Let It Go? Not Quite: Addressing Item Cold Start in Sequential Recommendations with Content-Based Initialization

Anton Pembek, Artem Fatkulin, Anton Klenitskiy, Alexey Vasilev

SBER AI Lab

Cold start problem

Sequential Recommender Systems predict a user's next action based on past interactions but struggle when items have few or no interaction data

Without sufficient interaction history, the model cannot learn effective item representations

This leads to poor recommendation quality for new or rarely interacted items. This is especially important where new items appear quite frequently (news, streaming services)

Approaches with initialized content embeddings

A common approach to this issue is using content-based Information. For all items (cold and warm) we can use information like product descriptions or audio features to generate embeddings and initialize by them the embedding layer of the backbone model.

Frozen embeddings

+ It's easy to add cold items because they are from the same space as warm ones
- Due to the small number of trainable weights, the model learns poorly

Trainable embeddings

+ Model quality metrics are better than for the original model
- The embeddings of warm items may differ significantly from content-based embeddings used for cold items

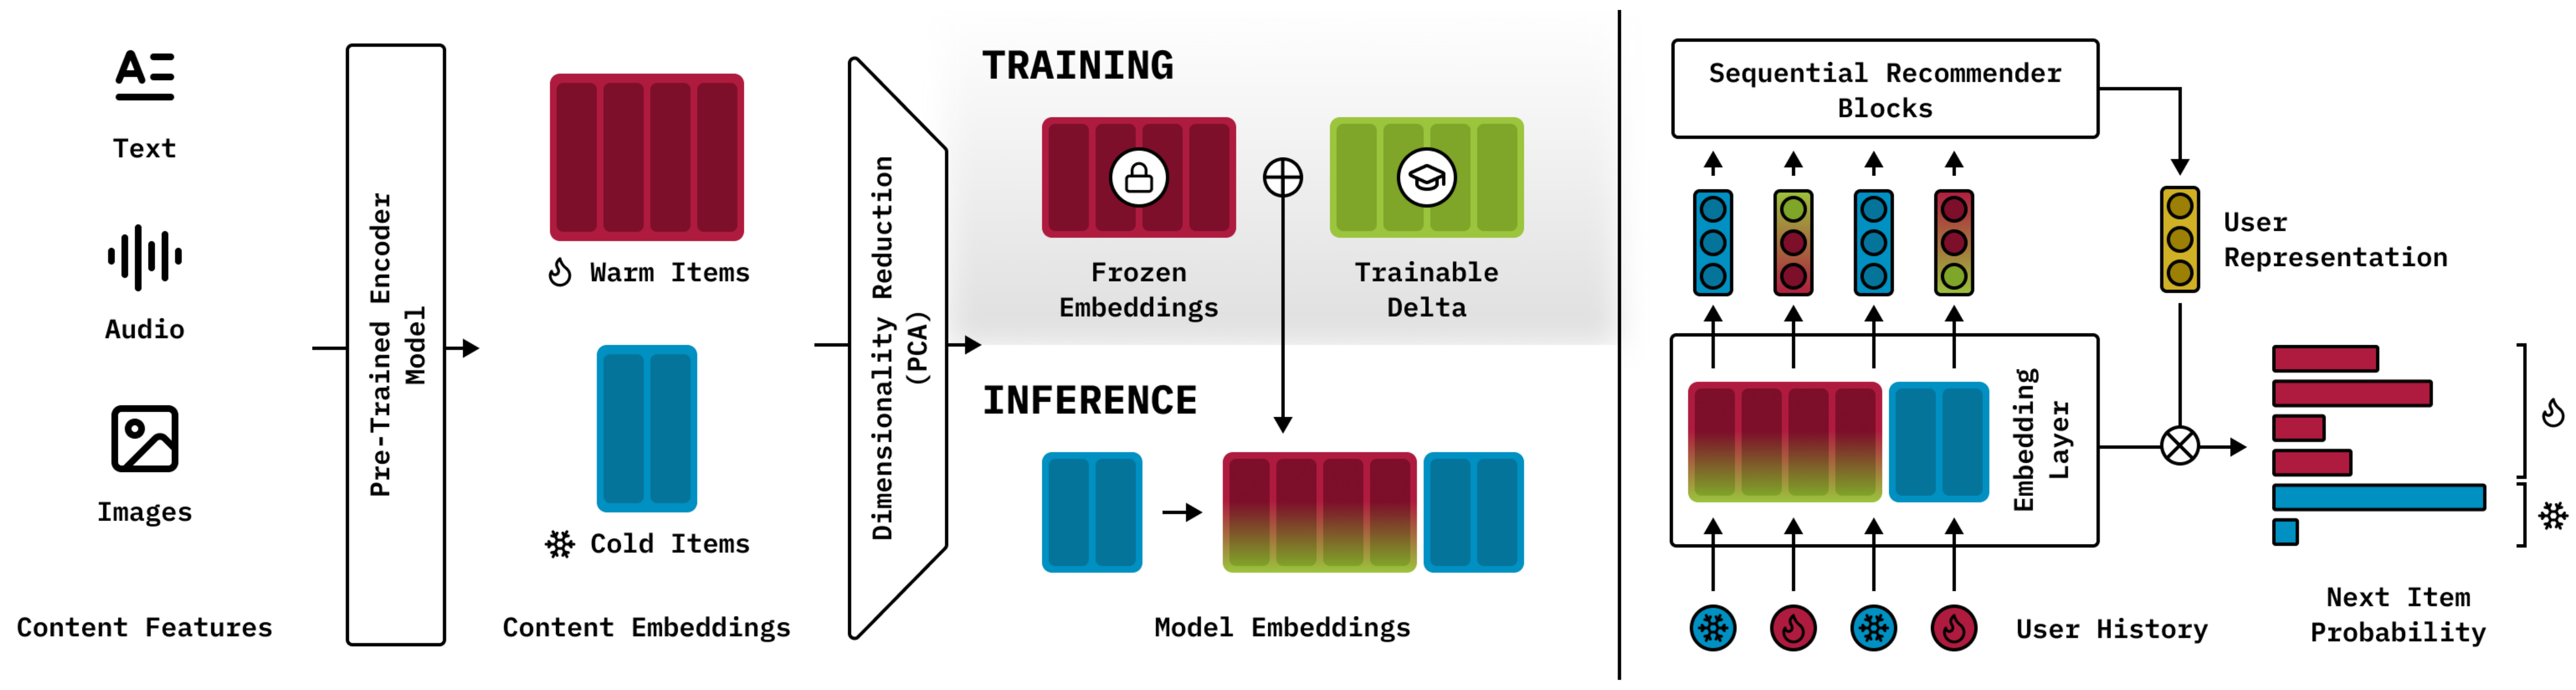

Our approach

Combining both approaches:

1. Freezing content-based item embeddings
2. Fixing their norm
3. Training only a small delta layer

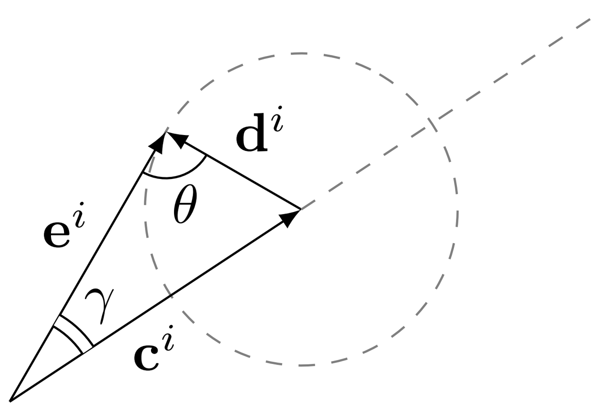

Training only a small delta layer allows the model to adjust item representations slightly while keeping them close to the original embeddings

