

UniTRec: A Unified Text-to-Text Transformer and Joint Contrastive Learning Framework for Text-based Recommendation

Zhiming Mao^{1,2}, Huimin Wang^{1,3}, Yiming Du^{1,2}, Kam-Fai Wong^{1,2}

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

²MoE Key Laboratory of High Confidence Software Technologies, China

³Jarvis Lab, Tencent, Shenzhen, China

{zmmao, ydu, kfwong}@se.cuhk.edu.hk

hmmmwang@tencent.com

Abstract

Prior study has shown that pretrained language models (PLM) can boost the performance of text-based recommendation. In contrast to previous works that either use PLM to encode user history as a whole input text, or impose an additional aggregation network to fuse multi-turn history representations, we propose a unified local- and global-attention Transformer encoder to better model two-level contexts of user history. Moreover, conditioned on user history encoded by Transformer encoders, our framework leverages Transformer decoders to estimate the language perplexity of candidate text items, which can serve as a straightforward yet significant contrastive signal for user-item text matching. Based on this, our framework, UniTRec, unifies the contrastive objectives of discriminative matching scores and candidate text perplexity to jointly enhance text-based recommendation. Extensive evaluation shows that UniTRec delivers SOTA performance on three text-based recommendation tasks.¹

1 Introduction

Text-based recommendation (Li et al., 2010; Gu et al., 2016; Okura et al., 2017; Malkiel et al., 2020) aims to recommend relevant textual content (e.g., news articles, Twitter posts) to people based on their behaviors as represented in historical log texts. For instance, engagement recommendation (Cheng et al., 2022) on social media (e.g., Twitter and Reddit) helps users discover and engage with interested threads by modeling their browsing history.

Pretrained language models (Devlin et al., 2019; Brown et al., 2020) have made waves in recent text-based recommendation research (Zhang et al., 2021; Qi et al., 2022; Geng et al., 2022). The most common practice is using PLM encoders (BERT family) to learn representations of user history and candidate item texts. Recommendation

matching scores are computed over the user and item representations and finally optimized by noise contrastive estimation (NCE) loss (Gutmann and Hyvärinen, 2010) for ranking multiple candidates.

Unlike encoding single text, using PLM to encode multi-turn texts of user history is nontrivial. Existing works (Malkiel et al., 2020; Qi et al., 2022; Geng et al., 2022) concatenate multi-turn history texts as a whole input text, then use one PLM encoder to learn the holistic user representation. This is a standard PLM encoding manner but ignores the relation among history turns, as all word tokens from different history turns are *equally attended*². In contrast, previous studies point out that learning the relation among user history turns is also beneficial (Zeng et al., 2020; Qi et al., 2021). Another approach is using PLM encoders to learn representations from multi-turn history texts, followed by an additional aggregation network to fuse the multi-turn representations (Wu et al., 2021; Li et al., 2022). However, the imposed aggregation networks (with newly initialized parameters) weaken the representation power of PLM encoders which are already pretrained on large-scale corpora.

This work introduces UniTRec, a **Unified** text-to-text **Transformer** framework for text-based **Recommendation**. In the encoder component of UniTRec, we design local- and global-attention to learn user history representations through tailored attention masking, which aims to jointly model word-level and turn-level relations of user history. UniTRec can utilize the full power of PLM encoders because it preserves the intact structure of PLM encoders without newly imposed parameters.

Different from most previous works that predict user-candidate matching scores solely based on the representations learned by Transformer encoders, we argue that conditioned on user representations

¹Our code is available at <https://github.com/Veason-silverbulet/UniTRec>.

²There is no inductive bias of turn-level and history-level relations introduced to Transformer self-attention computation, where each token plays an equal role.

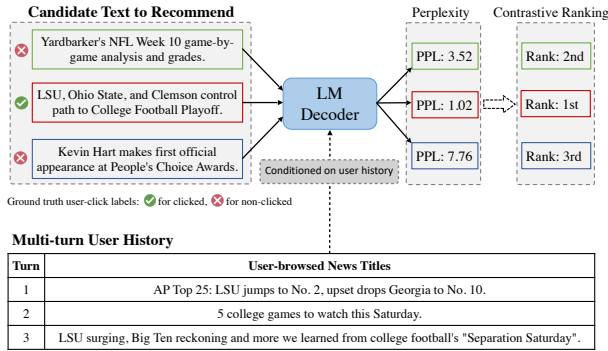


Figure 1: An example of perplexity-based ranking for candidate item texts, conditioned on user history. The illustrated task is text-based news recommendation.

learned by Transformer encoders, candidate text perplexity (PPL) estimated by pretrained Transformer decoders is also a straightforward yet significant signal for text-based recommendation. As shown in Figure 1, we hypothesize that the candidate text perplexity estimated by pretrained LM decoders can directly measure the text matching degree between user history and candidate texts. It is because the perplexity estimates the likelihood of candidate texts based on encoder outputs, which naturally indicates the probabilities of candidate texts given the user history. Besides, UniTRec can use the last hidden states of Transformer decoders to directly predict matching scores. Hence, this work unifies the contrastive objectives of discriminative matching scores and candidate text perplexity to jointly enhance text-based recommendation.

