

Customer Long Term Propensity Driven Prime Video Page Composition

Venkataramana Kini
Prime Video
venkini@amazon.com

Ravi Kiran Divvela
Prime Video
rdivvela@amazon.com

Devendra Yadav
Prime Video
yaddevn@amazon.com

Zhen Wen
Prime Video
zhenwen@amazon.com

Fei Wang
Prime Video
feiw@amazon.com

ABSTRACT

The Prime Video Homepage provides customers with several carousels to explore the diverse catalog. Each of these carousels is constructed around a certain theme. It's not only important to compose the page with individual carousels relevant to the customer, but also balance different customer and business aspects. The Prime Video business positions itself as an entertainment hub with diverse content types such as Movie/TV Shows/Sports etc. with different offers such as free with Prime membership, individual title purchase and channel subscriptions. From a customer perspective, we would like to recommend a certain offer/content type only if they have high propensity to that offer/content type. Also, when customer looks at a page, we would like them to see diverse content in the adjacent carousels. To address these requirements we first develop a long term customer propensity model. Then we linearly combine propensity with carousel relevance scores so that we balance between customer's immediate streaming needs and long term interests. On top of the propensities we apply upper confidence bound (a multi-armed bandit method) so that customers can explore the unexplored offer/content types. Next we make use of customized maximum marginal relevance criterion so that neighboring carousels are not from the same offer/content type. The balance between different offer/content types with customer relevance is analyzed using Pareto front so that we can select appropriate treatments with best balance for an A/B experiment. The best performing treatment from A/B experiment is found to have significant improvement (+4.6%) in customer engagement metrics as compared to control.

KEYWORDS

Customer Long Term Propensity, Transformers for Propensity Prediction, Share of Voice, Multi-objective Optimization & Pareto Front, Page Composition

ACM Reference Format:

Venkataramana Kini, Ravi Kiran Divvela, Devendra Yadav, Zhen Wen, and Fei Wang. 2023. Customer Long Term Propensity Driven Prime Video

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 ACM 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.

Conference'17, July 2017, Washington, DC, USA

© 2023 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

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

Page Composition. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

With the advent of streaming services, users now have a convenient and expeditious way of entertaining themselves. Among the streaming platforms, Amazon Prime Video has emerged as the preferred destination for most users. This is largely due to the extensive range of content available on the platform, including movies, TV shows, sports, linear TV, and more. Moreover, Amazon Prime Video hosts unique video offerings such as included with Prime (or just Prime), with add-on subscriptions such as Prime Video Channels (third-party (3P) channels) and available for rentals/purchases such as Transaction Video on Demand (TVOD), which enhance the viewing experience for customers. The 3P and TVOD together known as non-Prime content. As a result, Prime Video strives to be the ultimate "Entertainment Hub" and epitomize Amazon's offerings to customers.

On the Prime Video Homepage, titles are organized into thematic groups called carousels (eg. "Recommended for You", "Trending Titles", "Continue Watching", "Because You Watched" etc.), with a typical homepage consisting of approximately 20 carousels. The top 5-7 carousels receive the most attention from customers. However, the abundance of offers and content types on Prime Video poses unique challenges when recommending thematic carousels to customers. The top carousels should not only include what customers already enjoy, but also allow for exploring new offers and content types.

Since there is limited real estate (top 5-7 carousels) and many contenders, the practical solution should be fair to different offer/content types and customers as well. To simplify the process of page construction and to make it intuitive, we adhere to certain guiding principles while designing the page. First, we prioritize customer personalization aspects over business aspects. Second, we divide the functionalities of composing page into logical modules so that we that we can make progress for those modules independently. Third, we use offline evaluation / analysis for guiding us to select treatments for an online A/B experiment. The offline experiments used to produce various treatments for the online experiment, and the subsequent online A/B experiment is used to decide which treatment to productionalize. Fourth, we abstract the customer interaction at higher granularity to address the sparseness and business requirements. Finally, we prioritize functionalities and features to be tested in A/B experiments incrementally over time

so that its possible to understand and analyze the impact of the features on customers becomes explainable.

