a nonlinear transformation is performed to project the $r_{cl,i}^{intra}$, which is originally a linear combination of node-level features in cluster C_i , into cluster-level feature spaces.

$$\tilde{r}_{cl,i}^{intra} = \text{ReLU}(\tilde{W}r_{cl,i}^{intra} + \tilde{b}) + r_{cl,i}^{intra} \quad (11)$$

To derive inter-cluster features associated with candidate news, we design an inter-cluster attention layer, regarding the candidate news representation r_{can}^n as a query, and the i -th intra-cluster feature $\tilde{r}_{cl,i}^{intra}$ of graph \mathcal{G} as a key-value pair.

$$\alpha^{inter} = \text{Attention}(r_{can}^n, \{\tilde{r}_{cl}^{intra}\})$$

$$r_{cl}^{inter} = \sum_{i=1}^{|C|} \alpha_i^{inter} \tilde{r}_{cl,i}^{intra} \quad (12)$$

The inter-cluster feature r_{cl}^{inter} attends to the cluster-level features $\{\tilde{r}_{cl,i}^{intra}\}_{i=1}^{|C|}$ of graph \mathcal{G} , associated with the candidate news representation r_{can}^n . With intra-cluster and inter-cluster attention, r_{cl}^{inter} hierarchically aggregates user interest representations within the cluster graph \mathcal{G} . r_{cl}^{inter} is adopted as the user representation r^u , i.e., $r^u = r_{cl}^{inter}$.

3.3 Click Predictor and Model Training

Given the news and user representations r^n and r^u , the click predictor is employed to predict the probability that user u clicks on news n . Motivated by the previous works (Wang et al., 2018; Wu et al., 2019a), we compute the dot-product $\hat{y}_{n,u}$ of r^n and r^u , i.e., $\hat{y}_{n,u} = \langle r^n, r^u \rangle$, as the unnormalized matching score of news n and user u .

Following common practice of previous works (Huang et al., 2013; Wang et al., 2020), we employ negative sampling strategy to model training. For

# users	200,000	# users in train set	189,532
# news	78,520	# news in train set	75,858
# training logs	594,433	# positive samples	902,330
Avg. title len	11.67	Avg. content len	41.01

Table 1: Statistics of the 200K-MIND dataset.

each user click-impression, i.e., the i -th impression log that user u had clicked on news n , we compute its unnormalized matching score as \hat{y}_i^+ . Besides, we randomly sample K pieces of news, which are not clicked by the user u . Unnormalized matching scores $\hat{y}_{i,j}^-$ are computed for these K negative samples, where $j = 1, \dots, K$. By such means, it can be reformulated as a $(K + 1)$ -way classification problem. We employ softmax function to derive the normalized matching probabilities and sum up the negative log-likelihood of positive samples over the training dataset \mathcal{D} , as model training loss \mathcal{L} .

$$\mathcal{L} = - \sum_{i=1}^{|\mathcal{D}|} \log \frac{\exp(\hat{y}_i^+)}{\exp(\hat{y}_i^+) + \sum_{j=1}^K \exp(\hat{y}_{i,j}^-)} \quad (13)$$

4 Experiment Setup

4.1 Dataset and Experiment Settings

We conduct experiments on the MIND dataset (Wu et al., 2020). MIND is a large-scale English news recommendation dataset built from real-world MSN news and anonymized user click logs⁴. Since the MIND is quite large-scale, following existing works (Wu et al., 2019a,b,c; Wang et al., 2020)⁵, we randomly sample 200K users' click logs out of 1 million users from the user behavior logs of MIND training and validation sets. Since the MIND test set is not labeled, we half-split the original validation set into experimental validation and test sets. We employ the abstract column texts in MIND as the news content texts⁶. Detailed statistics of the 200K-MIND dataset are shown in Table 1.

We truncate news title and content with the maximum length of 32 and 128 respectively. The number of news in user browsing history is capped at 50. Following (An et al., 2019; Wu et al., 2019a,b,c), we perform negative sampling with the sampling ratio K of 4 (see Section 3.3). The word embedding is initialized from the pretrained 300-dimensional

⁴<https://msnews.github.io>

⁵These works used the MSN news dataset with 10K sampled users, as training on the full MIND dataset with 1 million users is very expensive in GPU time.