The contributions of this work are: (1) We propose local- and global-attention to model two-level relation of user history without additional parameters, which enjoys the full power of PLM encoders. (2) We introduce PLM perplexity to measure user-candidate text matching and unify the objectives of discriminative matching scores and candidate text perplexity to enhance text-based recommendation. (3) Experiments on three text-based recommendation datasets validate the effectiveness of UniTRec.

2 Approach

2.1 Unified User-history Modeling

Formally, multi-turn history of a user is represented as $H = [t_1, t_2, \dots, t_N]$, and each turn text t_i contains $|t_i|$ words as $t_i = [x_i^1, x_i^2, \dots, x_i^{|t_i|}]$. UniTRec aims to unify learning word- and turn-level context representations in one Transformer encoder.

Local attention on word-level context. We first

concatenate the multi-turn history texts as the input tokens $X = [x_1^1, x_1^2, \dots, x_1^{|t_1|}, \dots, x_N^1, x_N^2, \dots, x_N^{|t_N|}]$. Inspired by Dong et al. (2019), we tailor the attention masking in Transformer self-attention to learn the word-level context of each turn. Specifically, we allow word tokens from the same turn to attend to each other, while tokens from different turns are excluded from self-attention computation:

$$\mathbf{M}_{i,j} = \begin{cases} 0, & \text{token } x_i \text{ and } x_j \text{ in the same turn} \\ -\infty, & \text{otherwise} \end{cases}$$

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}} + \mathbf{M}\right)V \quad (1)$$

, where Q, K, V are self-attention query, key, and value in Vaswani et al. (2017), \mathbf{M} is the mask matrix to achieve local-attention inside each turn text. The local self-attention blocks consist of L_1 layers, by which original PLM encoders can be adapted to learn word-level context representations of turns.

Global attention on turn-level context. Over the local self-attention layers, we leverage global self-attention to model the relation among history turns. Specifically, tokens from all turns attend to each other in self-attention computation (by setting the mask matrix $\mathbf{M} = \mathbf{0}$). In this way, Transformer encoders can perform global interaction among each token (and turn) to learn turn-level context representations of user history. There are L_2 layers in the global self-attention blocks, which can also be inherited from PLM encoders directly.

2.2 Joint Contrastive Ranking Objectives

Conditioned on the history representation, we input the candidate text to Transformer decoders to predict how likely it should be recommended. It is worth noting that Transformer decoders can naturally perform effective **cross-attention** interaction between history and candidate hidden states.

2.2.1 Objective on Discriminative Scores

Motivated by Lewis et al. (2020), we feed the last hidden state of decoder output h_T to an MLP score-head which predicts the user-candidate matching score $S^d = \text{ScoreHead}(h_T)$. The matching score is discriminative, as higher scores indicate higher user-candidate matching probabilities.

Following previous works (Li et al., 2022; Qi et al., 2022), we adopt negative sampling with NCE loss to optimize matching score prediction. Given the user history and its ground truth matched candidate C_i , UniTRec predicts the matching score

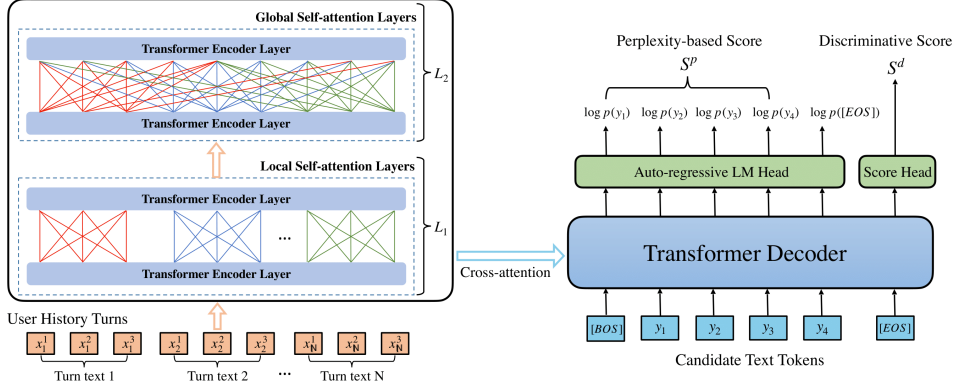


Figure 2: Overview of UniTRec. In training, matching scores S^d and S^p are optimized by the NCE loss, respectively. In inference, S^d and S^p are normalized and combined to derive the final output ranking.

