

# Breaking the Curse of Quality Saturation with User-Centric Ranking

Zhuokai Zhao\*  
zhuokai@uchicago.edu  
University of Chicago  
Chicago, IL, USA

Yang Yang  
zyyang@meta.com  
Meta AI  
Menlo Park, CA, USA

Wenyu Wang  
owenwang@meta.com  
Meta AI  
Menlo Park, CA, USA

Chihuang Liu  
chihuang@meta.com  
Meta AI  
Menlo Park, CA, USA

Yu Shi  
yushi2@meta.com  
Meta AI  
Menlo Park, CA, USA

Wenjie Hu  
wenjiehu@meta.com  
Meta AI  
Menlo Park, CA, USA

Haotian Zhang  
htzhang@meta.com  
Meta AI  
Menlo Park, CA, USA

Shuang Yang<sup>†</sup>  
shuangyang@meta.com  
Meta AI  
Menlo Park, CA, USA

## ABSTRACT

A key puzzle in search, ads, and recommendation is that the ranking model can only utilize a small portion of the vastly available user interaction data. As a result, increasing data volume, model size, or computation FLOPs will quickly suffer from diminishing returns. We examined this problem and found that one of the root causes may lie in the so-called “item-centric” formulation, which has an unbounded vocabulary and thus uncontrolled model complexity. To mitigate quality saturation, we introduce an alternative formulation named “user-centric ranking”, which is based on a transposed view of the dyadic user-item interaction data. We show that this formulation has a promising scaling property, enabling us to train better-converged models on substantially larger data sets.

## CCS CONCEPTS

• **Information systems** → **Recommender systems; Personalization.**

## KEYWORDS

Ranking; Recommendation Systems; Collaborative Filtering; Deep Learning; User-centric Ranking

## 1 INTRODUCTION

Scaling has been one of the main themes in deep learning and the key driving force behind many eye-opening breakthroughs in the past decade, especially in computer vision (CV) [9, 11, 25], natural language processing (NLP) [2, 6, 8], and multi-modality modeling [21, 22, 29]. In these areas, scaled-up big models were able to improve the corresponding quality metrics by orders of magnitude compared to the state-of-the-art of their previous generations. For example, on ImageNet [7], the ViT [9] model reduced the image classification error rate, compared to the first super-human model ResNet-152 [14], by more than half [9]. This scaling success, however, has not yet happened in ranking (e.g., search, ads, recommendation systems). This seems both surprising and mysterious given that ranking represents an important aspect of the AI industry.

In a typical scaling scenario, one important condition is that the model should have the capability to utilize more data, so that

increasing data volume and computing will continue to improve model quality. When it comes to ranking, we notice that even with an abundant or even infinite amount of data (i.e., massive user engagement activities constantly accumulating in systems like Google ads, Facebook news feed, YouTube video recommendation, etc.), the ranking models typically can only utilize a small portion (i.e., a few days to a few weeks of logged data). Increasing training data volume, model size, or computation FLOPs can only lead to very little quality improvement. This is known as the “quality saturation” problem.

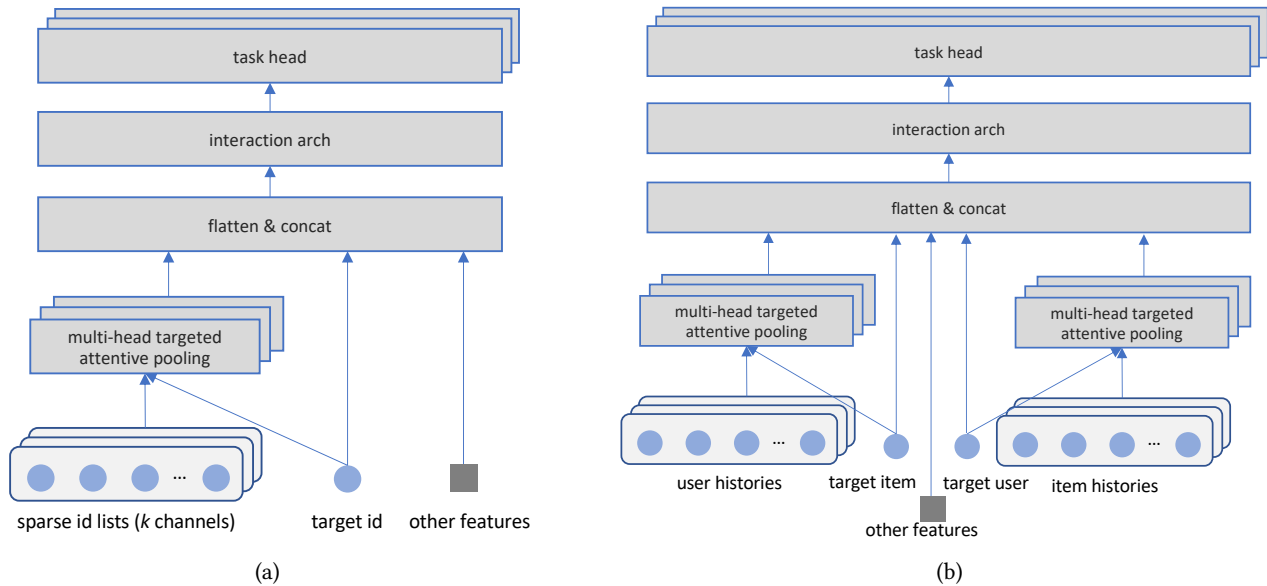
To be fair, the quality of every machine learning model will eventually saturate, sooner or later. What makes it unique in ranking is that the quality saturation happens too soon. Considering the important role that ranking models play and their business impact, a reasonable expectation is that a ranking model should be able to utilize at least a few months of training data.

