

CAM2: Conformity-Aware Multi-Task Ranking Model for Large-Scale Recommender Systems

Ameya Raul*
Meta Inc.
Menlo Park, USA
araul@meta.com

Amey Porobo Dharwadker*
Meta Inc.
Menlo Park, USA
ameydar@meta.com

Brad Schumitsch
Meta Inc.
Menlo Park, USA
bschumitsch@meta.com

ABSTRACT

Learning large-scale industrial recommender system models by fitting them to historical user interaction data makes them vulnerable to conformity bias. This may be due to a number of factors, including the fact that user interests may be difficult to determine and that many items are often interacted with based on ecosystem factors other than their relevance to the individual user. In this work, we introduce CAM2, a conformity-aware multi-task ranking model to serve relevant items to users on one of the largest industrial recommendation platforms. CAM2 addresses these challenges systematically by leveraging causal modeling to disentangle users' conformity to popular items from their true interests. This framework is generalizable and can be scaled to support multiple representations of conformity and user relevance in any large-scale recommender system. We provide deeper practical insights and demonstrate the effectiveness of the proposed model through improvements in offline evaluation metrics compared to our production multi-task ranking model. We also show through online experiments that the CAM2 model results in a significant 0.50% increase in aggregated user engagement, coupled with a 0.21% increase in daily active users on Facebook Watch, a popular video discovery and sharing platform serving billions of users.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; • **Computing methodologies** → **Ranking**; **Multi-task learning**.

KEYWORDS

Recommender system, Multi-task learning, Causal embedding, Conformity bias

ACM Reference Format:

Ameya Raul, Amey Porobo Dharwadker, and Brad Schumitsch. 2023. CAM2: Conformity-Aware Multi-Task Ranking Model for Large-Scale Recommender Systems. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*, April 30-May 4, 2023, Austin, TX, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3543873.3584657>

*Both authors contributed equally to this work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW '23 Companion, April 30-May 4, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9419-2/23/04...\$15.00

<https://doi.org/10.1145/3543873.3584657>

1 INTRODUCTION

Large-scale recommender systems provide personalized recommendations to users by considering their historical interactions and learning predictive models to fit that data. Such systems could result in strong conformity bias as they fail to take into account the fact that previous user interactions may have been influenced by multiple factors causing them to be inconsistent with user preferences. For example, a user may watch a video just because it has a lot of views, even if it doesn't align with their true interests. This, in turn, can result in the system recommending more popular items rather than those that may be more relevant to the individual user, thereby limiting users' exposure to long-tail content [3].

Existing methods to disentangle conformity and interest address this primarily by eliminating popularity bias using a static global term on the item side [2], while ignoring differing levels of conformity among users. To solve this issue, a recent work [19] decomposes user interactions into two factors - conformity and interest, considering both users and items, and learns separate embedding representations for them with cause-specific data. This helps model the personalized conformance effect and captures the fact that different users have different levels of influence on their judgment. It computes user-item matching scores for both causes and sums them up to estimate the overall score on whether a user will interact with the item.

In this work, we present CAM2, a conformity-aware multi-task ranking model to improve user recommendations on Facebook Watch¹ by extending the above idea of cause-specific embeddings. The main contributions of this work are as follows:

- This work disentangles conformity and relevance based representations learning in the same model over the full training data by separating statistical interaction-based features from rich attribute-based and content-based features.
- We present a novel loss formulation for our causal representation modules. In contrast to previous works that rely on aggregating all losses, we train these modules independent of the user interaction tasks' losses. We modify gradient back-propagation from the main network to each of the causal modules to ensure that the performance of the predicted user interaction tasks are optimized without ignoring the causal losses.
- We provide deeper insights into efficiently using the personalized conformity-aware embeddings in our multi-task ranking model to achieve the best predictive performance across tasks.
- Through experiments on our large-scale video recommendation platform, we validate the effectiveness of the proposed approach

¹<https://www.facebook.com/watch>

to improve prediction performance across user interaction tasks, leading to improvements in online metrics.

2 METHODOLOGY

2.1 Problem Formulation

Facebook Watch, with more than 1.25 billion monthly viewers [15] is one of the largest global destinations for discovering and sharing video content. We use a typical two-stage recommendation ranking system design with a candidate generation stage and a ranking stage [6, 11]. This paper focuses on the ranking stage, where the recommender has a few hundred promising candidates retrieved from the candidate generation stage. It applies a sophisticated, large-capacity model to rank these candidates and sort them in descending order of their relevance to users [14].