as S_i^{d+} . In addition, K unmatched negative candidates $\{C_j\}_{j=1}^K$ are sampled from the candidate set, and their matching scores are $\{S_j^{d-}\}_{j=1}^K$. The NCE loss is represented in a contrastive form:

$$\mathcal{L}_i^d = -\log \frac{\exp(S_i^{d+})}{\exp(S_i^{d+}) + \sum_{j=1}^K \exp(S_j^{d-})} \quad (2)$$

2.2.2 Objective on Candidate Text Perplexity

As aforementioned, UniTRec leverages perplexity to rank candidate texts. Since lower perplexity indicates higher user-candidate matching probability, regarding the candidate text $Y = [y_1, y_2, \dots, y_T]$, we define the perplexity-based matching score S^p as its negative perplexity³:

$$S^p = -\text{PPL}(Y) = \frac{1}{T} \sum_{i=1}^T \log p_{\theta}(y_i | y_{<i}) \quad (3)$$

, where $p_{\theta}(\cdot)$ denotes the target probability output from the UniTRec Transformer decoder. Similar to Eq. (2), we optimize the perplexity-based matching score S^p in the NCE loss form. As perplexity empirically varies in a wide range, we introduce a temperature parameter τ to balance the joint NCE loss gradients following Radford et al. (2021).

$$\mathcal{L}_i^p = -\log \frac{\exp(\tau \cdot S_i^{p+})}{\exp(\tau \cdot S_i^{p+}) + \sum_{j=1}^K \exp(\tau \cdot S_j^{p-})} \quad (4)$$

, where τ is learnable and initialized to 1. On the training dataset \mathcal{D} , the joint contrastive learning objective is formulated as:

$$\mathcal{L} = \sum_{i=1}^{|\mathcal{D}|} (\mathcal{L}_i^d + \mathcal{L}_i^p) \quad (5)$$

³Note <https://huggingface.co/docs/transformers/perplexity> for LM perplexity calculation. We empirically discard the outer exponential term in the PPL formula, because it already exists in NCE loss Eq. (4) and does not affect the final ranking.

2.3 Model Initialization and Inference

As UniTRec is a standard text-to-text Transformer, we initialize the parameters from pretrained BART (Lewis et al., 2020). In inference, UniTRec predicts the discriminative and perplexity-based scores for each candidate item, respectively. The two separate scores S^d and S^p are normalized, averaged, and finally ranked as the output. Detailed ranking process is provided in Appendix B.

3 Experiments

We evaluate UniTRec on three text-based recommendation tasks: 1) *NewsRec*, to recommend news articles to users based on their browsing history. We use the *MIND-small* dataset (Wu et al., 2020) for experiments. 2) *QuoteRec*, to recommend quotations to users based on their conversation history. We use the *Reddit-quotation* dataset (Wang et al., 2021) for experiments. 3) *EngageRec*, to recommend social media posts for users to engage with based on their comment history. We use the dataset released by Zeng et al. (2020) for experiments. Detailed dataset statistics is provided in Appendix A.

Implementation Details. The UniTRec encoder and decoder both consist of 6 Transformer layers with 768-dimensional hidden states and 12 attention heads. We set $L_1 = 3$ and $L_2 = 3$. We use AdamW optimizer (Loshchilov and Hutter, 2019) to train UniTRec with cosine learning rate decay.

Baselines. We compare UniTRec with competitive baselines: 1) *GRU4Rec* (Balázs et al., 2016) utilizes a GRU network to learn multi-turn history. 2) *SASRec* (Kang and McAuley, 2018) encodes user history with a self-attention based sequential model. 3) *BERT4Rec* (Sun et al., 2019) employs bidirectional self-attention to model user history. 4) *RoBERTa-Sim*, a simple yet strong baseline men-