We examined this problem and found that one of the root causes may lie in the formulation. With an analogy to NLP, the current ranking formulation predicts dyadic responses (e.g., ads click-through) by casting ‘items’ as ‘tokens’ and ‘users’ as ‘documents’, a paradigm called “item-centric ranking”. This is actually an ill-posed formulation because the model size or the number of parameters to learn will grow linearly as data volume increases. As a remedy, we introduce an alternative formulation called “user-centric ranking” based on a transposed view, which casts ‘users’ as ‘tokens’ and ‘items’ as ‘documents’ instead. We show that this formulation has a number of advantages and shows less sign of quality saturation when trained on substantially larger data sets.

The proposed methods have been tested in a variety of our production systems with significant metric wins, including search, ads, and recommendation. These systems are quite diverse in nature (e.g., different interaction interfaces, items of very different types) and can be regarded as representative of many ranking systems in the industry, yet our findings are quite consistent. Our reported experiment results are primarily based on one production surface, which has 6 different tasks (including both positive and negative engagements, and both immediate and deferred reward feedback), and the comparison and trend are consistent across all these tasks. In addition to offline results, we also report online live experiment results. Furthermore, to improve the reproducibility of our findings, we also include results on a public data set and plan to open-source our implementation code for public access.

\*Work done during an internship at Meta.

<sup>†</sup>Lead author; part of the work was done at NewsBreak.



**Figure 1: (a) An example of one-tower ranking model; (b) A hybrid ranking model containing both a user-centric and an item-centric sub-architecture.**

## 2 RELATED WORK

The past decade has witnessed tremendous successes achieved by deep learning models that are growing in scale exponentially over time. In computer vision, big model architectures have been widely used for image classification and object detection tasks. The neural architectures have evolved from Convolutional Neural Networks (CNNs) with a handful of layers [19], to ResNet who has more than 100 layers and 100 million parameters [14], to recent gigantic Transformer-based models that contain hundreds of billions of parameters [9]. The trend is even more prominent in NLP, especially in the few years of post-BERT era [8, 24]. A surge of state-of-the-art models are emerging with ever growing sizes, complexities, and new levels of capabilities, e.g. GPT-3 and GLaM [10] are among the largest language models to date and have demonstrated impressive performance in various NLP tasks [2, 6].

It is a bit surprising that, unlike the other areas, scaling has not gained much success in ranking, even though it is the biggest industry for AI and there is no shortage of training data [12]. Ranking models used to be dominated by the “two-tower” architectures, where the user-side and the item-side were modeled independently with separate architectures in the early stage known as the two towers; and fusion or interaction between the two sides happens at a relative late stage [5, 15, 18]. Recently, “single-tower” architectures based on Transformer emerged and quickly became the new state of the art [24, 30]. However, compared to other areas, these models are notably simpler, for example, they are using only a single (or a few, if Transformer is also used in interaction sub-arch) layer of Transformer block, and even though these models could be big in size (e.g. 1 trillion parameters), the majority of the parameters are

sparse-id based embeddings, only a tiny fraction of which are active for each prediction.

The current common practice in ranking is to model each user based on the sequence of historically interacted items. The representation of user interests can be learned from historical behaviors, and the likelihood of a potential engagement is assessed based on the affinity of the target item with respect to historical interactions. These models provide an item-centric perspective to utilize the dyadic user-item interaction data; we call it item-centric because learnable embeddings are allocated for items but not users. We show that this formulation could be the cause of quality saturation in ranking. The proposed user-centric ranking is the first to provide an alternative formulation based on a transposed view of the dyadic interactions. We show that it can help to alleviate quality saturation in ranking. We want to note that our contribution is to introduce this new formulation, not a specific neural architecture. These two are orthogonal, in fact, any SoTA item-centric ranking model can be converted to its user-centric counterpart using the new formulation.

It is important to capture the complex relationships between users and items to improve ranking accuracy in ranking systems. Using user information corresponding to a target item is a natural choice. One example is graph-based recommendation models [4, 27], which represents users and items as nodes in a bipartite graph. The graph model learns to generate user and item embeddings for recommendation through the process of embedding, propagation, and prediction. Our approach of user-centric ranking models user-item interaction in a different way and targets for replacing or complementing the current item-centric ranking models that suffer from quality saturation. There are other attempts to alleviate the changing inventory problem, such as meta learning approaches [3, 28]. The goal of meta learning for ranking is to improve robustness

and/or fairness of ranking models caused by unintended data biases. In contrast, we aim to address the quality saturation problem caused by inventory dynamics.

### 3 RANKING FORMULATIONS

In ranking, we are concerned with modeling *dyadic responses*. Given a set of users  $\mathcal{U}$  and a set of items  $\mathcal{I}$ , the goal is to predict  $y_t(u, i)$  for any given user  $u \in \mathcal{U}$  and item  $i \in \mathcal{I}$  at time  $t$ . In different contexts,  $y$  can have different semantic meanings, e.g., click-through of an ad, conversion of a transaction, following an account, or finishing watching a video. Ranking models are trained on historical interaction data in the format of  $\mathcal{D} = \{(u, i, t, y)\}$ , which can be thought of as a bipartite graph between  $\mathcal{U}$  and  $\mathcal{I}$ .