In literature, large scale page composition has been studied at industry scale by Amazon [1], LinkedIn [2], Netflix [3]. Page recommendation policies go beyond simple item click through rate prediction [4, 5] approaches as true user conversion is also affected by the variety of options shown to the user. [6, 7] show improved conversion rates by adding diversity and fairness constraints. With multiple objectives beyond pure user-item relevance, the problem is posed as multi-objective optimization. Pareto-efficient methods for recommendation [8–10] use offline objectives to jointly optimize user fairness, item fairness and diversity. The difference between these existing multi-objective methods and our industrial setting is that our target objectives cannot be estimated accurately in an offline setting. The target user conversion and long term retention has shown significant deviations in offline proxy metrics and online A/B testing performance. Hence, our approach prioritizes online user testing, adopting methods that trade-off multiple offline objectives to produce candidate solutions which will be accurately evaluated with online A/B tests. Moreover, the methods that target diversity or fairness objective do not consider a user's propensity towards the groups and categories, leading to inefficient recommendations. We discuss the differences between existing work and our problem setting in more detail in the Related Work section.

To address complexities of the Homepage carousel recommendation, the task of the carousel recommendation is separated into two modules: the Carousel Relevance (CR) module and the Page Composition (PC) module. The CR module focuses on determining relevance of carousel for a given customer under specific context. It's basically a machine learning model that predicts streaming probability of a title from the carousel given carousel, customer and contextual features. Whereas PC module focuses on composing page from carousels. It utilizes the relevance scores that are generated by the CR module and makes sure that the generated page resonates with the customer intents in terms of offers and content types. It has additional responsibility of making sure that the carousels are placed in the page in such a way that neighboring carousels are not from same offer / content types.

Currently, the PC module relies on many heuristics, rules on top of ML model to impose business requirements during page composition. Since there are multiple offers and content types, balancing between them has become a persistent problem. For example, fixating on offer balance will disturb the content type balance and vice versa. There's no best page defined as such, we have to experiment and figure out utility of the page to the customer. One of the main limitations of the current PC module is that it doesn't consider customer's overall interest towards certain offer/content type.

The customer's interest can be captured by using what customer has interacted in the past. As we look at sufficiently longer time period so that we can learn more about the customer, this is also called as *affinity*. However, affinity is something from the past not necessarily reflective of customer action in future. We can learn and extrapolate from the past long term customer interactions using ML models to predict future interactions. Such predicted affinity is called as *propensity*. We typically try to predict relatively longer term customer propensity. The reason to consider longer

history of the user for predicting propensity is because customer behavior can be seasonal and we need to incorporate data from all seasons for accurate predictions. For example, during holiday months customers may tend to buy certain TVOD content. Also, during summer vacation parents may subscribe to certain channels for kids to watch.

In Prime Video, we have millions of movies and TV shows across different offer types. If we intend to represent customer using interaction history at item level granularity, we end up with a very sparse vector representing customer interactions. Such sparse vector will not capture customer interests at offer/content type level adequately. So instead, we resort to aggregation of the data at offer/content and month level so that customer data is represented as monthly time series at offer/content type granularity. This enables us not only to address sparsity but also abstract out unnecessary complexity and solve problem both from customer and business perspective.

Different machine learning models can be utilized for predicting customer propensity. Since the customer interactions form a time series over period of time w.r.t. different offers, the sequential models such as recurrent neural networks [11, 12] can be utilized. In this paper, we propose to capture long range dependency in offer interaction sequence and future horizon offer/content type Propensity using Transformer [13, 14] models that have been found to be useful in diverse natural language processing sequence modeling tasks.

The propensity of the customers w.r.t. offers / content types predicted by model is broadly guided by historical interactions. E.g. for a customer who has never interacted with TV shows content, the model will predict very low propensity. In first place customer may not be aware of good TV Shows on Prime Video. To address this, we introduce multi-arm bandit exploration on top of propensity scores. Specifically, we propose to use Upper Confidence Bound (UCB) method to introduce exploration on long term customer propensity.