⁶Detailed MIND dataset format at <https://github.com/msnews/msnews.github.io/blob/master/assets/doc/introduction.md>

Glove embedding (Pennington et al., 2014). The number of GCN layers in SUE is set as $L = 4$ (investigated in Section 5.4). We use Adam optimizer (Kingma and Ba, 2015) with the learning rate of $1e-4$ to train our model with the dropout rate of 0.2. The area under the ROC curve (AUC), mean reciprocal rank (MRR), and normalized discounted cumulative gain (nDCG@5 and nDCG@10) are adopted as ranking metrics to evaluate model performance. We set the batch size to 64 and conduct early stopping if the validation AUC score had not improved over 5 epochs. We independently repeat each experiment for 10 times and report the average performance scores.

4.2 Comparison Methods

We compare our model with state-of-the-art general and news-specific recommendation methods.

General Recommendation Methods. General methods utilize handcrafted features to learn news and user representations. We use the TF-IDF features extracted from news and user history with the one-hot news and user IDs as input features for the experiments. The general methods include (1) **LibFM** (Rendle, 2012), a factorization machine estimating the sparse feature interaction between news and users; (2) **DSSM** (Huang et al., 2013), a deep structured semantic model, regarding the user history as a query and candidate news as key documents; (3) **Wide&Deep** (Cheng et al., 2016), a framework consisting of wide channels with a linear model and deep channels with a neural model.

Neural News Recommendation Methods. We compete with the state-of-the-art neural models, which are specifically designed for news recommendation, including (1) **DAE-GRU** (Okura et al., 2017), encoding news with a denoising autoencoder and users with a gated recurrent network; (2) **DFM** (Lian et al., 2018), a deep fusion model using multi-channel inception blocks to capture various interaction among news features; (3) **DKN** (Wang et al., 2018), utilizing knowledge-aware CNNs to fuse knowledge encoding and textual encoding of news title; (4) **LSTUR** (An et al., 2019), encoding news title with a CNN network, while jointly modeling long-term user preferences and short-term user interests with a GRU network; (5) **NAML** (Wu et al., 2019a), utilizing CNN networks to encode title and content texts, while encoding the category and subcategory topics with dense layers. The text and topic representations are incorporated by multi-

Methods	AUC	MRR	nDCG@5	nDCG@10
LibFM	61.16	27.88	30.06	36.44
DSSM	64.74	30.12	33.22	39.50
Wide&Deep	64.62	29.87	32.71	39.11
DAE-GRU	65.98	31.48	34.93	41.12
DFM	64.63	29.80	32.82	39.29
DKN	66.20	31.25	34.23	40.92
LSTUR	68.10	32.87	36.46	42.69
NAML	68.63 [†]	33.16	36.79	43.07
NPA	67.34	32.59	35.98	42.28
NRMS	68.61	33.46 [†]	37.02 [†]	43.30 [†]
FIM	68.44	32.95	36.58	42.97
CNE-SUE	69.55*	33.70*	37.54*	43.79*
CNE w/o CS	69.39	33.52	37.30	43.62
CNE w/o CA	69.48	33.61	37.39	43.68
SUE w/o GCN	69.31	33.48	37.25	43.53
SUE w/o HCA	69.40	33.52	37.37	43.65

Table 2: Performance comparison results ([†] denotes the highest baseline scores, * denotes that the performance improvements over all baseline methods are validated by Student’s unpaired t-test with p -value < 0.01).

view attention. NAML uses an attention network as its user encoder; (6) **NPA** (Wu et al., 2019b), attending to important words and news articles by personalized attention networks built with user embeddings; (7) **NRMS** (Wu et al., 2019c), utilizing effective multi-head self-attention networks (Vaswani et al., 2017) to extract fine-grained representations from the news title and user history respectively; (8) **FIM** (Wang et al., 2020), encoding news titles with dilated convolution networks and modeling the interaction between candidate news and user history with 3D convolutional matching networks.

Variants of Our Model. To further verify the efficacy of our model design, we also experiment with the ablation variants of our model by respectively removing the cross-selection module (**CNE w/o CS**), cross-attention module (**CNE w/o CA**), GCN layers (**SUE w/o GCN**), and hierarchical cluster attention module (**SUE w/o HCA**).

5 Experiment Results and Analyses

5.1 Main Comparison Results

Table 2 shows the performance comparison results. Our model CNE-SUE achieves the highest performance consistently in all evaluation metrics. Detailed observations can be obtained as follows.

