method

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

task

Advisor : Jia-Ling, Koh

Speaker : Hsuan, Lu

Source : ACL’2023

Date : 2023/11/28

Outline

- Introduction
- Method
- Experiment
- Conclusion

Text-based Recommendation

Testing data

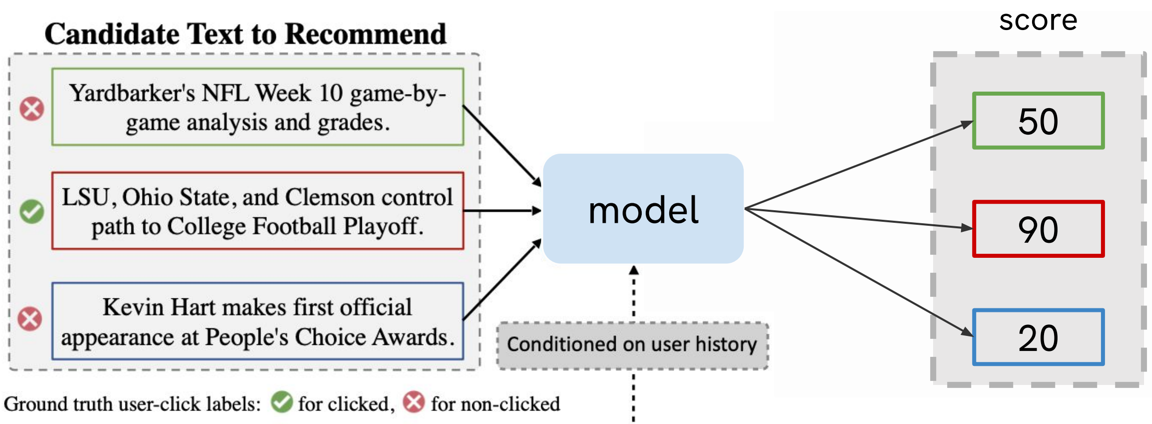

Input
Training data

Multi-turn User History

Turn	User-browsed News Titles
1	AP Top 25: LSU jumps to No. 2, upset drops Georgia to No. 10.
2	5 college games to watch this Saturday.
3	LSU surging, Big Ten reckoning and more we learned from college football's "Separation Saturday".