Once the top carousels are selected and ranked w.r.t. content aware relevance, we propose to perform the diversification of neighboring Carousels w.r.t. offer/content types. We adopt a simple approach based on Maximal Marginal Relevance (MMR). We extend the MMR method so that we not only diversify the carousels at overall page level, we also diversify carousels in the neighborhood of carousels.

Contributions: We summarize contributions of this work as follows:

- We propose an approach for modeling customer long term propensity using Transformers
- We utilize customer long term propensity in driving the PV homepage composition
- To account for customer not having a chance to explore certain offer/content type, we perform an UCB based exploration
- We introduce an MMR inspired diversification approach so that neighboring carousels are not from same offer/content type.
- Finally, we conduct both offline and online A/B experiments to ascertain how it can benefit the customers.

The paper is organized as follows. In Section 2, we review the previous work on page composition. In Section 3, we propose approaches for content agnostic/based page composition in Prime Video. In Section 4, we perform extensive validation of our proposed approaches on a proprietary Prime Video dataset and showcase our offline and online results, and finally draw our conclusions in Section 5.

2 RELATED WORK

The problem of page composition is studied in different industrial settings with large scale recommender systems such as Amazon [1], LinkedIn [2], Netflix [3]. Page Composition or Slate Recommendation differs from general user-item click through rate prediction which has been widely studied in literature [4, 5, 15]. Specifically, the page as a whole needs to be diverse to maximize probability of conversion rather than showing multiple similar items with highest relevance scores. It has been shown that diversity in recommendations improves user conversion [6, 7, 16].

Similar to our problem setting, [1] addresses the problem of recommending a diverse set of items in multiple categories while enforcing certain business constraints. To account for diversity, it applies global constraints for achieving predefined exposure for each category without considering individual customer propensity towards each category. It uses Bayesian linear model [17] together with the submodular diversity to model relevance. This approach uses diversity as a feature input but doesn't enforce it. Enforcing diversity across the list of recommended items has been studied in [7, 10, 16, 18] using Maximal Marginal Relevance and Determinantal Point Process. However, these methods treat each category equally and do not consider a user's preference towards each diversity category. In our problem setting, the relationship between items and diversity category (offer types) is not fixed and changes constantly. As such, the relevance score of the item does not reflect relevance of the offer type for a user. Previous works have considered diversifying with user preferences such as genres [6] and personality [19] but these categories have relevance that is tied to item relevance directly. Hence, we extend these methods to account for user's propensity towards each diversity category, removing unpersonalized recommendations generated after enforcing diversity.

Learning optimal recommendation policies with multiple objectives has been studied in the context of Pareto-efficient optimization and fairness. In our problem setting, apart from customer relevance, we also need to consider appropriate exposure for content producers of each offer type. A similar problem was tackled by [8] where they use offline objectives to jointly optimize user fairness, item fairness and diversity. Similarly, [9, 20] propose a Pareto-efficient method to optimize for objective such as fairness across user subgroups and business revenue, in addition to relevance. These methods work well to estimate optimal solution on the Pareto front when the target objectives are known and calculated offline. In our problem setting, the target metrics of business revenue and long term user retention cannot be estimated offline with acceptable accuracy. While measures of relevance, fairness and diversity can serve as proxies, we have observed large deviations in offline estimates and online user behavior. This is due to presence of biases in offline logs and

the fact that the recommender system's output is processed further before presentation to the customer. As such, we aim to learn a method that can produce several solutions on the multi-objective Pareto front, allowing us to narrow down potential candidates for A/B testing.

In previous approaches at Prime Video, the Page Composition problem was posed as a constrained optimization problem [1]. Given relevance, monetary value and diversity measure (IDG) we had minimized deviation between new relevance and given relevance under monetary/diversity constraints. Since monetary value and diversity measure are not based on customer propensity, our online experiments showed that we improved customer's paid transactions and subscriptions but resulted in significant loss in overall engagement metric, mainly due lesser TV show streams.