Dataset	# Users	# Items	# Interact.	Avg. length	Cold items in GT
Amazon-M2 FR [9]	129,983	44,049	566,806	4.3	7%
Beauty [14]	21,029	11,733	149,147	7.1	25%
Zvuk [16]	9,076	131,085	3,236,653	356.6	13%

Results

Performance on cold, warm, and all ground-truth items

Metric	Model	Amazon-M2			Beauty			Zvuk		
		Cold GT	Warm GT	Total	Cold GT	Warm GT	Total	Cold GT	Warm GT	Total
HR@10	Content-based KNN	0.454±0.000	0.383±0.000	0.388±0.000	**0.043±0.000**	0.044±0.000	0.044±0.000	0.009±0.000	0.000±0.000	0.001±0.000
	SASRec	0.000±0.000	0.610±0.003	0.567±0.002	0.000±0.000	0.072±0.001	0.054±0.001	0.000±0.000	**0.094±0.003**	0.082±0.003
	SASRec with c.i.	0.435±0.015	**0.620±0.002**	0.607±0.001	0.032±0.004	0.088±0.002	0.074±0.002	0.023±0.003	0.088±0.002	0.080±0.002
	SASRec with t.d. (ours)	**0.509±0.005**	0.617±0.002	**0.609±0.002**	0.038±0.008	**0.092±0.002**	**0.078±0.003**	**0.034±0.002**	**0.094±0.002**	**0.087±0.002**
NDCG@10	Content-based KNN	0.312±0.000	0.232±0.000	0.238±0.000	**0.022±0.000**	0.024±0.000	0.024±0.000	0.006±0.000	0.000±0.000	0.001±0.000
	SASRec	0.000±0.000	0.438±0.002	0.407±0.002	0.000±0.000	0.043±0.001	0.033±0.001	0.000±0.000	**0.063±0.001**	**0.055±0.001**
	SASRec with c.i.	0.297±0.010	**0.448±0.001**	0.438±0.000	0.018±0.002	0.053±0.001	0.044±0.001	0.014±0.002	0.058±0.002	0.052±0.001
	SASRec with t.d. (ours)	**0.359±0.000**	0.446±0.002	**0.440±0.002**	**0.022±0.004**	**0.054±0.001**	**0.046±0.002**	**0.021±0.002**	0.060±0.002	**0.055±0.002**

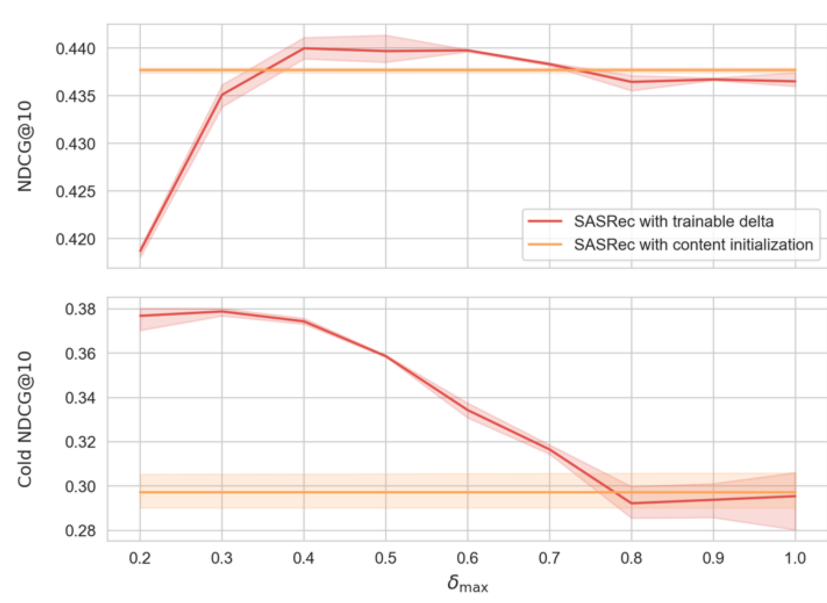

Mean total (top) and cold (bottom) NDCG@10 for SASRec with trainable delta evaluated against δ_{max} on the Amazon-M2 dataset