⋮

PLM Empowered News Recommendation

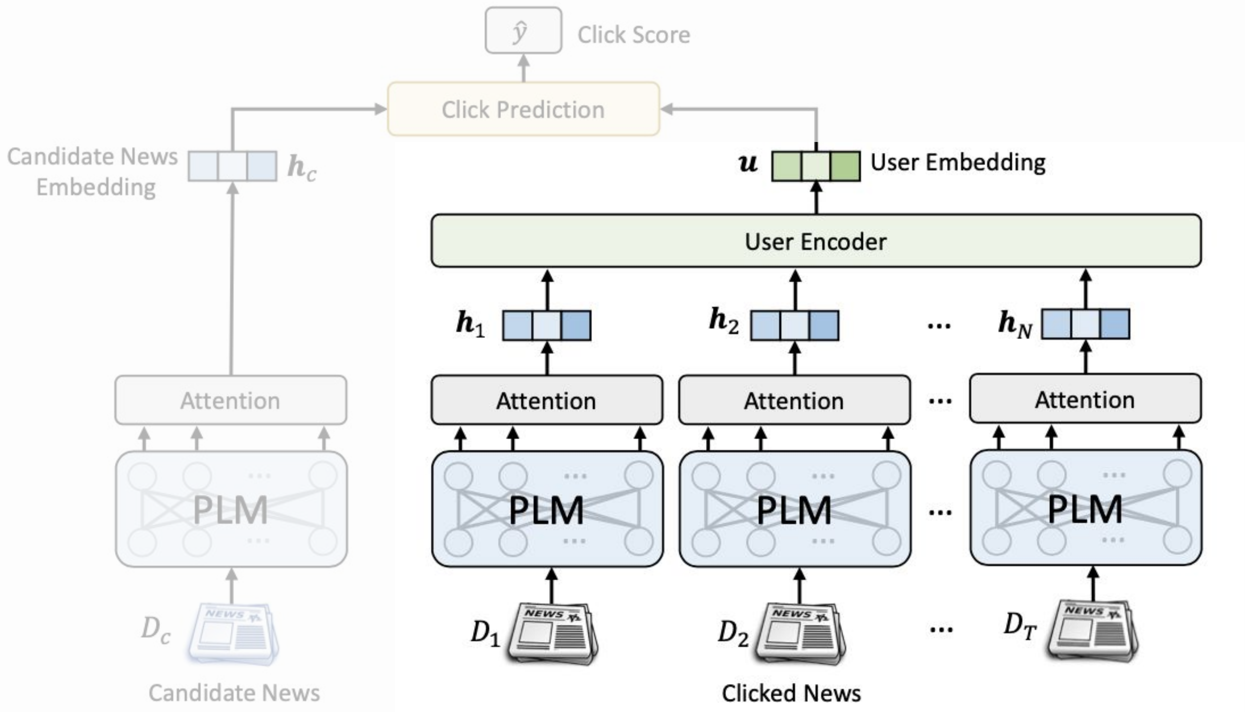

Outline

- Introduction
- Method
- Experiment
- Conclusion

Text-based Recommendation

Perplexity困惑度越低
代表跟使用者匹配的機率越高

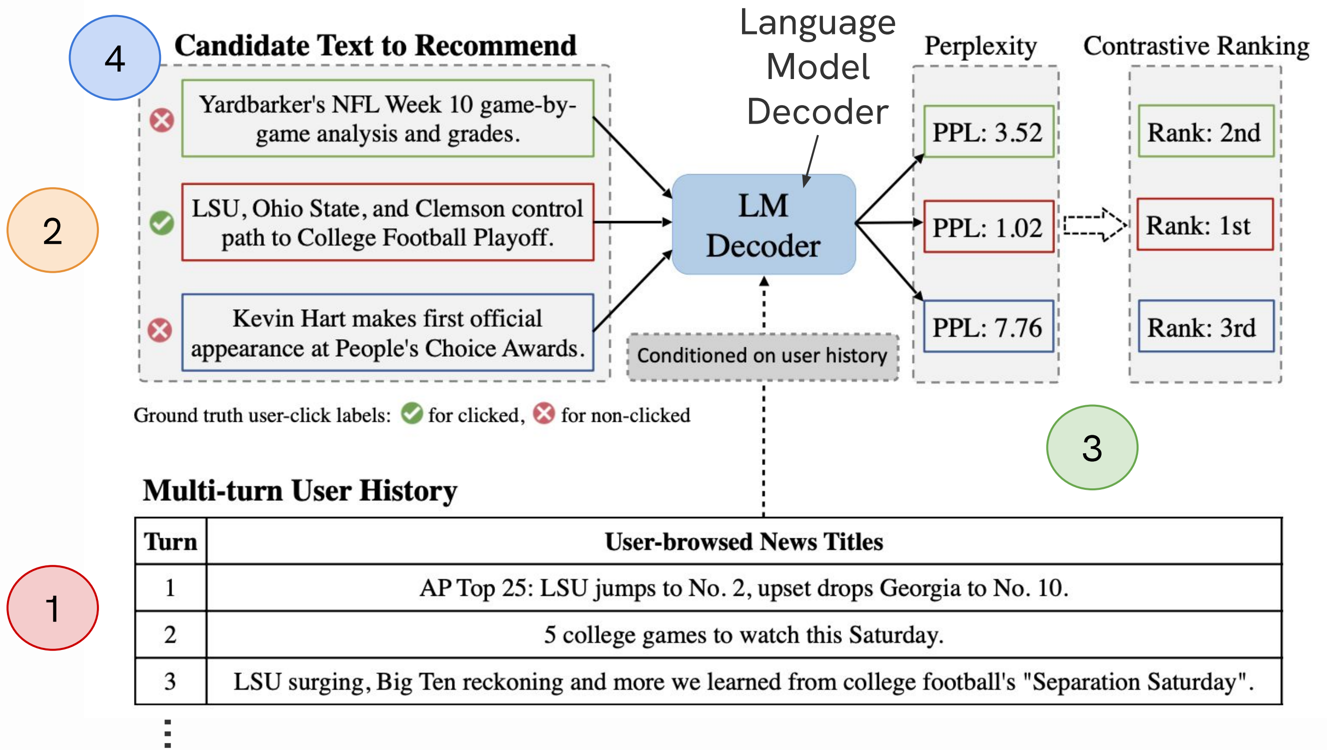

Overview of UniTRec

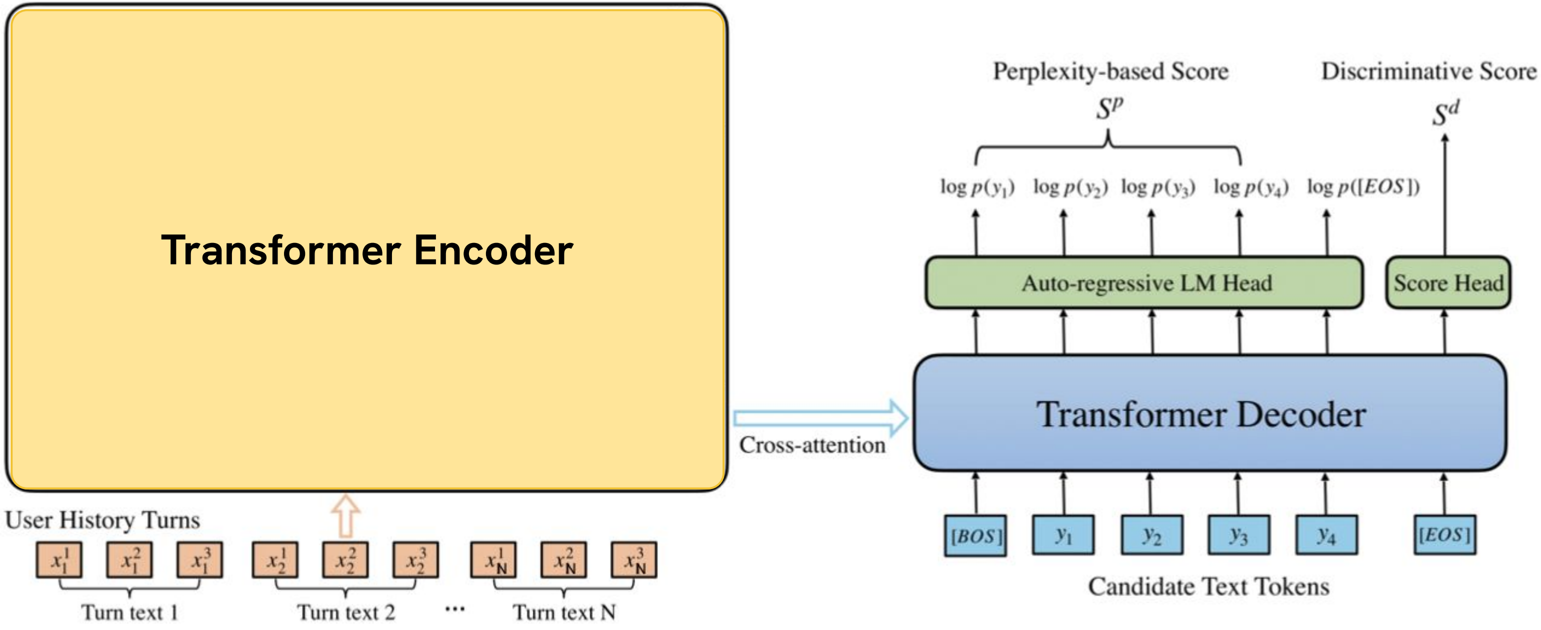

Unified User-history Modeling