An interesting note is that ranking bears a lot of similarities with NLP, because NLP data can be thought of as dyadic interactions between ‘documents’ and ‘tokens’. In fact, a lot of ranking techniques are inspired by progresses in NLP [17, 23, 30].

#### 3.1 Item-Centric Ranking

Figure 1 shows one example of single-tower item-centric architectures. The key idea, with an analogy to NLP, is to think of items as tokens and users as documents, i.e., each user is modeled by a list of items that they engaged with, in chronological order according to the time of engagements. When multiple types of engagements are involved (e.g., in video recommendations, engagements could include clicks, video completion, likes, follow-author, etc.), they can be organized into multiple channels, one for each engagement type.

For each channel, items in the engagement history are first mapped to their embeddings, positions are encoded based on relative time-stamps, and multi-head attentions are applied on top. The aggregation output is then concatenated with all other features, on top of which an interaction sub-architecture (e.g., Deep & Cross Network (DCN) [26] or self-attention [24]) is employed to encode higher-order nonlinear interactions among different feature groups. And finally, a number of task heads (e.g., one MLP for each engagement prediction task) provide the output probabilities. Because of the daunting scale in ranking, these ranking architectures are highly-simplified versions compared to what are commonly used in NLP, noticeably: 1) only one layer of attention is typically used; 2) instead of full-sized self-attention, the aggregation is based on the so-called “targeted attentive pooling”, i.e., when predicting  $y_t(u, i)$ , the engagement history of user  $u$  is aggregated by attending only w.r.t. the target item  $i$  (i.e., the embedding of item  $i$  is used as query in the attention function). The latter is similar to document/paragraph representation in NLP, where the aggregation is by attending to the special symbol ‘CLS’.

This formulation is called “Item-Centric Ranking” (ICR) to reflect that items are allocated free-parameter embeddings to be learned in training whereas user embeddings are derived by aggregating item embeddings.

#### 3.2 User-Centric Ranking

Why do ranking models saturate so fast? Why doesn’t this happen to NLP models given that they bear lots of similarities? When we carefully compare these two settings, we notice an important

difference. In NLP, the vocabulary size (i.e., total number of tokens) is often fixed; given a neural architecture, the number of parameters is constant when we increase the training data. This is, however, not the case in ranking when item-centric formulation is used.

In particular, especially in the so-called “creator economy”, where the inventory of items are highly dynamic: new items are being created constantly (e.g., tens of millions of posts/videos are created on Facebook/Instagram every day) and items are time-sensitive and ephemeral (e.g., each post/video has a short life-span ranging from a few days to a few weeks). In this setting, because the item inventory grows linearly over time  $|\mathcal{I}| = O(t)$ , for any given neural architecture, the number of model parameters will grow unboundedly in  $O(t)$  (due to the use of per-item embeddings). As a result, when we increase the training data (e.g., to use more days of logged interactions), because of the linear growth in model size, the per-parameter data density will not grow, and hence using more data will not make the model converge better (e.g., lower the variance). In fact, this is a setting that we rarely see elsewhere.

Based on this observation, we propose an alternative formulation called “User-Centric Ranking” (UCR), which is based on a transposed view of the user-item interactions. Using the NLP analogy again, UCR casts ‘users’ as ‘tokens’ and ‘items’ as ‘documents’; free-parameter embeddings are learned for users, and item embeddings are derived by aggregation. For mature ranking systems in double-sided markets, it is typical to see an increase in inventory, while the user set  $\mathcal{U}$  remains relatively consistent; thus, the model size (i.e., the number of parameters) of these ranking systems will stay stable as we increase training data. Our expectation is that with this formulation, when we scale up training data the consistent growth of per-parameter data density should translate to better model convergence.

In a typical setting where user set is capped while both the inventory size and the training data set size grow linearly over time, it can be shown the asymptotic error rate (i.e., the expected distance between the optimal value of model parameter  $\theta^*$  and its actual value  $\hat{\theta}$ ) for each of the formulations is as follows [20]:

- Item-centric ranking:  $\mathbb{E}[|\theta^* - \hat{\theta}_t|^2] = \text{Const}$
- User-centric ranking:  $\mathbb{E}[|\theta^* - \hat{\theta}_t|^2] = O(\frac{1}{t})$

As training data grow, asymptotically UCR converges at a sublinear rate (at most), while ICR cannot be improved further, which explains the quality saturation we have observed.

From an intuitive perspective, UCR could be advantageous over ICR. In ICR, because items are ephemeral, so are their embeddings (i.e., an item embedding will soon become irrelevant and useless as that item exits the system). In UCR, we are continuously accumulating and improving our knowledge about every user by refining its embedding over time as long as that user keeps on interacting with the system.

Any SoTA item-centric ranking model can be converted to its user-centric counterpart using the new formulation. Note that the example architecture in Figure 1(a) applies to both item-centric and user-centric. The key difference is whether users or items are used as keys for embedding look-ups (i.e., the ‘sparse-id’ and ‘target-id’ in the figure).

### 3.3 Hybrid Models

It is also possible and actually straightforward to have a hybrid formulation, i.e., to implement models that include both a user-centric and an item-centric attentive pooling components. Figure 1(b) shows how the example architecture in Figure 1(a) looks like in the hybrid formulation. Such hybrid models will have similar “parameter explosion” problem as item-centric models. We will compare all these different model formulations in our experiments.