Model	NewsRec			QuoteRec			EngageRec		
	MRR	NDCG@5/10	HR@5/10	MRR	NDCG@5/10	HR@5/10	MRR	NDCG@5/10	HR@5/10
GRU4Rec	32.91	36.20/42.53	50.33/68.35	34.08	34.65/37.93	44.45/54.63	2.12	1.04/1.51	1.27/2.65
SASRec	32.60	36.03/42.37	50.63/68.64	33.63	34.30/37.49	44.32/54.20	2.40	1.49/1.95	2.16/3.47
BERT4Rec	32.87	36.18/42.40	50.21/67.97	33.59	34.26/37.27	43.76/53.05	3.04	1.98/3.23	2.81/6.67
RoBERTa-Sim	32.96	36.47/42.81	51.06/69.08	37.13	37.96/41.18	48.14/58.06	3.74	2.66/3.75	4.42/7.70
UNBERT	33.09	36.53/42.84	50.87/68.82	39.75	40.74/43.69	50.90/60.04	2.83	1.96/2.67	3.11/5.24
UniTRec	33.76	37.63/43.74	52.61/69.89	41.24	42.38/45.31	52.87/61.88	4.06	3.23/4.29	4.58/7.68

Table 1: Experiment results on three text-based recommendation tasks. MRR denotes mean reciprocal rank, NDCG denotes normalized discounted cumulative gain, and HR denotes hit ratio (presented in percentage). The overall performance of UniTRec is better than other baseline models with p -value < 0.05 , validated by unpaired t-test.

Model	NewsRec			QuoteRec			EngageRec		
	MRR	NDCG@5/10	HR@5/10	MRR	NDCG@5/10	HR@5/10	MRR	NDCG@5/10	HR@5/10
UniTRec	33.76	37.63/43.74	52.61/69.89	41.24	42.38/45.31	52.87/61.88	4.06	3.23/4.29	4.58/7.68
w/o BART Init	30.31	33.32/39.69	47.55/65.78	19.02	17.66/20.80	22.45/32.16	2.24	0.86/1.61	1.27/3.62
w/o Local-Att	33.34	37.22/43.32	52.28/69.54	40.44	41.63/44.56	52.09/61.15	3.92	3.19/4.15	4.38/7.36
w/o Global-Att	33.22	37.06/43.17	52.14/69.47	40.25	41.47/44.26	52.07/60.76	3.64	2.78/3.59	3.89/6.35
Disc-Score only	33.07	36.76/43.03	51.68/69.46	40.59	41.81/44.65	52.39/61.14	3.82	2.99/3.60	4.49/6.85
PPL-Score only	32.83	36.39/42.59	51.05/68.67	40.31	41.43/44.47	52.13/61.20	3.29	2.39/3.03	3.86/5.66

Table 2: Recommendation performance of ablation model variants.

tioned in Qi et al. (2022), uses the hidden states of [CLS] tokens to measure user-candidate similarity. 5) UNBERT, implemented as Zhang et al. (2021), concatenates history and candidate texts as the input to BERT and predicts matching scores from the final hidden states of [CLS] tokens.

Note that we do not consider other methods that use non-text inputs (e.g., user profile, text topic labels). For fair comparison, all baseline models use pretrained 12-layer RoBERTa-base (Liu et al., 2019) as text encoders to learn embeddings of texts.

3.1 Main Results

Table 1 shows the performance of experiment models. From the results of *NewsRec* and *QuoteRec*, we can see that UniTRec outperforms all baseline models by a clear margin. Also, RoBERTa-Sim and UNBERT that directly use the [CLS] hidden states to represent user history, surpass other baselines that build additional aggregation networks upon the whole RoBERTa outputs. As displayed in the results, *EngageRec* is the most difficult task. We inspect the dataset and find that the texts on social media contain too much noise (e.g., URL and emoji), and the user history contains less number of turns. Nevertheless, UniTRec achieves better overall performance than other baseline models, validating its robustness on noisy text inputs and limited user history.

3.2 Ablation Studies and Analyses

We further conduct ablation studies on UniTRec. The experiment results are reported in Table 2.

Initialization of UniTRec. We train UniTRec from scratch without initialization from pretrained BART (refer to w/o BART Init). The recommendation performance significantly drops in all three tasks, which indicates that acquiring effective text understanding ability from PLM is a necessary key to UniTRec performance.

Local and global attention. We investigate the function of two-level attention modules of the UniTRec history encoder. Concretely, we set $L_1 = 0$ in w/o Local-Att and $L_2 = 0$ in w/o Global-Att, where $L_1 + L_2 = 6$. We can observe that removing local and global attention from the original UniTRec history encoder both lead to suboptimal performance, while the performance drop is more significant in w/o Global-Att. The results justify the effectiveness of jointly modeling two-level history contexts through adapted Transformer attention masking without additional parameters.