General recommendation methods yield much lower performance than most neural news recommendation methods. This is due to that deep neural models can learn refined representations adaptively, which are more effective than general feature engi-

News Enc.	AUC	MRR	nDCG@5	nDCG@10
CNN	68.31	33.07	36.65	42.92
KCNN	65.27	30.72	33.67	40.28
Per-CNN	68.37	33.12	36.68	42.90
MHSA	68.26	33.07	36.59	42.83
NAML	68.63 [†]	33.16 [†]	36.79 [†]	43.07 [†]
CNE	69.21	33.32	37.16	43.44
NAML-T	68.49	33.21	36.89	43.11
NAML-C	67.22	32.08	35.62	41.76
CNE-T	68.25	33.18	36.80	43.12
CNE-C	67.96	32.65	36.34	42.44

Table 3: Ablation study of news encoders (KCNN denotes knowledge-aware CNN in DKN, Per-CNN denotes CNN with personalized attention in NPA, MHSA denotes multi-head self-attention networks in NRMS, *-T(C) denotes title (content) encoding only).

neering with fixed handcrafted features.

In all evaluation metrics, our model CNE-SUE outperforms all comparison methods by significant margins (+0.92% AUC, +0.24% MRR, +0.52% nDCG@5, and +0.49% nDCG@10 compared to the best baseline performance). This performance improvement derives from the collaborative news representations and structural user representations extracted by our model. Specifically, CNE can extract more accurate news semantics by leveraging the title-content semantic interaction, compared to title-encoding (e.g., LSTUR, NRMS, and FIM) and separate encodings of title and content (e.g., NAML). SUE modeling diverse user interests with hierarchical cluster structure is more powerful than the comparison methods formulating user history as a linear sequence of news, which employ recurrent neural networks (e.g., LSTUR) or attention networks (e.g., NPA, NAML, and NRMS).

From table 2, we can observe varying degrees of performance decreases on the ablation variants compared to our full model. It suggests the usefulness of different components in our model. CNE w/o CA performs the best among all variants. This is because the news representations learned by CNE are composed of self- and cross-attentive representations (refer to Eq. (8)), and the remaining self-attention can achieve suboptimal performance. Removing GCN layers leads to the most significant impact on performance, indicating the efficacy of structural modeling on user history.

5.2 Effectiveness of Collaborative Encoding

We conduct ablation experiments on news encoders. For fair comparison and excluding the influence of

User Enc.	AUC	MRR	nDCG@5	nDCG@10
LSTUR	68.10	32.87	36.46	42.69
ATT	68.31	33.07	36.65	42.92
Per-ATT	66.97	32.22	35.48	41.82
Can-ATT	68.48	33.23	36.73	43.04
MHSA	68.75 [†]	33.34 [†]	36.97 [†]	43.26 [†]
SUE	69.03	33.53	37.26	43.48

Table 4: Ablation study of user encoders (ATT denotes vanilla attention networks, Per-ATT denotes personalized attention in NPA, Can-ATT denotes candidate-aware attention in DKN, MHSA denotes multi-head self-attention networks in NRMS).

SUE, all ablation models apply the same attention user encoders. We also examine the title and content encodings. Table 3 shows the ablation results⁷.

From the ablation results, we can observe that CNE significantly outperforms other existing news encoding methods. NAML is also competitive, as it can incorporate informative representations from news texts and topic categories. Table 3 also shows that title-encoding (*-T) is much more effective than content-encoding (*-C), though content texts are theoretically more informative. This confirms the “*semantic encoding dilemma*” in news encoding and may explain why many existing works (e.g., LSTUR, NPA, NRMS, and FIM) employ title-encoding only. Comparing NAML to NAML-T, there is no significant performance enhancement. In contrast, CNE achieves much higher scores than the individual title and content encodings (CNE-T and CNE-C). It validates the necessity of encoding news title and content with word-level semantic interaction to enhance news representation learning.

5.3 Effectiveness of Structural Encoding

We conduct ablation experiments on user encoders. For fair comparison and excluding the influence of CNE, all ablation models apply the same CNN title encoders. Table 4 shows the ablation results.

According to Table 4, MHSA performs much better than other baseline user encoders. This is because MHSA (Vaswani et al., 2017) can model the correlation of each pair of news in user history. It validates the necessity of modeling the correlation of historical news in user encoding. Moreover, Table 4 shows that SUE significantly outperforms MHSA. This is because the manner of encoding historical news correlation with hierarchical clusters in

⁷We did not include the ablation of FIM, because the FIM news encoder produces special hierarchical 3D-sized representations, which are incompatible with other ablation encoders.