To account for customer propensity, we develop a novel method to capture customer propensity towards different offers. We propose to make use of Transformers which is usually applied for Natural Language Processing (NLP) tasks [13]. There are some attempts towards making use of Transformers for time series forecasting [14]. However, we propose to make use of transformers for predicting customer's long term propensity and combine it with relevance so that we can balance between different offers. Also, we perform exploration on top of propensities using UCB [21] so that customers can explore new offers. Consequently, we diversify the page in the neighborhood of carousels using customized MMR criterion [18].

3 PROPOSED APPROACH

The flow of proposed overall approach is shown in Figure 1. There are two main parts to the flow. The first part is content agnostic part shown in blue color and in contrary second part is shown in red color and is after filling the carousels with specific titles. This differentiation is due technical and budget constraints as it will not be cost effective to fill in thousands of possible carousels with title information. The infrastructure cost optimization leads to two relevance models: Content Agnostic Relevance Retrieval (CARR) and Content Based Relevance Ranking (CBRR). Each of these relevance modules is followed by "Page Composition" modules, which performs trade-off between relevance and some of the business aspects. These modules are logically separated in such a way that we can perform improvements to each of these modules simultaneously.

When we perform any business trade-off, its important that we do not hurt the customer page relevance. In our past online experiments, we observed that increasing non-prime content (TVOD/PV Channels) positively impact PV Channel sign ups but adversely impact customer engagement metrics. Hence, we have learned that it is essential to consider customer propensity towards non-Prime content while designing the page. In this work we hypothesize that higher customer's long term propensity towards specific offer, higher should be the customer's Share of Voice (SoV) towards that offer.

The first Page Composition (PC) module performs selection of carousels such as way that we achieve optimal Share of Voice of offers/content types by considering customer's long term offer propensity. The second PC module addresses the diversification with respect to different offers on the page. The diversification makes sure that neighboring carousels are diverse in terms of offers.

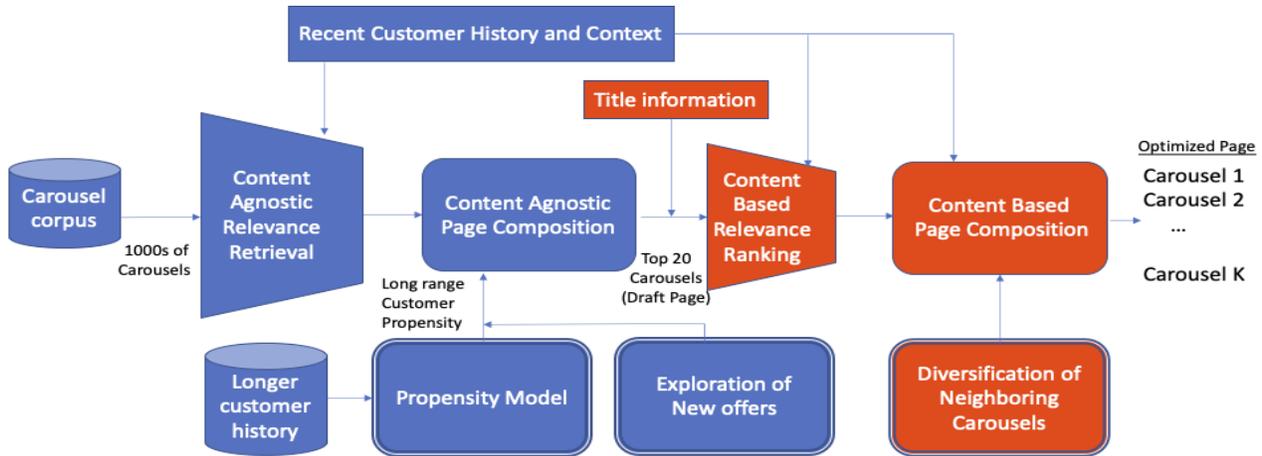


Figure 1: Flow of proposed approach for page composition

The relevance modules (CARR and CBRR) are developed independently with the main focus being "relevance" of carousels to the customers and is not part of this study. In this work, we focus on page composition modules: Content Agnostic Page Composition (CAPC) and Content Based Page Composition (CBPC) and proposed modules are shown by double boarder in Figure 1.