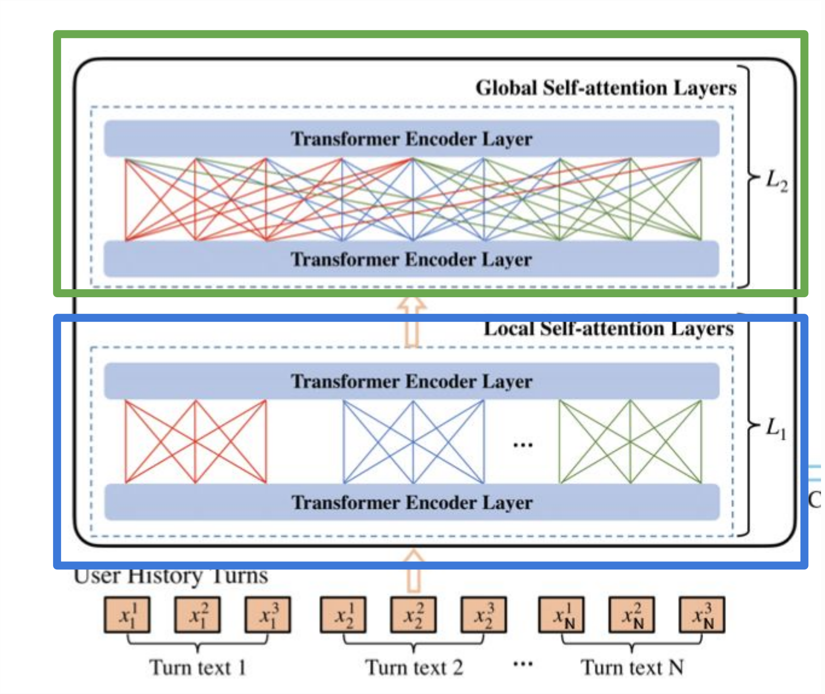

turn-level

word-level

Unified User-history Modeling — Input

Multi-turn history of a user $H = [t_1, t_2, ..., t_N]$

Each turn text $t_i = [x_i^1, x_i^2, ..., x_i^{|t_i|}]$

Input token $X = [x_1^1, x_1^2, ..., x_1^{|t_1|}, ..., x_N^1, x_N^2, ..., x_N^{|t_N|}]$