To effectively learn multiple types of user behavior and relevance dimensions, we use a multi-task deep neural network ranking model to predict multiple binary user interaction objectives such as user clicks, user interaction level with recommended video, user comments, etc. Similar to other multi-task ranking models, our system subsequently combines these predictions to compute a final relevance score for ranking.

Suppose we have a dataset of n i.i.d. training data samples and T is the number of user interaction tasks to model. All tasks share an input space $\{x_i\}$ consisting of user, item and user-item interaction features for example i and $\{y_i^t\}$ is the binary user interaction label for the t^{th} task for example i . The goal of our multi-task model is to learn a causal-aware scoring function $f(x, y, t|\theta)$ which predicts the probability of user interaction t occurring on each example the model is applied on, parameterized by shared and task-specific parameters θ . Our goal is to optimize predictions of the model for each interaction task objective by considering both conformity and relevance, to maximize recommendation performance metrics on future test data, that is not i.i.d. with the training data.

2.2 Conformity-Aware Model Training

Our existing multi-task deep neural network recommendation model serving production traffic on Facebook Watch consists of a sequence of multi-layer perceptrons, composed of a sequence of fully connected layers and an activation function applied component-wise. The model has a combination of hard and soft parameter sharing [16] in the lower layers to exploit task relatedness better [1, 5] and enable some form of regularization as different tasks share model parameters. We input simple binary and continuous features directly into the model as real normalized values. In addition, embeddings are used to process the input of sparse categorical features into appropriate fixed-width dense representations. The model is trained with logistic regression under normalized cross-entropy (NE) loss. The blue box in Figure 1 shows a simplified view of our existing production model architecture.

In order to improve relevance of recommendations by learning better representations, we need to better understand the underlying cause of user interactions. Our goal is to disentangle interactions due to user's alignment to conformity from those primarily due to user relevance that provide higher user value. We learn personalized causal embeddings in our multi-task ranking model corresponding to conformity and relevance based components using the steps

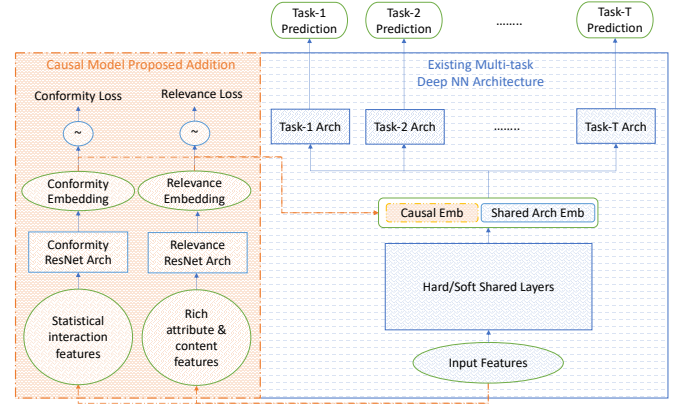


Figure 1: Proposed conformity-aware multi-task ranking model architecture with existing prod model architecture (blue box on the right) and proposed causal model addition to it (orange box on left).

outlined below. The orange box in Figure 1 shows the proposed causal architecture addition to the production model.

2.3 Causal Embeddings Model Architecture

In this work, we train two separate cause-specific residual network (ResNet) modules [10] as auxiliary tasks to learn decoupled causal embeddings indicating conformity alignment and relevance. We also introduce dedicated conformity and relevance losses to train these modules. They are only used to train the causal architectures for learning embedding representations jointly with the main model parameters, and not for model inference. We also disable gradient back-propagation from the task architectures to the causal modules to decouple the causal modules and prevent incorporating any bias in them from the task architectures. This allows us to train our causal embeddings on all available data without partitioning data into factor-specific data subsets based on the underlying reasons of user engagement.

First, the conformity module models the probability of user engagement due to conformity. There are two types of conformity:

- User Conformity, denoted by \bar{u} , is the global likelihood that a user will engage with any given item.
- Item Conformity, denoted by \bar{i} , is the global likelihood that an item will be engaged with by any user.

The conformity module predicts a combined conformity \bar{c} by combining the losses from both of these types of conformity. The total conformity loss (L_C) is calculated as:

$$L_C(u, i) = \|\bar{c} - \|\hat{u} + \hat{i}\|\|$$

where $\|\cdot\|$ denotes L2-norm and \hat{u} and \hat{i} are the predictions for user and item conformity respectively.