## 4 IMPLEMENTATION

### 4.1 Item-Centric Ranking

Item-centric id-lists represent the engagement history of each user. Although the number of items that one user can interact within one day is hardly over a few hundreds, the list of distinctive items and their embeddings gets accumulated very quickly over time, especially considering that the same item is rarely recommended to the same user again. A sampling strategy is needed in order for each engagement list to not exceed certain length. In our implementation, we limit the length to 1024 at max, by only including the most recent engagements. In our experiment, this method is referred to as “IC-Sampling”.

### 4.2 User-Centric Ranking

One of the challenges for implementing UCR is to handle the distribution skewness. In an item-centric setting, the number of items one user can interact with tends to be evenly distributed (e.g., daily engagements range from a few to a few hundred), whereas in the new setting, the distribution is more irregular, e.g., some items can attract millions of users to engage with while others can get only a few. This means that for some items it is no longer feasible to fit the entire list of engaged users in memory during training/inference. We explore three different approaches:

- **Sampling.** In this implementation, we simply down-sample the list of engaged users of an item to a fixed-size sub-list uniformly using reservoir sampling. Note that in practice, if we sample for each item only once, instead of resampling for each user-item interaction, this will introduce an artificial bias. This method is referred to as “UC-Sampling”.
- **Aggregation.** Another approach is to summarize a long sequence of engaged users to a shorter list, e.g., by clustering the users and using cluster-id in replacement of user-id. In our implementation, the clusters are obtained by applying the Louvain algorithm [1] to the user-item interaction graph. Our in-house implementation provides the functionality to incrementally update the clustering structure over time with constraints on cluster size and re-mapping ratio. This method is referred to as “UC-Clustering”.
- **Retrieval.** Alternatively, we can pre-index the engagement history and use retrieval (e.g., max inner-product search) to identify the subset of most relevant users (w.r.t. the target user), on which attentive pooling is then applied. Since attention is of quadratic complexity, the overhead of retrieval can be compensated by the speedup due to a shorter and more selective attention window. A sparsified attention distribution also means an improved signal-to-noise

ratio (i.e., long-tail less relevant candidates are pruned and excluded from the attentive aggregation) and can further improve model quality. We leave this method for future investigation.

Note that this problem is only a concern for a very small subset of the most popular items, for which most ranking models already have good prediction accuracy. For the vast majority of items in our case, the engagement users are below the 1024 length limit.

### 4.3 Parameter Hashing

Another technical challenge is memory management when working with large-scale ID spaces such as user-ids  $\mathcal{U}$  and item-ids  $\mathcal{I}$ . Considering that we are learning embedding vectors, one for each distinctive ID, the extremely large cardinalities (i.e., in the order of billions) of these ID spaces imply that the memory requirement as well as the index to map IDs to their address can be quite a challenge. Especially for item-centric ranking, the number of item IDs can grow unboundedly to infinity.

One common approach to address this problem is to implement feature hashing, i.e., to maintain a constant hash space for these IDs and allocate one embedding vector for each distinctive “hashed ID”. This is of course not ideal. The existence of hash collisions means that we are forcing certain random IDs to share the same embedding vectors. This is not necessarily a bad thing when the collision rate is at a reasonable level, because feature hashing provides a type of regularization effect to the embedding parameters similar to dropout. However, for unbounded ID spaces such as  $\mathcal{I}$  in user-centric ranking, the collision rate is expected to grow linearly over time (i.e.,  $O(t)$ ), and can be arbitrarily large and no longer negligible. In contrast, in user-centric ranking, the ID space  $\mathcal{U}$  is bounded and hence collision rate is under control.

### 4.4 Aggregation Operators

We implement two aggregation operators, sum-pooling and targeted attentive pooling. The former aggregates the list of associated IDs by the sum or mean of their corresponding embeddings. Sum-pooling is computationally inexpensive and easy to implement. However, it has very limited expressive capability (e.g., the operator itself is parameter-less) and needs to rely on the interaction arch to encode complex interactions. Moreover, especially when the list is long, using an unweighted sum could deteriorate the signal-to-noise ratio and make the prediction less accurate. By attending to the target user (item), attentive pooling can adaptively adjust how much weight an embedding could get based on not only the relevancy of the current item (user) at hand but also the relevance of other competing entities. This aggregation is especially powerful when the list contains entities of diverse topics (e.g., a user’s engagement history could contain items in different categories), for which the multiple distribution modes would be inevitably collapsed into one if sum-pooling is used. Attentive pooling is also more robust and tolerant to noises, outliers or corruptions in the ID list.

**Table 1: Evaluation results (AUC) on MovieLens data.**

	ICR	UCR	Hybrid
DIN with Attentive Pooling	0.712	0.731	0.737

## 5 EXPERIMENTS

### 5.1 On Public Data

A major goal of this paper is to improve the scaling capability of ranking models due to the curse of quality saturation caused by growing item inventories. To test our findings, data sets need to be both (1) substantially large-scale and (2) based on dynamic inventory as in real-world systems. Unfortunately, public data cannot meet the requirement: they do not have the desired scale, nor do they have the needed dynamics (matrix completion settings with fixed users & items). We notice that this is a common issue in the community. Notably, recent works on scaling, including those in NLP and CV are based on dedicated data sets. The matter is even worse in the area of ranking, because published data is not only too small in scale but also lacks many vital characteristics that real-world systems possess, making findings on such toy data sets less reliable when being generalized to real world. However, to improve the reproducibility of our results, we tested our methods on one public data set for demonstration purposes.