Discriminative and perplexity-based objectives. We probe into training UniTRec with standalone discriminative (Disc-Score only) and perplexity-based (PPL-Score only) contrastive objectives, respectively. We can see that the discriminative objective yields better performance than the perplexity-based objective. Besides, the model performance on both standalone objectives declines compared to the original joint objective. The results indicate that the discriminative and perplexity-based matching scores are complementary and can jointly provide more accurate signals of user history and candidate text matching for text-based recommendation.

4 Conclusion

We present a unified Transformer UniTRec for text-based recommendation. UniTRec learns two-level contexts of multi-turn user history and jointly exploits discriminative matching scores and candidate text perplexity as matching objectives. Empirical experiments on three text-based recommendation datasets corroborate the effectiveness of UniTRec.

5 Limitations

Our model only focuses on utilizing text information for recommendation, which is a key limitation of this work. In real-world settings, recommender systems are usually required to handle heterogeneous information inputs. UniTRec is a pure text-based recommender modeling user history and candidate texts as inputs. However, incorporating additional side information (e.g., user profile, text topic labels, and dwell time of user behaviors) could further improve the recommendation performance and alleviate the *cold start* problem. Furthermore, UniTRec only models two-level relations of user behavior history. Nonetheless, incorporating more user behavior information, such as implicit and negative feedback, could further enhance the recommendation performance.

Acknowledgements

We appreciate constructive comments from anonymous reviewers. The research described in this paper is partially supported by CUHK under Project No. 3230366.

References

Hidasi Balázs, Karatzoglou Alexandros, Baltrunas Linas, and Tikk Domonkos. 2016. [Session-based recommendations with recurrent neural networks](#). In *4th International Conference on Learning Representations ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*,

volume 33, pages 1877–1901. Curran Associates, Inc.

Daniel Cheng, Kyle Yan, Phillip Keung, and Noah A. Smith. 2022. [The engage corpus: A social media dataset for text-based recommender systems](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1885–1889, Marseille, France. European Language Resources Association.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. [Unified language model pre-training for natural language understanding and generation](#). In *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*.

Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. [Recommendation as language processing \(rlp\): A unified pretrain, personalized prompt & predict paradigm \(p5\)](#). In *Proceedings of the 16th ACM Conference on Recommender Systems, RecSys '22*, page 299–315, New York, NY, USA. Association for Computing Machinery.

Youyang Gu, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. [Learning to refine text based recommendations](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2103–2108, Austin, Texas. Association for Computational Linguistics.

Michael Gutmann and Aapo Hyvärinen. 2010. [Noise-contrastive estimation: A new estimation principle for unnormalized statistical models](#). In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 297–304, Chia Laguna Resort, Sardinia, Italy. PMLR.

Wang-Cheng Kang and Julian McAuley. 2018. [Self-attentive sequential recommendation](#). In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 197–206.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.

- Jian Li, Jieming Zhu, Qiwei Bi, Guohao Cai, Lifeng Shang, Zhenhua Dong, Xin Jiang, and Qun Liu. 2022. [MINER: Multi-interest matching network for news recommendation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 343–352, Dublin, Ireland. Association for Computational Linguistics.
- Yize Li, Jiazhong Nie, Yi Zhang, Bingqing Wang, Baoshi Yan, and Fuliang Weng. 2010. [Contextual recommendation based on text mining](#). In *Coling 2010: Posters*, pages 692–700, Beijing, China. Coling 2010 Organizing Committee.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). In *arXiv preprint arXiv: 1907.11692*. arXiv.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Itzik Malkiel, Oren Barkan, Avi Caciularu, Noam Razin, Ori Katz, and Noam Koenigstein. 2020. [RecoBERT: A catalog language model for text-based recommendations](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Online. Association for Computational Linguistics.
- 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.
- Fanchao Qi, Yanhui Yang, Jing Yi, Zhili Cheng, Zhiyuan Liu, and Maosong Sun. 2022. [QuoteR: A benchmark of quote recommendation for writing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 336–348, Dublin, Ireland. Association for Computational Linguistics.
- Tao Qi, Fangzhao Wu, Chuhan Wu, Peiru Yang, Yang Yu, Xing Xie, and Yongfeng Huang. 2021. [HieRec: Hierarchical user interest modeling for personalized news recommendation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5446–5456, Online. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. [Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, page 1441–1450, 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.
- Lingzhi Wang, Xingshan Zeng, and Kam-Fai Wong. 2021. [Quotation recommendation and interpretation based on transformation from queries to quotations](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 754–758, Online. Association for Computational Linguistics.
- Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2021. [Empowering news recommendation with pre-trained language models](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 1652–1656, New York, NY, USA. Association for Computing Machinery.
- 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.
- Xingshan Zeng, Jing Li, Lu Wang, Zhiming Mao, and Kam-Fai Wong. 2020. [Dynamic online conversation recommendation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3331–3341, Online. Association for Computational Linguistics.
- Qi Zhang, Jingjie Li, Qinglin Jia, Chuyuan Wang, Jieming Zhu, Zhaowei Wang, and Xiuqiang He. 2021. [Unbert: User-news matching bert for news recommendation](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3356–3362. International Joint Conferences on Artificial Intelligence Organization. Main Track.