Unified User-history Modeling — Self-Attention

$$X = [x_1^1, x_1^2, ..., x_1^{|t_1|}, ..., x_N^1, x_N^2, ..., x_N^{|t_N|}]$$

- Local attention on word-level context

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

- Global attention on turn-level context => M = 0

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

Self-attention

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

key

query

value

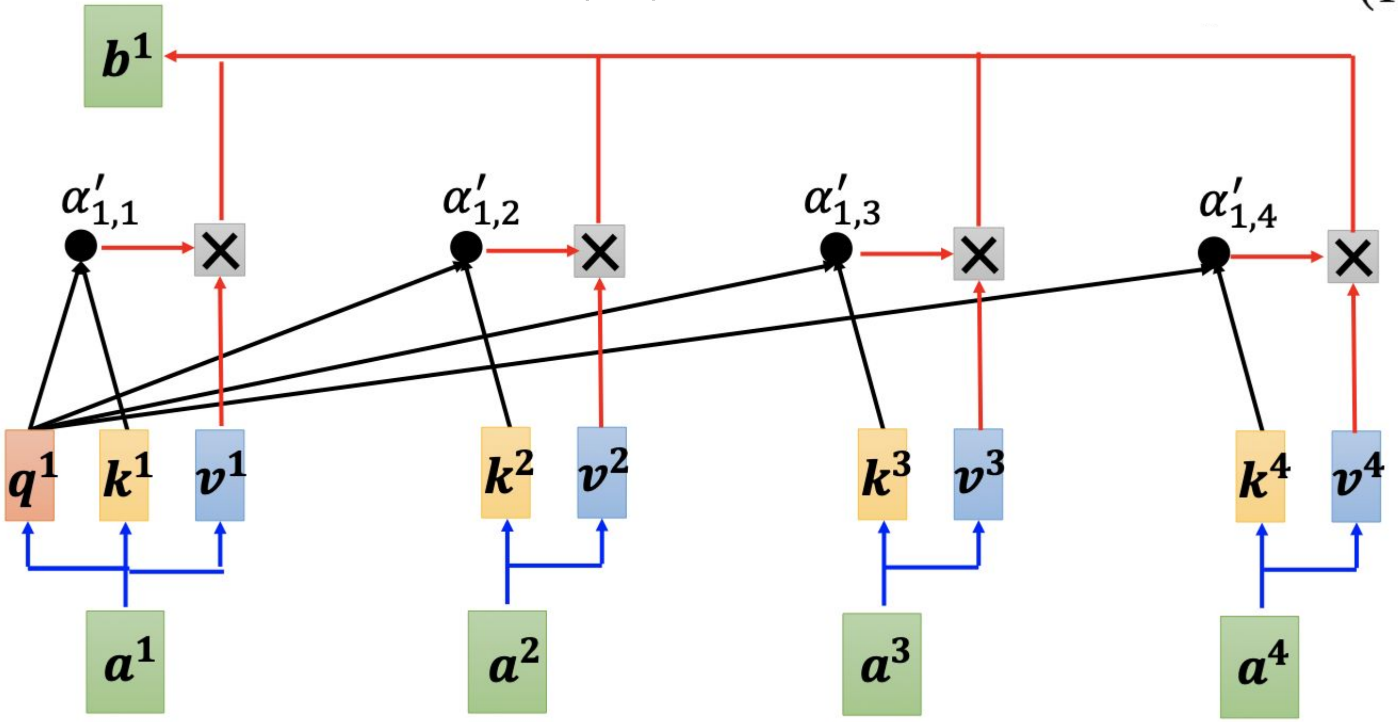

Joint Contrastive Ranking Objectives

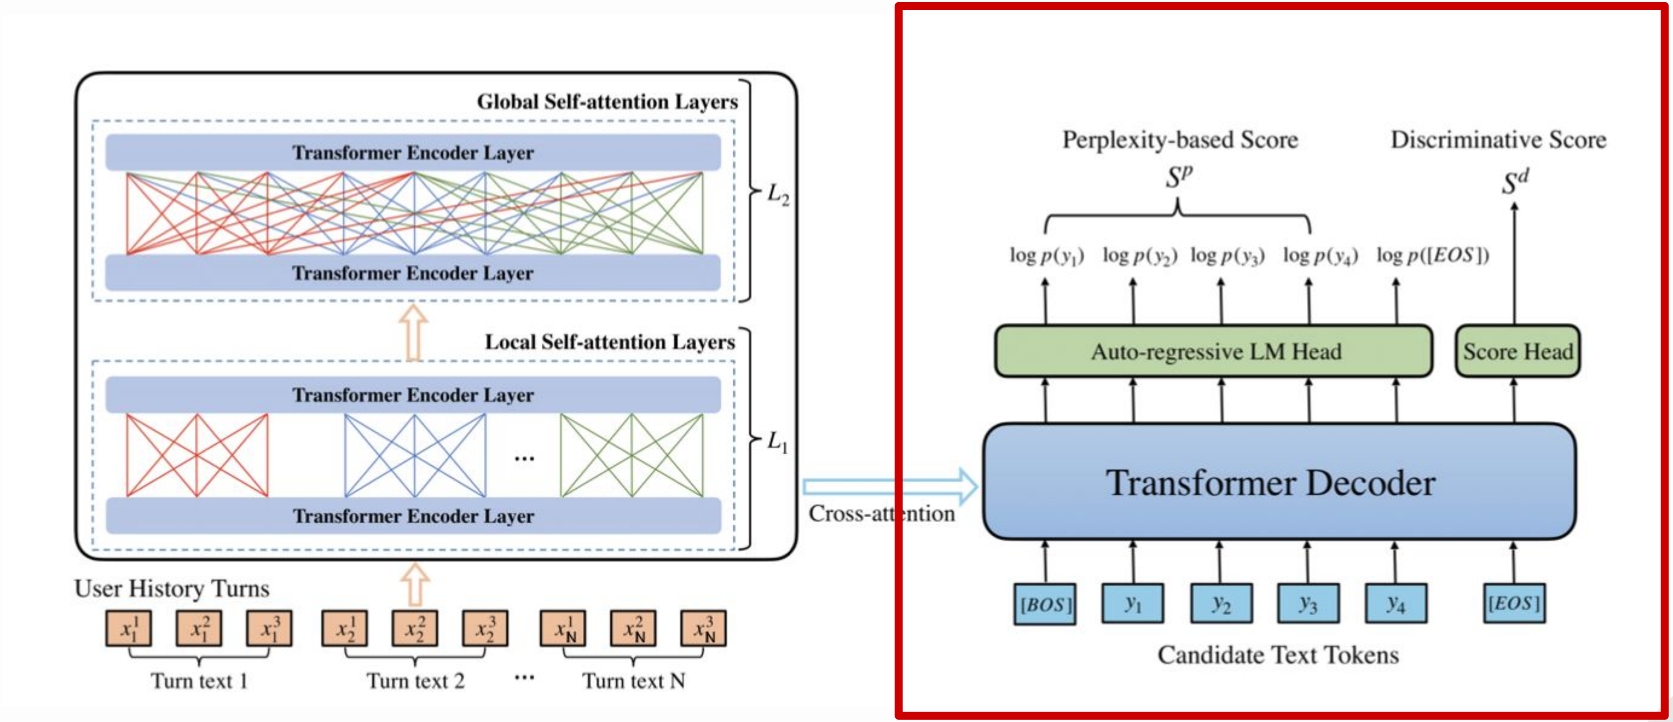

Objective on Discriminative Scores

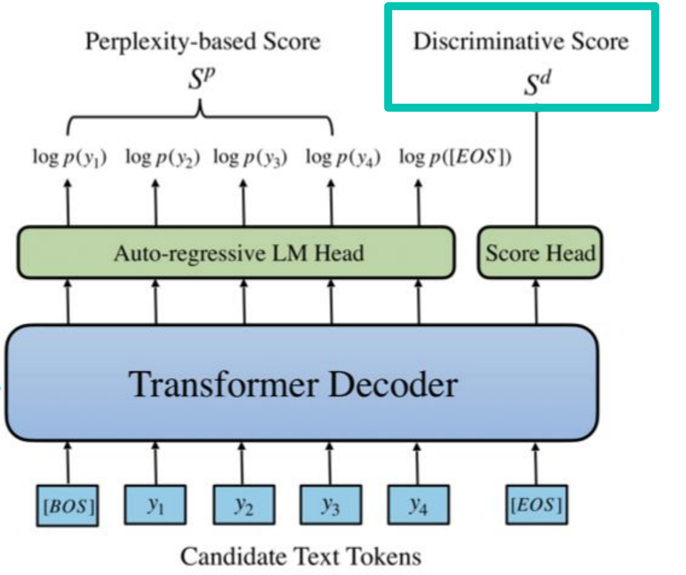