**5.1.1 Data.** The MovieLens-20M data set is a popular benchmark in recommendation systems [13]. It contains 20-million ratings from 138,493 users on 27,278 movies. In our experiments, we follow a protocol similar to that of [30]: ratings of 4-star or above are treated as positive and the rest as negative; for each user, the most recent  $N$  ( $N = 512$ ) positively-rated movies are used as item-centric channels of that user; similarly, the  $M$  ( $M = 512$ ) users who historically rated a movie positively are used as user-centric channels of that movie. As we mainly compare the difference between ICR and UCR, we do not include other categorical features, such as genre.

**5.1.2 Results.** We tested the DIN [30] architecture (Figure 1(a)) in the three different formulations (i.e. ICR, UCR, hybrid) with ‘Attention-pooling’ as aggregation operator. A 4:1 split is used for training and testing. The evaluation results in terms of AUC (i.e. area under ROC curve) are reported in Table 1.

Note that MovieLens is a static data set. It does not have the inventory dynamics that real-world systems have, and hence we will not be able to see parameter explosion on this data set. From Table 1, our observation is that UCR is at least on par with or slightly better than ICR, while hybrid performs the best possibly because it uses more signals than either of them.

### 5.2 On Real-World Production Data

**5.2.1 Data.** We further experiment on real-world production data. For offline evaluation, we created a ‘lab data set’ by sampling the production log of a real-world short-form video recommendation system. Our data set contains about 24 million users and their engagement activities in the time range of 60 days (from late July to early October of 2022). In total, the data set contains about 28 billion

examples (engagement activities) involving 1 type of negative and 5 types of positive engagements.

**5.2.2 Metric.** We use Normalized Cross-Entropy (NCE) as the primary evaluation metric [16]. NCE is defined as the cross-entropy loss of the model prediction  $p$  normalized by the entropy of the label  $y$ .


$$NCE(p, y) = \frac{\text{CrossEntropy}(p, y)}{\text{Entropy}(y)} \quad (1)$$

NCE is widely used as the gold standard offline metric for engagement probability (e.g. CTR) prediction tasks because of its high consistency with online engagement metrics.

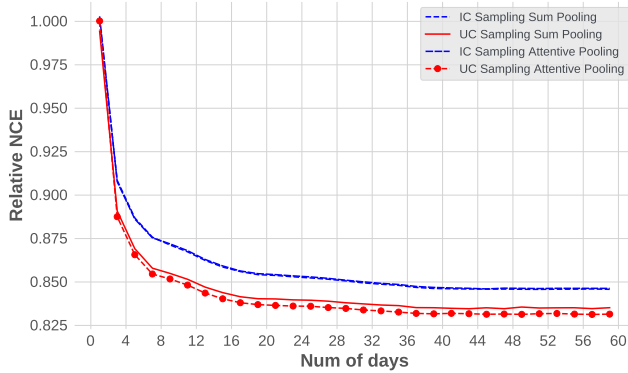
**5.2.3 Parameter Growth.** In both ICR and UCR, the total number of parameters that a model has can be expressed as  $const + n \times d$ , where the constant part is mostly related to model architectures, while  $n$  and  $d$  denote the total number of distinctive sparse-ids and the dimensionality of each embedding vector. In our data set, as is common in most ranking systems, the cardinality of the user set tends to be bigger than that of the item set for any given day,  $|\mathcal{U}| > |\mathcal{I}_{t+1}| - |\mathcal{I}_t|$ , where  $\mathcal{I}_t$  is the accumulative item set on day  $t$ . However, that comparison is quickly reversed as time goes by because  $|\mathcal{I}_t|$  grows linearly in  $O(t)$ .

Figure 2 shows the model size growth over time for both ICR and UCR models. We only plotted the curves for the case with sampling and attentive pooling, but the trend is similar for all other variants. While it is true that for the first few days the ICR model has fewer parameters, it constantly adds parameters every day as new item IDs emerge. As a result, the ICR model size grows almost linearly over time. In contrast, the UCR model, although has a bit more parameters initially, the model size stays relatively stable over time.

Considering these two models are trained using the same amount of dyadic interaction data, the drastic contrast of the parameter growth can have profound impacts on model quality. For example, at the end of the 60-day window, the ICR model is 21x larger in size than its UCR counterpart. This means that ICR consumes 21x more memory, or when parameter hashing is used the collision



**Figure 2: The growths of model size (the total number of parameters) over time for ICR and UCR models.**



**Figure 3: Comparison of ICR and UCR models in offline evaluation. Models are trained recurrently on a daily basis and evaluated on future 10K activities using NCE. (lower is better)**

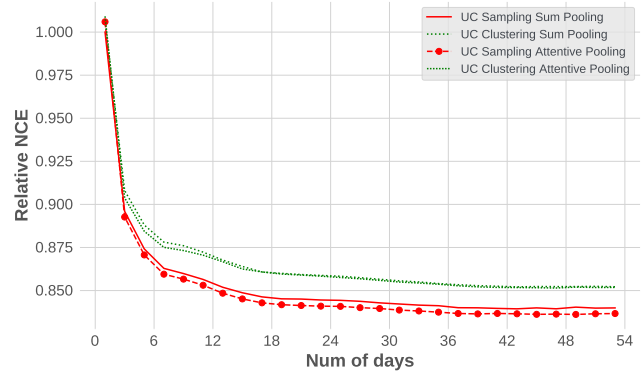
rate is 21x higher; at the same time, on average, each ID embedding receives 21x less training data in ICR as compared to in UCR.