3.1 Content Agnostic Page Composition

This is the first layer of the page composition and main goal is to select the carousels for a page based on customer's overall long term propensity. As the propensity is key towards deciding the share of voice, we develop a machine learning model that considers customer's long (past one year) historical behavior towards offers and content types and extrapolates it to future time period (next one month). Once we estimate the propensity of the customers, we need to combine it with relevance score provided by CARR so that we take into account both customer's carousel relevance and customer's long term offer propensities. Towards this we devise a simple approach by linearly combining relevance and propensity.

3.1.1 Customer Long Term Propensity Model. We develop Customer Long Term Propensity Model using customer interaction history. We pose Customer Long Term Propensity Model as a multivariate time series forecasting model. I.e. given customer's last 12 months of interaction history, we would like to predict next one month interactions. However, the traditional time series models such as ARIMA model will not be able to capture long term customer interaction history. Also, models like ARIMA will not be able to train on time series from multiple customer time series data.

The traditional machine learning models such as recurrent neural networks (RNNs) implemented using LSTMs are not capable of capturing the long term dependencies in customer behavior. We propose to develop generative Transformer model that uses attention mechanism to model non-Markovian dynamics. We make use of Transformer model as a multivariate time series forecasting model.

The proposed Transformer model architecture for customer propensity prediction is shown in Figure 2. The Transformer has

two parts: encoder and decoder similar to [13]. But unlike machine translation task, we do not have two separate datasets as source and target. In our case, we split first 6 months' (month 1-6) customer history as encoder input and next 6 months' (month 7-12) data as decoder input and last month's (month 13) as decoder output. The estimated propensity of the customer is based on 13th month's prediction. This formulation is similar to [22], but we perform multi-output regression with MSE loss function and we do not form moving window. This helps us in terms of representing the first 6 months of time series effectively and then feed it to decoder. The decoder in turn focuses on more recent data and encoder generated embedding to predict the last month's data point.

Let's suppose we have K different offer/content type combinations. At each time step (monthly), we aggregate number of streams w.r.t. K different offer/content types. We represent a customer u 's streams for a month t by $y_{ut} = [y_{ukt}]_{k=1}^K$. The propensity model predicts customer's streams from different offer/content type for next month. Next we normalize individual offer/content propensity by sum of propensities of all offers so that propensities sum to 1.

3.1.2 Combining Relevance and Long range Customer Propensity. Having estimated the customer's long term propensity, next we propose an approach to linearly combine customer's long term propensity and relevance score of a carousel. This approach provides flexibility to combine the relevance and propensity to different degrees so that we can influence the appropriate SoV for a customer based on the their propensity.

$$P_{r,p} = (1 - \lambda)P_r + \lambda P_p \quad (1)$$

Here, P_r represents the relevance score of a carousel and P_p propensity score from the corresponding offer/content type. λ is the parameter that controls the influence of propensity and relevance. Based on the combined score $P_{r,p}$, the top 20 carousels are selected.

Another important dimension to consider is content type (Movie / TV Shows). Typically customers watch movies which doesn't require lot of commitment in terms of time. TV Shows customers watch over days or even weeks/months leading to higher engagement with platform. The propensities generated by the model are