loss

matching score

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

unmatched negative candidates' matching scores

= -log1 = 0 ← ideal

Objective on Candidate Text Perplexity

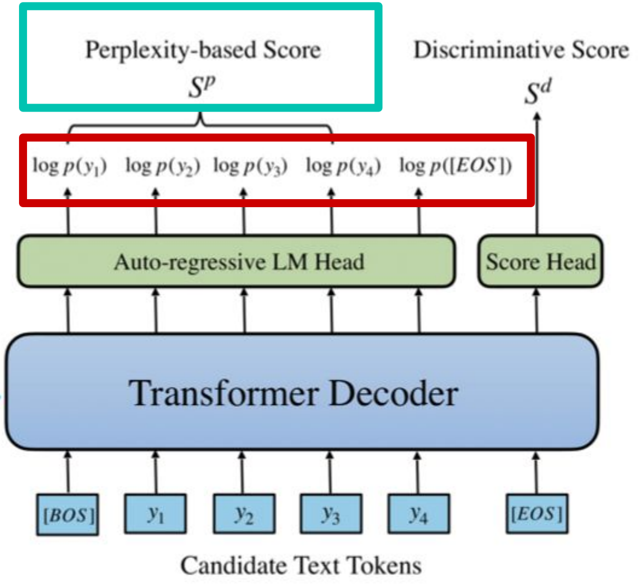

Candidate text: $Y = [y_1, y_2, ..., y_T]$

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

wide range

learnable and initialized to 1

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

loss

= -log1 = 0 ← ideal

Joint Contrastive Ranking Objectives

D: training dataset / 1 batch

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

越接近0越好

Model Initialization and Inference