Second, the relevance module models the probability of a user engaging with an item based on the user's interest affinity with the item. Our model predicts whether the item is aligned with the user's interests or not. The total relevance loss (L_R) is calculated as:

$$L_R(u, i) = \sum_{x=0}^k \|\bar{r}_x - (\hat{u}_x \cdot \hat{i}_x)\|$$

where $\|\cdot\|$ denotes L2-norm, \bar{r}_x denotes the true engagement rate of a user with interest x , while \hat{u}_x is the predicted probability $\Pr(\text{User Engagement} \mid \text{Interest } x)$ and \hat{i}_x is the predicted probability that the item belongs to interest x . Since each item can be associated with multiple interests k , we aggregate over all these interests to obtain the final loss for a user-item pair.

CAM2 predicts multiple user engagement tasks in the multi-task model, hence we take a weighted sum of these individual cause-specific losses along with our dedicated task losses to get the total loss L .

$$L = \sum_{t=1}^T (w_t \cdot L_t) + w_C \cdot L_C + w_R \cdot L_R$$

where L_t is the loss for task t and w_t are tunable weights. To simplify the above formulation, we model this as a classification problem and define our causal labels by decomposing the probability of user engagement task t into conformity and relevance as:

$$\begin{aligned} \Pr(t) &= w_1 \cdot \Pr(t \mid \text{Conformity}) + w_2 \cdot \Pr(t \mid \text{Relevance}) \\ &= w_1 \cdot \Pr(t \mid X \geq \text{thresh}) + w_2 \cdot \Pr(t \mid X < \text{thresh}) \end{aligned}$$

where X is a scalar derived from user video historical engagement, used to distinguish the conformity and relevance separation with a static threshold thresh . Parameters w_1 and w_2 denote $\Pr(\text{Conformity})$ and $\Pr(\text{Relevance})$ respectively and are learnt jointly with the tasks. The conformity label is 1 if a user engagement happens and $X \geq \text{thresh}$ and 0 otherwise, while the relevance label is 1 if user engagement happens and $X < \text{thresh}$ and 0 otherwise. Hence for the relevance loss, \bar{r}_x is our label as defined above and our model predicts the product ($\hat{u}_x \cdot \hat{i}_x$). For conformity loss, \bar{c} is our label as defined above and $\|\hat{u} + \hat{i}\|$ is our model's prediction.

This causal embedding training approach improves the main (non-auxiliary) tasks' predictions in the model, which are then combined to compute a final relevance score for ranking and recommending top-k items to the given user. The framework is generalizable and can be scaled to support multiple representations of conformity and user relevance using the above formulations.

2.4 Input Features Partitioning

The DICE framework [19] disentangled representations for conformity and relevance by partitioning the training data into cause-specific parts, and trained different embeddings with cause-specific data. In contrast, we present a unique way to achieve this using all available training data samples by partitioning our features into two types:

- Statistical engagement-based features such as number of video impressions, number of video views, social engagement rate of the user, click-through rate on the video, etc.
- Rich attribute-based and content-based features such as user's age, historical engaged videos, video interest topics, content type, video quality features, etc.

We hypothesize that statistical engagement-based features on both the video and user side are sufficient to learn user conformity alignment, while rich attribute-based and content-based features characterizing user and video properties enable user relevance learning. Combining these features results in a higher degree of entanglement and tends to skew the model towards popular and

over-represented training data samples. Our input features partitioning approach allows the ranking model to learn more accurate embedding representations for both popular and long-tail groups by separating each cause of user engagement.

3 EXPERIMENTS

3.1 Experimental Setup

We perform both offline and online evaluations on Facebook Watch, a real-world video recommendation system serving billions of users. We compare results against our production multi-task deep neural network ranking model, which is comparable to other large-scale industry video recommender ranking models [6]. The offline dataset contains 1.2B+ users, 100B+ instances and 35B+ user video engagements. It is split chronologically into the train set for model training and the next-day holdout test set for model performance evaluation. The models are trained recurrently on each day of additional data, starting with the previous network weights and optimizer state.

3.2 Offline Experiments

We use normalized cross-entropy (NE) as our primary offline evaluation metric as we expect accurate predictions, rather than just getting the optimal ranking order [11]. The lower the NE value, the better the model prediction.