Figure 3: The performance of our model on validation set with respect to the number of GCN layers (the trend of ndcg@10 is similar to ndcg@5 and hence omitted).

Curse of the No. 3 Seed in the initial CFP rankings How Would That Work ?

(a) Attention weights of our model over the title words.

The team ranked third in the initial College Football Playoff rankings has never made the playoffs . does that mean Alabama is toast this year ? we take a look at the trend and ask How Would That Work ?

(b) Attention weights of our model over the content words.

Figure 4: Attention weight visualization on the news N_6 (darker colors denote higher attention weights).

SUE is more fine-grained than MHSA. Concretely, modeling intra-cluster news interaction is more effective to reflect aspects of user interests, while modeling inter-cluster user interests interaction is more effective to encode overall user representations. The ablation results indicate that structural modeling of the hierarchical *user-interest-news* correlation can effectively enhance user encoding.

5.4 Parameter Analysis

We investigate the influence of the number of GCN layers L in our model. Figure 3 shows the results. The model performance on validation set increases and reaches a peak, as L increases from 1 to 4. This is because equipped with deeper GCN, the model can capture more fine-grained information of user browsing behaviors by modeling higher-order interaction of browsed news. Nevertheless, as L continues to increase, the model performance begins to decline. This may be because deep GCN always suffers from the *over-smoothing* issue (Chen et al., 2020). As GCN becomes too deep, the user history representations r^h tend to be indistinguishable and impair the ultimate user representations. Herein, $L = 4$ is optimal for our model.




Figure 5: The AUC scores of different models with respect to the number of user interest clusters.

5.5 Case Study

We then probe into how our model processes news texts. According to Eq. (6), (7), and (8), we define the word attention weights of title (content) as $\alpha^{t(c)} = (\alpha_{self}^{t(c)} + \alpha_{cross}^{t(c)})/2$, where $\alpha^{t(c)} \in [0, 1]$. As shown in Figure 4, we visualize our model’s output title (content) attention weights $\alpha^{t(c)}$ over the title (content) words of the news N_6 in Figure 1.

From Figure 4(a), we observe that our model mainly attends to the words “curse” and “CFP”, which contain the core information of the news N_6 . It validates that our model can distill the most informative words from the news title. As the content visualization shown in Figure 4(b), our model mostly attends to the words “Alabama” and “toast”, which interpret the specific semantics of the word “curse” in the context of the news N_6 . Besides, our model also attends to the important contextual words, such as “initial”, “rankings”, and “trend”. These title and content attention weights indicate that our model can accurately encode the news N_6 .

5.6 Analysis on User Interest Modeling

We analyze how our model performs with different numbers of user interest clusters $|C|$ (refer to Section 3.2). The results are shown in Figure 5. When $|C| = 1$, our model slightly underperforms NAML. This is because our model is overfitted to represent single user interest with cluster graphs. All models’ performance increases with the growth of $|C|$. This is because the models can learn more precise user representations as more news information is incorporated. In cases of $|C| > 1$, our model consistently outperforms all baselines. It validates the usefulness of encoding user history with hierarchical cluster-structure in cases of modeling diverse

user interests. Moreover, the performance of all models decreases when $|C|$ becomes too large (i.e., $|C| > 8$). This reveals the challenge of predicting a user’s news-clicking behavior when her browsing history covers too many kinds of news topics.

6 Conclusion

In this work, we present a neural news recommendation model with collaborative news encoding and structural user encoding. CNE leverages the title-content semantic interaction to enhance news encoding. SUE exploits the correlation of browsed news and represents user interests with hierarchical cluster graphs to enhance user encoding. Experiment results show that our model achieves significant performance enhancement compared to the existing state-of-the-art methods. We also further analyze our model and validate its effectiveness.

Acknowledgements

We appreciate some insightful comments from the anonymous reviewers. The research described in this paper is partially supported by Hong Kong RGC-GRF #14204118 and Hong Kong RSFS #3133237.