- Initialize the parameters from pretrained BART

Discriminative scores $S^d = \{S_1^d, S_2^d, ..., S_M^d\}$

Perplexity-based scores $S^p = \{S_1^p, S_2^p, ..., S_M^p\}$

S = log(S^d) + log(S^p)

Rank(S) = Rank({0.2, 0.6, 0.7, 0.4}) = [4, 2, 1, 3]

Outline

- Introduction
- Method
- Experiment
- Conclusion

Dataset

篇數

總字數

Dataset	*NewsRec*	*QuoteRec*	*EngageRec*
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

平均一篇的字數

Baseline

- GRU4Rec
- SASRec
- BERT4Rec

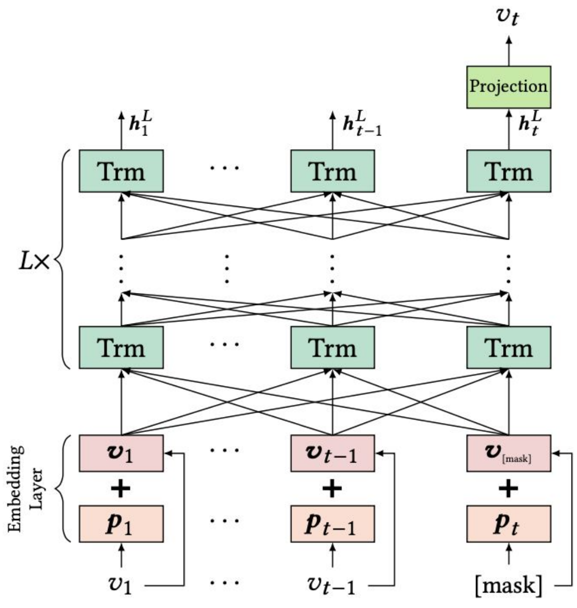

(b) BERT4Rec model architecture.

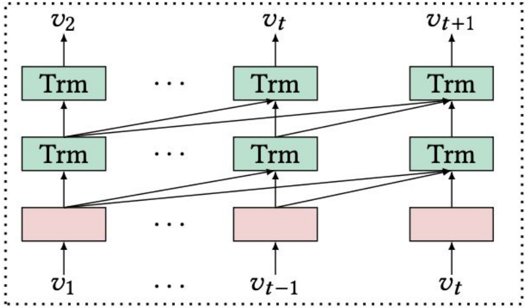

(c) SASRec model architecture.