Dataset	<i>NewsRec</i>	<i>QuoteRec</i>	<i>EngageRec</i>
Avg. history turns	26.09	4.24	3.29
Avg. history tokens	414.40	279.82	286.82
Avg. candidates	37.23	1111	7163
Avg. candidate tokens	16.15	19.11	102.42

Table 3: Statistics of three text-based recommendation training datasets. History and candidate **tokens** denote the number of BPE-tokenized tokens. The test set distribution is closed to the training sets (except candidates of *EngageRec*) and hence omitted. Note that the max length of each history log is truncated to 1024 tokens.

A Dataset Statistics

The detailed statistics of the three text-based recommendation datasets are displayed in Table 3. Note that we use news titles as the text inputs for *NewsRec* following Qi et al. (2021). *NewsRec* regards the user clicked and non-clicked news as candidate texts, while *QuoteRec* and *EngageRec* regard all potential quotation texts and post texts as candidates. Different from Zeng et al. (2020) that formulates the task as recommending candidate users to given posts based on post content, we formulate the task as recommending candidate posts to given users based on user history.

Algorithm 1 Candidate Ranking Process

Input: discriminative scores $S^d = \{S_1^d, S_2^d, \dots, S_M^d\}$,
perplexity-based scores $S^p = \{S_1^p, S_2^p, \dots, S_M^p\}$.

Output: final averaged ranking \bar{R} .

- 1: Derive the normalized discriminative scores $S_{norm}^d = \text{softmax}(S^d)$.
 - 2: Derive the normalized perplexity-based scores $S_{norm}^p = \text{softmax}(S^p)$.
 - 3: Derive the geometric average scores $\bar{S} = \log(S_{norm}^d) + \log(S_{norm}^p)$.
 - 4: Sort the averaged scores \bar{S} by descending order to derive the final ranking: $\bar{R} \leftarrow \text{Rank}_{\text{des}}(\bar{S})$.
 - 5: **return** \bar{R}
-

B Inference Ranking

Given the user history and M candidate texts, UniTRec first predicts the discriminative ranking scores $S^d = \{S_1^d, S_2^d, \dots, S_M^d\}$ and perplexity-based ranking scores $S^p = \{S_1^p, S_2^p, \dots, S_M^p\}$ of the candidates. Algorithm 1 outlines an approach to aggregate the final ranking based on S^d and S^p . Note that the function $\text{Rank}(S)$ ⁴ denotes outputting the sorted order of elements in a score list S . There exist other ways to average the ranking of S^d and S^p , which we leave for future work to explore.

⁴ $\text{Rank}(S)$ works similarly to `scipy.stats.rankdata()`. For example in ascending order, $\text{Rank}_{\text{asc}}(\{0.2, 0.6, 0.7, 0.4\}) = \text{scipy.stats.rankdata}([0.2, 0.6, 0.7, 0.4]) = [1, 3, 4, 2]$

C Qualitative Analysis

We show randomly sampled outputs of UniTRec, for instance, demonstrated on the news recommendation and quote recommendation tasks. Table 4 and 5 showcase the qualitative samples.