The basic conformity-aware multi-task ranking model as illustrated in Figure 1, consists of feature decomposition at the input layer, and two dedicated cause-specific conformity and relevance residual network (ResNet) modules trained with dedicated loss functions. The disentangled embeddings extracted from the corresponding ResNets are used to train the task architectures after concatenating with the main network's shared bottom arch outputs. For simplicity, we refer to this variant as **CAM2-Proposed**. We conduct ablation studies to comparatively evaluate our modeling choices.

- **Embedding Usage:** Instead of concatenating the embeddings with the shared bottom arch outputs, variant **CAM2-TaskArch** experiments with concatenating them directly at the final layer of the main engagement task architectures.
- **Causal Losses:** In order to understand the importance of task independent causal losses, we experiment with a joint loss defined using a combination of existing engagement task labels and our causal labels in variant **CAM2-JointLoss**.
- **Input Features Partitioning:** Variant **CAM2-AllFeats** experiments with adding all features to conformity and relevance modules. It serves as an ablation study on the input features partitioning we propose.

The ablation studies in Table 1 show that our proposed CAM2 model that leverages causal embeddings to learn every task by concatenating with the shared bottom arch outputs, improves over the existing production model. Additionally, using input features partitioning and separate auxiliary conformity and relevance losses improve the model performance further. Our hypothesis is that leveraging disentangled representations in all stages of the model learning can encode different aspects of the reason for user engagement, resulting in better prediction performance across tasks.

Variant	Tasks NE (Aggregated)
CAM2-TaskArch	-0.060%
CAM2-JointLoss	-0.016%
CAM2-AllFeats	+0.029%
CAM2-Proposed	-0.139%

Table 1: Offline Results for ablation studies relative to production. Best results for the metric are indicated in bold.

[0-1 day)	[1-3 days)	[3-10 days)	[10+ days)
+2.43%	+1.17%	-0.28%	Neutral

Table 2: User engagement metric by video age relative to production.

3.3 Online Experiments

We deploy CAM2 on Meta’s dedicated inference cloud [9] and conduct three weeks of A/B testing to rank videos for users on Facebook Watch. The model predictions are used for scoring and generating a ranked list of videos displayed to users. The same ranking strategy and business logic is applied on control and treatment groups for fair comparison. Online A/B test shows that the causal model results in substantial wins across various metrics, all statistically significant with 95% confidence interval.

Compared to production multi-task ranking model control, the treatment group using CAM2 for ranking videos in Watch shows 0.50% increase in aggregated user engagement metric and 0.21% increase in daily active users on Facebook Watch (Figure 2), advancing our billion-scale video recommender system by a large margin.




Figure 2: Overall user engagement metric and daily active users improvements on Facebook Watch in online A/B test.

Our conformity-aware multi-task model reduces popularity bias by serving more long-tail and new videos. The number of videos resulting in 50% and 75% aggregated user engagements increased by 2.35% and 2.06% respectively. As shown in Table 2, we also notice significant increase in user engagement metric on new videos compared to control. This demonstrates that the causal embedding

can learn better long-tail representation and capture relevant user interests by separating conformity information, instead of recommending irrelevant popular videos.

Casual (low activity) users are defined as those users who have been active on Facebook Watch for 1 to 2 days in the last 28 days. To measure engagement as well as retention metrics on casual users, we freeze user activity levels just before the start of the experiment. Our model can better generalize to casual users by leveraging cause-specific embeddings, resulting in significant 0.83% aggregated user engagement metric gains and 0.62% increase in daily active users on Facebook Watch on this user cohort. This indicates that casual users have a better overall experience and are encouraged to return to the platform due to better recommendation quality.

4 RELATED WORK

In recommender systems, it is important to consider that conformity can distort user ratings [12]. When users are influenced by other factors in the ecosystem, their interactions may not be an accurate representation of their own interests [13]. This can lead to biases in user ratings, and ultimately irrelevant recommendations. Existing approaches address this issue as an attempt to eliminate popularity bias through inverse propensity scoring (IPS) [4, 7] or causal approaches [8, 17, 18]. IPS-based methods re-weight popular items by the reciprocal of their popularity score to compensate for the fact that more popular items are more likely to be recommended. However, these methods are very sensitive to the weighting and don’t consider that not all popularity biases are bad, for example, viral and trending videos could be relevant to users and deserve to be recommended more. The causal methods mainly focus on confounding effects but lack fine-grained consideration of how to leverage popularity bias systematically to improve recommendation relevance.