**5.2.4 ICR vs. UCR.** We compare IC-Sampling and UC-Sampling with the two aggregation operator options. All the models are trained recurrently and evaluated on a daily basis using the first ~10K examples of the next day. Because we have 6 tasks (and correspondingly 6 engagement history channels) in our data set, each task (and the engagement channel) is evaluated independently. The results are reported in Figure 3, where only the results on ‘Task 1’ are shown (results on other tasks are very similar); all the NCE numbers are normalized by the NCE of the IC-Sampling sum pooling model on day 1, and relative NCEs are used in the plot.

We can observe that UC-Sampling demonstrates a clear gain over IC-Sampling, with the gap increasing rapidly from day 1 to day 10, and then slowly converging till the end. The performance matches our hypothesis that UCR accumulates and refines the understanding of each user, which helps with better recommendations as the data scales up. However, we did not notice the gain increase through the end of the experiments. We believe that this is because UCR excels more on active users due to its nature of aggregating user embeddings to profile engaged items, but falls short on less active users. We will come back to address more about this issue in Section 5.2.9.

**5.2.5 Sum Pooling vs. Attentive Pooling.** We also compare the impact of the two aggregation operators in ICR and UCR. As shown in Figure 3, attentive pooling consistently performs better than sum pooling in UCR. With more data, the gap is also increasing. After 60 days of training, UCR attentive pooling get 0.44% gain over the sum pooling alternative. In contrast, the advantage of attentive pooling in ICR is very minimal.

This also proves our hypothesis in Section 4.4. In ICR, the item ID is not well trained due to the linearly increased ID space. As a result the attention score between history item and target item does not learn useful signals, and attentive pooling falls back to mean (sum) pooling. In UCR, user ID space is stable, and all ID embeddings




**Figure 4: Comparison of the two implementation methods for UCR: sampling vs clustering.**

could be optimized. This finding verifies the potential to solve the quality saturation problem using UCR with more training data.

**5.2.6 Sampling vs. Clustering.** In UCR, one of the key aspects to ensure good performance is to construct better and more representative engaged user lists for each item, especially for those extremely popular items that gain millions of user interactions. We implemented two of the approaches presented in Section 4.2, namely UC-Sampling and UC-Clustering. Figure 4 shows the comparison between these two approaches. As can be seen, UC-Sampling seems to dominate UC-Clustering in terms of NCE consistently across the entire time span and all the tasks involved. We want to point out that this may not be definite as the performance highly depends on the choice of implementation, e.g., the incremental Louvain algorithm [1] used in our experiments. If a better algorithm is used, the result can be different. We leave such investigation for future research.

**5.2.7 Hybrid Method.** We also compare the hybrid method with UCR and ICR. Because the consistently superior performance of



**Figure 5: Comparison of the hybrid model with its UCR and ICR counterparts.**


**Table 2: Multi-task relative NCE percentage (%) change between ICR (baseline), UCR and Hybrid models implemented with attention pooling. Baseline setting is denoted as “-”.**

Task	Day 7			Day 14			Day 30			Day 60		
	IC	UC	Hybrid	IC	UC	Hybrid	IC	UC	Hybrid	IC	UC	Hybrid
1	-	-2.58	<b>-2.88</b>	-1.73	-4.01	<b>-4.32</b>	-3.01	-4.90	<b>-5.21</b>	-3.48	-5.18	<b>-5.31</b>
2	-	-2.71	<b>-2.94</b>	-0.46	-3.23	<b>-3.42</b>	-0.45	-3.24	<b>-3.44</b>	-0.44	-3.19	<b>-3.32</b>
3	-	+1.84	<b>-2.04</b>	-7.52	-7.60	<b>-10.38</b>	-10.96	-12.64	<b>-14.29</b>	-11.98	<b>-13.98</b>	-12.78
4	-	-2.86	<b>-3.08</b>	-0.31	-3.23	<b>-3.41</b>	-0.08	-3.03	<b>-3.23</b>	-0.05	-2.96	<b>-3.08</b>
5	-	-2.88	<b>-3.14</b>	-0.66	-3.61	<b>-3.88</b>	-0.79	-3.77	<b>-4.00</b>	-0.81	-3.73	<b>-3.88</b>
6	-	-3.09	<b>-3.28</b>	-0.86	-3.97	<b>-4.14</b>	-1.19	-4.34	<b>-4.53</b>	-1.28	-4.38	<b>-4.47</b>

sampling over clustering as reported before, we only experimented with the sampling implementation. The results are shown in Figure 5. It seems that the hybrid method has very similar performance as the UCR counterpart, albeit slightly better. This phenomenon is pretty consistent. We observe that the hybrid method achieves the best NCE results across all the tasks. Considering that the hybrid architecture, as shown in Figure 1, includes both an UCR sparse sub-arch and an ICR sparse sub-arch, the results are partly as expected (i.e., it should have the advantages of both UCR and ICR) and partly surprising (i.e., it has the same parameter explosion problem as ICR).