References

- Mingxiao An, Fangzhao Wu, Chuhan Wu, Kun Zhang, Zheng Liu, and Xing Xie. 2019. [Neural news recommendation with long- and short-term user representations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 336–345, Florence, Italy. Association for Computational Linguistics.
- Trapit Bansal, Mrinal Das, and Chiranjib Bhat-tacharyya. 2015. [Content driven user profiling for comment-worthy recommendations of news and blog articles](#). In *Proceedings of the 9th ACM Conference on Recommender Systems*, page 195–202, New York, USA. Association for Computing Machinery.
- Y. Bengio, A. Courville, and P. Vincent. 2013. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828.
- Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. 2020. [Measuring and relieving the over-smoothing problem for graph neural networks from the topological view](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):3438–3445.
- Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishi Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, Rohan Anil, Zakaria Haque, Lichan Hong, Vihan Jain, Xiaobing Liu, and Hemal Shah. 2016. [Wide & deep learning for recommender systems](#). In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*, page 7–10, New York, NY, USA. Association for Computing Machinery.
- Abhinandan S. Das, Mayur Datar, Ashutosh Garg, and Shyam Rajaram. 2007. [Google news personalization: Scalable online collaborative filtering](#). In *Proceedings of the 16th International Conference on World Wide Web*, page 271–280, New York, NY, USA. Association for Computing Machinery.
- Xinwei Geng, Longyue Wang, Xing Wang, Bing Qin, Ting Liu, and Zhaopeng Tu. 2020. [How does selective mechanism improve self-attention networks?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2986–2995, Online. Association for Computational Linguistics.
- J. Hu, L. Shen, and G. Sun. 2018. [Squeeze-and-excitation networks](#). In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7132–7141.
- 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*, pages 4255–4264, Online. Association for Computational Linguistics.
- Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. [Learning deep structured semantic models for web search using clickthrough data](#). In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*, page 2333–2338, New York, NY, USA. Association for Computing Machinery.
- Wouter IJntema, Frank Goossen, Flavius Frasincar, and Frederik Hogenboom. 2010. [Ontology-based news recommendation](#). In *Proceedings of the 2010 EDBT/ICDT Workshops*, New York, NY, USA. Association for Computing Machinery.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Guohao Li, Matthias Müller, Ali Thabet, and Bernard Ghanem. 2019. Deepgcns: Can gcns go as deep as cnns? In *The IEEE International Conference on Computer Vision (ICCV)*.

- Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, page 661–670, New York, NY, USA. Association for Computing Machinery.
- Jianxun Lian, Fuzheng Zhang, Xing Xie, and Guangzhong Sun. 2018. Towards better representation learning for personalized news recommendation: a multi-channel deep fusion approach. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 3805–3811. International Joint Conferences on Artificial Intelligence Organization.
- Shumpei Okura, Yukihiro Tagami, Shingo Ono, and Akira Tajima. 2017. Embedding-based news recommendation for millions of users. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 1933–1942, New York, NY, USA. Association for Computing Machinery.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Steffen Rendle. 2012. Factorization machines with libFM. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3):57:1–57:22.
- Jaya Shree, Emily Liu, Andrew Gordon, and Jerry Hobbs. 2019. Deep natural language understanding of news text. In *Proceedings of the First Workshop on Narrative Understanding*, pages 19–27, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jeong-Woo Son, A-Yeong Kim, and Seong-Bae Park. 2013. A location-based news article recommendation with explicit localized semantic analysis. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '13*, page 293–302, New York, NY, USA. Association for Computing Machinery.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. Curran Associates, Inc.
- Chong Wang and David Blei. 2011. Collaborative topic modeling for recommending scientific articles. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 448–456.
- 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*, pages 836–845, Online. Association for Computational Linguistics.
- 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, WWW '18*, page 1835–1844, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Chuhan Wu, Fangzhao Wu, Mingxiao An, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019a. Neural news recommendation with attentive multi-view learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 3863–3869. International Joint Conferences on Artificial Intelligence Organization.
- Chuhan Wu, Fangzhao Wu, Mingxiao An, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019b. Npa: Neural news recommendation with personalized attention. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 2576–2584, New York, NY, USA. Association for Computing Machinery.
- Chuhan Wu, Fangzhao Wu, Suyu Ge, Tao Qi, Yongfeng Huang, and Xing Xie. 2019c. Neural news recommendation with multi-head self-attention. 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)*, pages 6389–6394, Hong Kong, China. Association for Computational Linguistics.
- Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, and Ming Zhou. 2020. MIND: A large-scale dataset for news recommendation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3597–3606, Online. Association for Computational Linguistics.
- Qiannan Zhu, Xiaofei Zhou, Zeliang Song, Jianlong Tan, and Li Guo. 2019. Dan: Deep attention neural network for news recommendation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):5973–5980.