Recent work on DICE [19] which disentangles user conformity and interest representations is most relevant to our work. However, our proposed formulation of training personalized conformity-aware embeddings in the model as auxiliary tasks on the entire training data with input features partitioning, independent of the user interaction task losses has not been explored to the best of our knowledge.

5 CONCLUSION

In this work, we describe conformity bias as a real-world challenge in the design and development of large-scale recommender systems. We present a conformity-aware multi-task ranking model to address this issue and serve the most relevant recommendations. Our scalable and generalizable framework disentangles representations of user conformity and relevance in the model over the entire training data with input features partitioning and dedicated auxiliary losses. We demonstrate the effectiveness of our proposed method to improve user engagement and ecosystem value metrics on Facebook Watch, an industrial video discovery and sharing platform.

CAM2 develops a unique understanding of users by leveraging its input feature partitioning and loss formulation techniques. In the future, we will explore incorporating other user interaction causes in the model and investigate their impact on long-term metrics.

REFERENCES

- [1] Jonathan Baxter. 2000. A model of inductive bias learning. *Journal of artificial intelligence research* 12 (2000), 149–198.
- [2] Robert M Bell, Yehuda Koren, and Chris Volinsky. 2008. The bellkor 2008 solution to the netflix prize. *Statistics Research Department at AT&T Research* 1, 1 (2008).
- [3] Jesús Bobadilla, Fernando Ortega, Antonio Hernando, and Abraham Gutiérrez. 2013. Recommender systems survey. *Knowledge-based systems* 46 (2013), 109–132.
- [4] Léon Bottou, Jonas Peters, Joaquin Quiñero-Candela, Denis X Charles, D Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. 2013. Counterfactual Reasoning and Learning Systems: The Example of Computational Advertising. *Journal of Machine Learning Research* 14, 11 (2013).
- [5] Rich Caruana. 1997. Multitask learning. *Machine learning* 28, 1 (1997), 41–75.
- [6] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 191–198.
- [7] Alois Gruson, Praveen Chandar, Christophe Charbuillet, James McInerney, Samantha Hansen, Damien Tardieu, and Ben Carterette. 2019. Offline evaluation to make decisions about playlist recommendation algorithms. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. 420–428.
- [8] Priyanka Gupta, Ankit Sharma, Pankaj Malhotra, Lovekesh Vig, and Gautam Shroff. 2021. CauSeR: Causal Session-based Recommendations for Handling Popularity Bias. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 3048–3052.
- [9] Kim Hazelwood, Sarah Bird, David Brooks, Soumith Chintala, Utku Diril, Dmytro Dzhulgakov, Mohamed Fawzy, Bill Jia, Yangqing Jia, Aditya Kalro, et al. 2018. Applied machine learning at facebook: A datacenter infrastructure perspective. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 620–629.
- [10] 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.
- [11] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, et al. 2014. Practical lessons from predicting clicks on ads at facebook. In *Proceedings of the eighth international workshop on data mining for online advertising*. 1–9.
- [12] Gael Lederrey and Robert West. 2018. When sheep shop: measuring herding effects in product ratings with natural experiments. In *Proceedings of the 2018 World Wide Web Conference*. 793–802.
- [13] Yiming Liu, Xuezhao Cao, and Yong Yu. 2016. Are you influenced by others when rating? Improve rating prediction by conformity modeling. In *Proceedings of the 10th ACM conference on recommender systems*. 269–272.
- [14] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G Azzolini, et al. 2019. Deep learning recommendation model for personalization and recommendation systems. *arXiv preprint arXiv:1906.00091* (2019).
- [15] Paresht Rajwat. 2020. The Evolution of Facebook Watch. Retrieved Oct 27, 2022 from <https://about.fb.com/news/2020/09/the-evolution-of-facebook-watch/>
- [16] Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098* (2017).
- [17] Ayan Sinha, David F Gleich, and Karthik Ramani. 2016. Deconvolving feedback loops in recommender systems. *Advances in neural information processing systems* 29 (2016).
- [18] Yixin Wang, Dawen Liang, Laurent Charlin, and David M Blei. 2020. Causal inference for recommender systems. In *Fourteenth ACM Conference on Recommender Systems*. 426–431.
- [19] Yu Zheng, Chen Gao, Xiang Li, Xiangnan He, Yong Li, and Depeng Jin. 2021. Disentangling user interest and conformity for recommendation with causal embedding. In *Proceedings of the Web Conference 2021*. 2980–2991.