**5.2.8 Multi-Task Evaluation.** In our previous evaluations, we use 1 single task and 1 single engagement history channel. In this section, for both ICR and UCR, we use all the available engagement signal channels (one for each engagement type) and jointly train the model on all of the 6 tasks. This multi-channel and multi-task setting allows the model to capture correlations among different tasks as well as between the signal channel and the task loss corresponding to different engagement types, which cannot be done in the previous setting. The results are reported in Table 2, where the NCE is calculated relative to the NCE of the ICR model at day 7. We observe that overall UCR models show clear gains when compared to ICR counterparts across all the tasks; moreover, the hybrid model consistently performs the best at all of the tasks, although the difference with the UCR models is very marginal.

**5.2.9 Segment Analysis.** We segment users into five buckets based on their activeness (e.g., number of engagements within a given

**Figure 6: Distribution of NCE gains over ICR on different user activeness segments (negative means better).**

time window). In Figure 6(a), we show the NCE differences between one UCR model (UC-Sampling) and one ICR model (IC-Sampling) for each user segment. We can see that, although UCR performs better than ICR overall, the gain mostly come from more active users. For less active users (e.g., engagement counts < 10), UCR actually performs worse than the ICR baseline. This explains why the hybrid methods tend to perform the best because it leverages both components to provide the better of the two worlds. As a validation, Figure 6(b) shows the similar analysis of the Hybrid model over ICR, and we can see it provides gains across all the user segments.

**5.2.10 Ablation Study.** To better understand how different configurations impact model performance, we conduct a set of parameter sweep experiments. For this analysis, we set the number of training data to be 30 days for all the runs. In addition, we use IC-Sampling and UC-Sampling with the same single-task setting in our experiments.

**Hash Size.** Parameter hashing maps user IDs or item IDs to embedding vectors by applying a hash function. Though being space-efficient, it is essential to have a large enough hash space so that a high collision rate between these IDs can be avoided. In this experiment, we further examined how hash size affects model performance by varying it from the default value of 20 million. As hash size affects both IC and UC ranking, we test both IC-Sampling and UC-Sampling as well as using both sum pooling and attentive pooling model architectures. The results are reported in Table 3. Overall, increasing the hash size leads to a better model performance. This trend is more evident for UCR. For example, increasing the hash size from 1M to 30M for UC-Attn results in a 1.71% reduction in relative NCE. One reason why UCR benefits more than ICR is that UCR has much fewer embedding vectors, the reduction in hash collision is more dramatic for UCR when increasing hash size.

**Table 3: Relative NCE percentage (%) change from different models with varying hash sizes. Baseline setting is denoted as “-”.**

	1M	5M	10M	20M	30M
IC Sum	+0.08	+0.04	+0.01	-	+0.02
IC Attn	+0.12	+0.05	+0.05	+0.07	+0.06
UC Sum	-0.04	-0.73	-1.07	-1.43	<b>-1.53</b>
UC Attn	-0.24	-1.13	-1.48	-1.84	<b>-1.95</b>

**Table 4: Relative NCE percentage (%) change from different models with varying feature dimensions. Baseline setting is denoted as “-”.**

	96	192	384
IC-Sampling	-0.05	-	+0.01
UC-Sampling	-1.39	-1.91	-2.37

**Embedding Dimensionality.** We conduct another ablation study on the dimensionality of the embedding vectors. Our default embedding dimension is 192, and we tune it between 96 and 384. Results are illustrated in Table 4. We can see that IC-Sampling is not able to utilize a larger embedding dimension, and its performance is worse when the largest dimensionality is used. On the other hand, UC-Sampling shows consistent improvements when higher dimensional embeddings are used.

### 5.3 Online Results


Based on the encouraging results on the sampled lab data, we took the step forward to productionize the proposed techniques in our recommendation system. **On the full-scale production data, we observed up to 0.6% NCE gains** compared to the production ICR model when UCR models were trained with the standard workflow using a few days of training data without any architecture changes. The best version was then tested live in the production system.

A number of infrastructure optimizations were done to make this happen. For example, we optimize the batching algorithm to put the same user’s data in one batch for ICR, so the sum (attention) pooling of the item-centric features only needs to be computed once and then could be shared within the batch. For UCR, we do the similar operation to batch the same video’s data together. With the improvement on data locality, we can lower down the memory consumption, and in turn improve the throughput for both training and serving. Also, by using full-precision for training and lower-precision (e.g., FP16) for inference, we were able to improve the inference performance (both throughput and latency) without significant regression in prediction quality (e.g., NCE) and reduce the number of GPUs required for serving by almost half. The online A/B experiments showed that quite significant wins were achieved across a wide range of topline metrics, in particular, one of the key business metrics, **video watch time was improved by 3.24%**.

An important observation during our productionization process is that the offline NCE gain can be further enlarged when we increase the amount of training data. In addition, if we scale up both training data and model complexity, we could potentially obtain an outsized gain in terms of NCE in offline evaluation. This investigation is currently in progress.

### 5.4 Open Questions and Discussions

We are motivated to address the quality saturation problem in ranking. Our expectation is that the UCR formulation should provide somewhat a remedy. However, from our experiment results, this is only partially validated. In particular, we did see UCR models lead



**Figure 7: Prediction quality (NCE) of pre-trained models over the next 24 hours indicates there is a strong distribution drift in the data.**

to consistently better NCE than their ICR counterparts; we also saw a tendency of improving NCE gain as we increase the training data. Nonetheless, the NCE gap between UCR and ICR is not as big as we expected, and also that gap is being enlarged at a much slower speed, far too slow if we compare it with the model parameter or collision rate growth curves. This is kind of surprising.