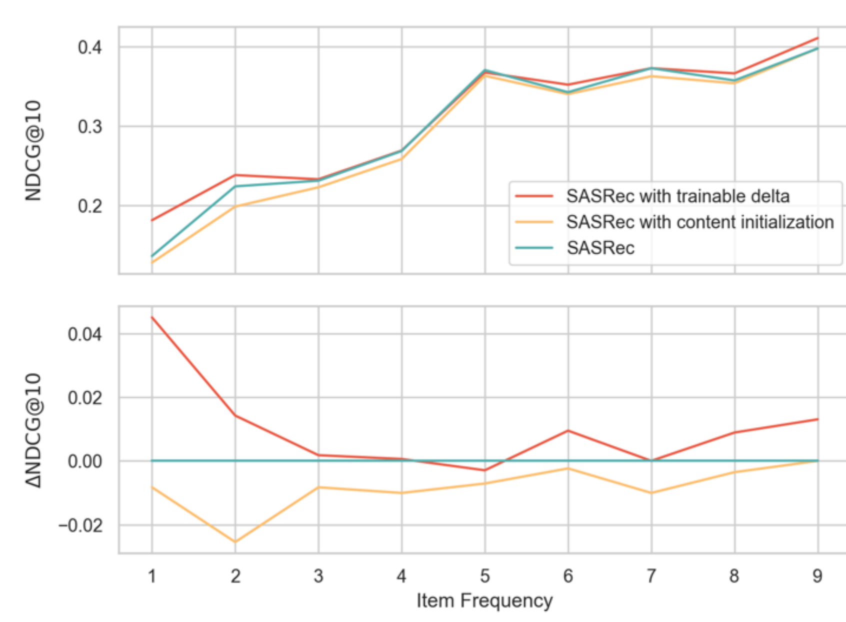

Mean NDCG@10 (top) and NDCG@10 evaluated against the frequency of ground-truth items in the training set on Amazon-M2 dataset

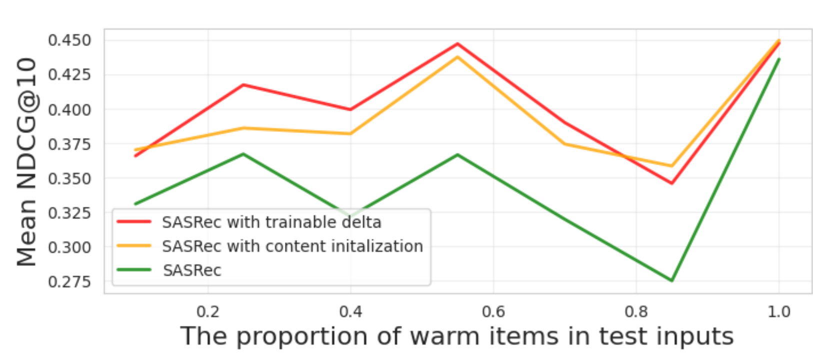

Mean NDCG@10 evaluated across different proportions of warm items in test input sequences on Amazon-M2 dataset

Conclusion

We proposed a simple yet effective method to tackle the item cold start problem in sequential recommendation systems by adding a small trainable delta with bounded norm to frozen content-based embeddings.

Our approach significantly improves recommendation performance on cold items across diverse data modalities, including textual and audio-based item content, while maintaining stable performance on warm items.

apembek@bk.ru
artem42fatkulin@gmail.com
antklen@gmail.com
alexxl.vasilev@yandex.ru

Paper

Repository