Turn	<i>History News Texts</i>				
#1	Mac Engel: As long as these results are acceptable, Dallas Cowboys will continue to be losers				
#2	NFL world reacts to officials handing Packers win over Lions				
#3	Maryland Congressman Elijah Cummings, a Democrat and Chair of House Oversight and Reform Committee, has died: CNN				
#4	Unprecedented movement detected on California earthquake fault capable of 8.0 temblor				
#5	Bag Explodes While Being Loaded On Volaris Flight At Midway Airport				
#6	Orlando Scandrick rips Eagles: They have "accountability issues"				
#7	Meghan King Edmonds, Jim Edmonds' Nanny Denies Cheating Allegations				
#8	Nearly \$400M worth of cocaine and marijuana intercepted by US Coast Guard				
#9	Former NBA first-round pick arrested in sex sting operation				
#10	China's trade with US shrinks in October despite optimism				
<i>Candidate News Texts</i>		S^d	S^p	\bar{R}	Clicked
	Taylor Swift Rep Hits Back at Big Machine, Claims She's Actually Owed \$7.9 Million in Unpaid Royalties	0.095	0.069	4	✗
	Former North Carolina State, NBA player Anthony Grundy dies in stabbing, police say	0.172	0.155	3	✗
	13 Reasons Why's Christian Navarro Slams Disney for Casting "the White Guy" in The Little Mermaid	0.048	0.065	7	✗
	Opinion: Colin Kaepernick is about to get what he deserves: a chance	0.303	0.250	1	✓
	3 Indiana judges suspended after a night of drinking turned into a White Castle brawl	0.076	0.059	5	✗
	66 Cool Tech Gifts Anyone Would Be Thrilled to Receive	0.009	0.005	9	✗
	Police find 26 children behind false wall at Colorado day care	0.034	0.116	6	✗
	I've been writing about tiny homes for a year and spent 2 nights in a 300-foot home to see what it is all about	0.029	0.019	8	✗
	Report: Police investigating woman's death after Redskins' player Montae Nicholson took her to hospital	0.235	0.261	2	✓

(i) Qualitative Example-A from news recommendation.

Turn	<i>History News Texts</i>				
#1	Toddler dancing to celebrate 11 months cancer-free goes viral				
#2	NFL Week 8 Power Rankings: Old-school football rules the day				
#3	The 25 US cities where it's easiest to get a mortgage				
#4	Burning questions for Cowboys vs Giants on "Monday Night Football"				
#5	Who's the favorite to win 2019 NFL rushing title?				
#6	Grading all 32 NFL teams heading into the last eight weeks of the 2019 season				
#7	Jennifer Aniston looks amazing in a makeup-free selfie, plus more news				
#8	This \$12 million "mansion yacht" is made entirely of stainless steel and it's a first for the industry. Take a peek inside				
<i>Candidate News Texts</i>		S^d	S^p	\bar{R}	Clicked
	Opinion: Colin Kaepernick is about to get what he deserves: a chance	0.330	0.400	1	✓
	U.S. Troops Will Die If They Remain in Syria, Bashar Al-Assad Warns	0.024	0.011	10	✗
	Pete Davidson, Kaia Gerber Are Dating, Trying to Stay "Low Profile"	0.064	0.033	6	✗
	The Hottest Tech Gifts This Holiday Season	0.050	0.027	8	✗
	Taylor Swift Rep Hits Back at Big Machine, Claims She's Actually Owed \$7.9 Million in Unpaid Royalties	0.046	0.038	7	✗
	13 Reasons Why's Christian Navarro Slams Disney for Casting "the White Guy" in The Little Mermaid	0.060	0.096	4	✓
	Some believe Mason Rudolph, hit in head with his own helmet, isn't getting enough blame	0.154	0.179	2	✓
	South Carolina teen gets life in prison for deadly elementary school shooting	0.066	0.046	5	✗
	The Unlikely Star of My Family's Thanksgiving Table	0.047	0.021	9	✗
	Police investigating woman's death after Redskins' player Montae Nicholson took her to hospital	0.158	0.149	3	✗

(ii) Qualitative Example-B from news recommendation.

Table 4: Case analyses of news recommendation. *History News Texts* are sorted by user-clicked timestamps. S^d , S^p , and \bar{R} are normalized discriminative, perplexity-based scores, and average ranking as described in Appendix B. **Clicked** denotes the ground truth user-click labels. Note that the experiment history logs are anonymized and delinked, which is always the first priority of the recommendation study.

Turn	Conversation Threading History
#1	I own an FJ. It's a great car and even on stockies. It s great offroad.
#2	I feel bad for you that you run the risk of being associated with the typical FJ owner.
#3	What is a typical FJ owner? I've not heard anything bad about FJ owners.
#4	It's like someone who drives a jeep wrangler in NYC. There's no need. Tons of FJ owners do that have it and not use it for what it's made for.
#5	God forbid someone likes the design of a car and doesn't use it offroad.
#6	Then buy a much more economic environmentalist friendly version. If you buy something and always use it for much less than it's purpose, why buy it?
#7	Or people can buy whatever the hell they want because it's their money and not yours.
#8	You're entirely right. Just like people can be rude just because you can do it, because you have the ability but why should you ass.
#9	I wasn't aware that somebody buying a vehicle that they like and you don't was morally wrong.
#10	I love FJs. It's perfectly fine to buy whatever you think looks nice.