In an attempt to understand the discrepancies, we have a few plausible explanations.

Firstly, we notice there’s a nontrivial discrepancy between the full-scaled production data and our sampled lab data. The scaling characteristics of UCR models are significantly better on production data than what we observed. This is partly related to the sampling algorithm we used to generate this data set, and partly related to the nonlinearity between the complexity that the data manifests and the scale at which the problem is examined.

Secondly, in the aforementioned areas where scaling has led to tremendous success, including CV and NLP, the concepts we try to model are often static. In other words, there’s usually a ground-truth model in hindsight and the goal of training is to approach that ground-truth. However, in ranking it is fundamentally different. There is drastic and frequent distribution drift due to the highly dynamic two-sided ecosystem and the interactive highly counterfactual nature of the engagement process. Because of the distribution drift, there is no ground-truth model (or you could say the optimal model is a moving target instead of static). For example, Figure 7 shows how a pre-trained static model performs in the next 24 hours after it was trained. We can see a very significant deterioration of the prediction NCE as the model becomes increasingly outdated. In a situation where the distribution is drifting dynamically, a model that scales well and does not saturate quickly in a static context may not always scale well. To fully combat the obstacles for scaling ranking models, deep understanding of and the ability to control such dynamics are critical.

Last but not the least, our current study is limited, without any changes to the model architecture. We observed, especially for the smaller-scale lab data set, the absolute NCE values are quite small and may be close to their limits for the architecture we used. At



the same time, we noticed that ranking model's architectures are significantly simpler than what are commonly used in NLP and CV, which is of course a practical choice given the scales in ranking. We believe that by using significantly more expressive architectures, we will be able to improve the scaling property further.

We leave these investigations for future study.

## 6 SUMMARY

We suspected that the item-centric formulation of ranking models may be contributing to the quality saturation problems. We introduced user-centric ranking as an alternative formulation. We showed that in general, UCR models have a stable model size (i.e., total number of parameters) that will not grow as we increase training data. On a lab data set of sampled production data, we observed that UCR models yield consistently better prediction quality and have slightly better scaling property. We did not believe that this fundamental problem in ranking has been fully solved. We listed a number of open problems from our study and hope they can spark further investigations.

## ACKNOWLEDGMENTS

We would like to thank the following individuals from Meta for the collaboration and support: Pei Yin, Hui Zhang, Jason Liu, Xianjie Chen, Mingze Gao, Jiyan Yang, Hitesh Kumar, Mert Terzihan, Nathan Berrebbi, Liang Xiong, Jiaqi Zhai, Shilin Ding, Shuang Yang is grateful to Jeff Zheng and Junhua Wang from Newsbreak for many helpful discussions.

## REFERENCES

- [1] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefevre. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* 2008, 10 (2008), P10008.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [3] Vitor R Carvalho, Jonathan L Elsas, William W Cohen, and Jaime G Carbonell. 2008. A meta-learning approach for robust rank learning. In *SIGIR 2008 workshop on learning to rank for information retrieval*, Vol. 1.
- [4] Lei Chen, Le Wu, Richang Hong, Kun Zhang, and Meng Wang. 2020. Revisiting graph based collaborative filtering: A linear residual graph convolutional network approach. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34, 27–34.
- [5] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishu Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*, 7–10.
- [6] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311* (2022).
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2021).
- [10] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. 2022. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*. PMLR, 5547–5569.
- [11] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. 2022. Masked Autoencoders As Spatiotemporal Learners. *arXiv preprint arXiv:2205.09113* (2022).
- [12] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. 2019. Are We Really Making Much Progress? A Worrying Analysis of Recent Neural Recommendation Approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems* (Copenhagen, Denmark) (RecSys '19). Association for Computing Machinery, New York, NY, USA, 101–109. <https://doi.org/10.1145/3298689.3347058>
- [13] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (dec 2015), 19 pages. <https://doi.org/10.1145/2827872>
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [15] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [16] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, and Joaquin Quiñero Candela. 2014. Practical Lessons from Predicting Clicks on Ads at Facebook. In *ADKDD'14*.
- [17] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [18] 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*. 2333–2338.
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger (Eds.), Vol. 25. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>
- [20] Lam Nguyen, Phuong Ha Nguyen, Marten Dijk, Peter Richtárik, Katya Scheinberg, and Martin Takáč. 2018. SGD and Hogwild! convergence without the bounded gradients assumption. In *International Conference on Machine Learning*. PMLR, 3750–3758.
- [21] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [22] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*. PMLR, 8821–8831.
- [23] 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*. 1441–1450.
- [24] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [25] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. 2022. GIT: A Generative Image-to-text Transformer for Vision and Language. *arXiv preprint arXiv:2205.14100* (2022).
- [26] Ruoxi Wang, Rakesh Shivanna, Derek Cheng, Sagar Jain, Dong Lin, Lichan Hong, and Ed Chi. 2021. Dcn v2: Improved deep & cross network and practical lessons for web-scale learning to rank systems. In *Proceedings of the Web Conference 2021*. 1785–1797.
- [27] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *Proceedings of the 42nd international ACM SIGIR conference on Research and development in Information Retrieval*. 165–174.
- [28] Yuan Wang, Zhiqiang Tao, and Yi Fang. 2022. A Meta-learning Approach to Fair Ranking. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2539–2544.
- [29] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. 2022. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917* (2022).
- [30] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1059–1068.