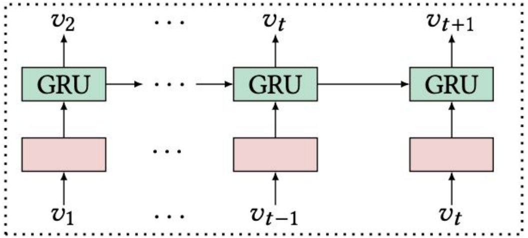

(d) RNN based sequential recommendation methods.

Baseline

- RoBERTa-Sim

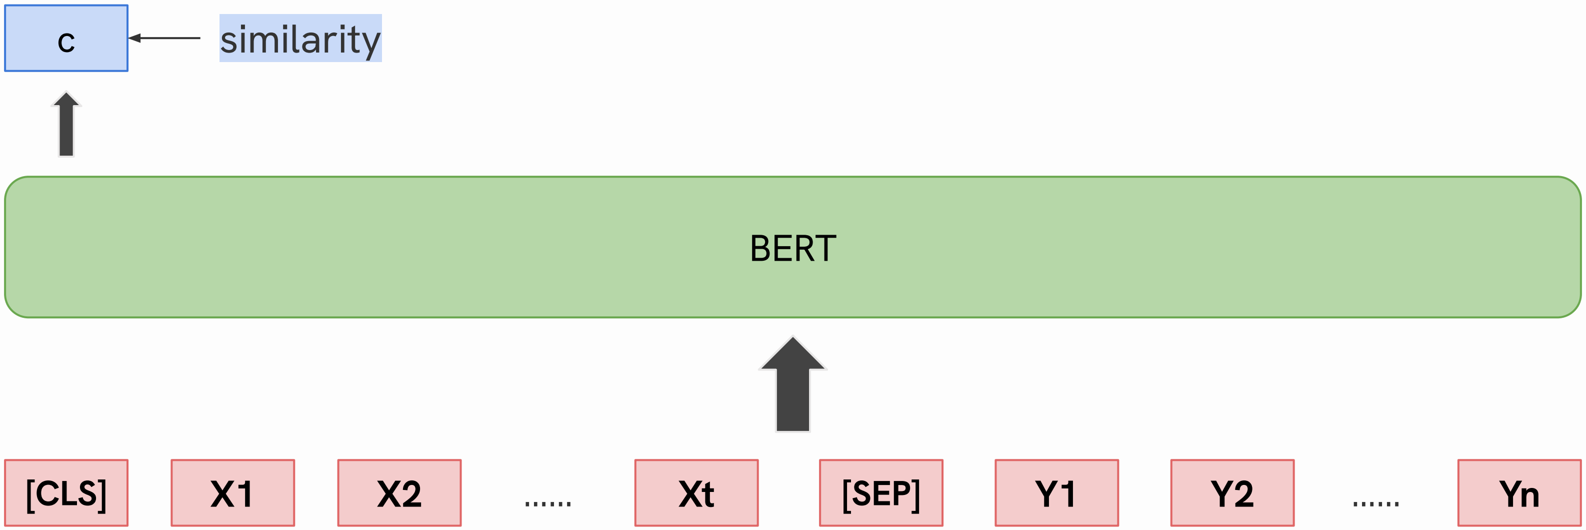

Baseline

- UNBERT

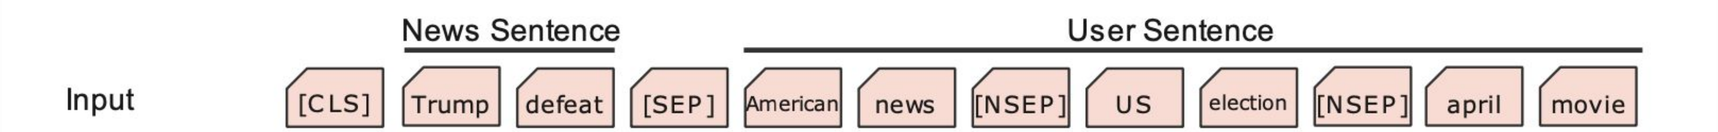

Baseline – UNBERT

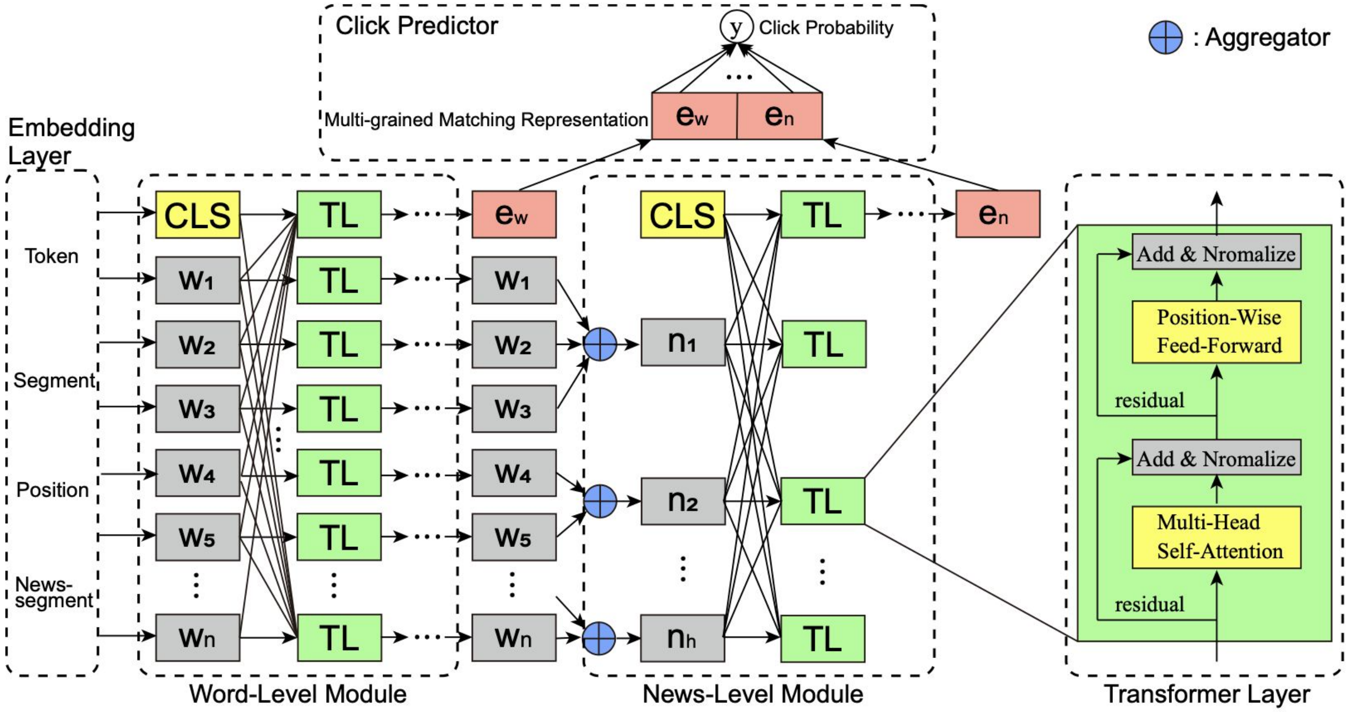

Evaluation

- MRR (Mean Reciprocal Rank): $\mathrm{MRR} = \frac{1}{\mathrm{N}} \sum_{\mathrm{i}=1}^{\mathrm{N}} \frac{1}{\mathrm{p_i}}$
- NDCG (Normalized Discounted Cumulative Gain):

$$\mathrm{NDCG} = \frac{1}{\mathrm{N}} \sum_{\mathrm{i}=1}^{\mathrm{N}} \frac{1}{\log_2(\mathrm{p_i} + 1)}$$

- HR (Hits Ratio): $\mathrm{HR} = \frac{1}{\mathrm{N}} \sum_{\mathrm{i}=1}^{\mathrm{N}} \mathrm{hits(i)}$

N: 總數量
Pi: 第i個文章/物品在推薦列表中的順位

Experiment

- EngageRec dataset contains too much noise (e.g. URL, emoji), and the user history contains less number of turns

Model	*NewsRec* MRR	*NewsRec* NDCG@5/10	*NewsRec* HR@5/10	*QuoteRec* MRR	*QuoteRec* NDCG@5/10	*QuoteRec* HR@5/10	*EngageRec* MRR	*EngageRec* NDCG@5/10	*EngageRec* 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

Ablation Study

- w/o BART Init

Model	*NewsRec* MRR	NDCG@5/10	HR@5/10	*QuoteRec* MRR	NDCG@5/10	HR@5/10	*EngageRec* 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

Ablation Study

- w/o Local-Att => L1 = 3 → 0
- w/o Global-Att => L2 = 3 → 0

	NewsRec			*QuoteRec*			*EngageRec*		
Model	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 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

Ablation Study

- Dis-Score only
- PPL-Score only

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
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

Outline

- Introduction
- Method
- Experiment
- Conclusion

Conclusion

- UniTRec learns two-level contexts of multi-turn user history and jointly exploits discriminative matching scores and candidate text perplexity as matching objectives.