Candidate Quote Texts	S^d	S^p	\bar{R}	Ground truth
Beauty is in the eye of the beholder.	0.480	0.471	1	✓
A fool and his money are soon parted.	0.176	0.140	2	
Form follows function.	0.051	0.046	3	
Everything is worth what its purchaser will pay for it.	0.040	0.058	4	
Because it's there.	0.038	0.029	5	
You can't fix stupid.	0.021	0.034	6	
The lady doth protest too much, methinks.	0.022	0.013	7	
It's all about the money.	0.020	0.013	8	
Anybody driving slower than you is an idiot, and anyone going faster than you is a maniac?	0.012	0.018	9	
Opportunity is missed by most people.	0.018	0.008	10	

(iii) Qualitative Example-C from quote recommendation.

Turn	Conversation Threading History
#1	Society is becoming more efficient, which is a good thing. People should realize there's no point in holding back this technology just for the sake of keeping people employed. If this were beneficial, then calculators and computers shouldn't exist either.
#2	One small problem is that people need to pay rent and eat.
#3	So we should ditch computers and go back to the typing pool? Should we get rid of heavy earth moving equipment and just use hundreds of guys with hand tools to build everything? It would employ a hell of a lot more people.
#4	No one's saying that. I don't think anyone is really against automation, but as it increases, there are soon going to be more people that there are jobs that actually need doing. I actually believe we've already passed this point. So what do we do with the people, who can't get jobs simply because there are none? It's an issue that need assessed immediately.
#5	Tons and tons and tons of American jobs have been replaced by new jobs created by technology or in support of technology years ago. An office might have needed people to handle filing paperwork, keeping it in order, and retrieving, where now a document management system has made them completely redundant. The upshot is that to access that DMS, people are out there selling computers, installing computers, servicing computers, and supporting end users building the servers installing, supporting monitoring backing them up, and all that jobs that come in support of those progress is progress. And it advances human efficiency and knowledge. These are just one or two examples, but the answer is not to kill progress. Other countries simply won't. The answer is to push education to the forefront, so people are prepared for these jobs and whatever other challenges the future may bring.
#6	This is true. But it s unfortunate technological advances tend to reduce low skill jobs and replace them with high skill jobs. It would feel more fair if the low skilled workers could all do training programs and become high skilled workers. But this isn't really the case. Those jobs end up being taken by someone who had better educational opportunities or someone younger who still has time to take advantage of education.
#7	The reality is the reality. Unfortunate or not educating people will create more educated people to handle high skill jobs, and I'll tell you being a desktop support technician isn't high skill. As that's where we push in the future, any amount of hand wringing won't change the facts. We must educate our people if we want to be a global leader in more than homelessness poverty.
#8	Education won't matter. We are at the end of the job age at some point in the near future. We are going to have to deal with the fact that getting a job isn't a reality for a significant percentage of the population. Society will have to radically change as it did during the industrial revolution.
#9	Much cheaper to heavily discourage having more children free abortions. Then in years there won't be so many useless people who can apparently be replaced by a simple robot.
#10	Virtually every job will be replaced by automation name skilled trades that can't be automated. I imagine you'd be surprised at how hard this is. Are pharmacists useless, surgeons, accountants? I'd bet that your job is just as replaceable as these.

Candidate Quote Texts	S^d	S^p	\bar{R}	Ground truth
There's no such thing as a free lunch.	0.365	0.417	1	
I can't predict the future.	0.185	0.210	2	✓
I have never let my schooling interfere with my education.	0.104	0.059	3	
Prevention is better than cure.	0.044	0.083	4	
Knowledge is power.	0.059	0.052	5	
Don't let schooling interfere with your education.	0.044	0.043	6	
Nature abhors a vacuum.	0.036	0.024	7	
There is no substitute for hard work.	0.024	0.017	8	
There are three kinds of lies: lies, damned lies, and statistics.	0.022	0.013	9	
You can't fix stupid.	0.019	0.010	10	

(iv) Qualitative Example-D from quote recommendation.

Table 5: Case analyses of quote recommendation. We demonstrate the candidate quotes of the top 10 rankings out of all candidates. Note that there is only one ground truth note for each conversation history.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 5
- A2. Did you discuss any potential risks of your work?
We see no concern about potential risks.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

The Abstract provides the link to our code.

- B1. Did you cite the creators of artifacts you used?
Not applicable. Left blank.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
In the Abstract, a Github repository with documentation is released.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Appendix A

C Did you run computational experiments?

Section 3

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 3

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 3

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 3

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Not applicable. Left blank.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.