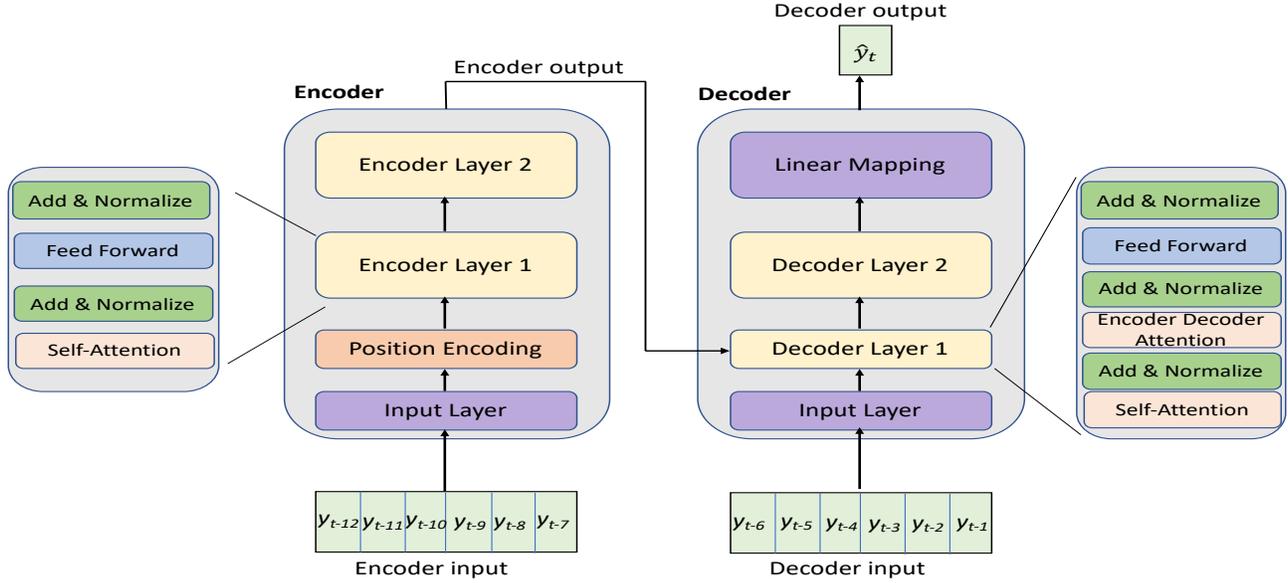


Figure 2: Transformer based propensity model

aggregated at the content type level and linearly combined with both relevance and offer level propensity:

$$P_{r,p} = (1 - \lambda_1 - \lambda_2)P_r + \lambda_1 P_p^o + \lambda_2 P_p^c \quad (2)$$

P_p^o and P_p^c are offer and content type propensities respectively. One of the limitations of using propensity models is that they rely on extrapolating observed customer behaviors. If an offer (say TVOD) was never interacted by customer, propensity model will predict low score for TVOD. As customer has not explored TVOD content in the past, customer may not be aware of TVOD offer so there's finite chance that customer may consider streaming from TVOD carousel. To address this we propose an approach based on Upper Confidence Bound [21].

$$\tilde{P}_p = P_p + c\sqrt{\frac{1}{N}} \quad (3)$$

Here, P_p is propensity predicted by model for specific offer/content type and N is the number of times the same offer/content type recommended to the customers in the last month. c is the hyper-parameter to control the amount of exploration.

Once the carousels are selected, the titles for each of the carousels is filled by calling respective services. Also, CBRR module (ResNet based model) is used to predict the score for each of the 20 carousels by using title and customer/contextual features. The resulting carousel relevance scores are represented using $P_{r'}$.

3.2 Content Based Page Composition

Having selected the carousels in the content agnostic layer, next we diversify the neighboring carousels. The objective is to make sure that neighboring carousels are not from same offer / content type. Towards this we propose an Maximum Marginal Relevance (MMR) [18] based approach while iteratively constructing page by adding one carousel at a time. We maintain a priority queue for

each offer/content type by sorting carousels based on the relevance. We pop candidates from each of the queues and apply:

$$P_{r,n} = \gamma P_{r'} + (1 - \gamma)(1 - P_{page})(1 - P_{nbr}) \quad (4)$$

Here, γ is the parameter to control the amount of diversity and relevance. The higher value of γ favors relevancy and lower value of γ favors diversity. We use two terms for diversity: one at page level, P_{page} representing % of specific offer-content type carousels already added to page and other at neighboring carousel level (P_{nbr} : representing % of specific offer-content type carousels already added in previous 2 slots). The new term P_{nbr} introduced helps in favoring diversity not only at page level but also at carousel neighborhood level.

4 EXPERIMENTS AND RESULTS

In this section, we validate both content agnostic and content based page composition models. The relevance models are validated separately and is not part of this study. However, we make use of scores generated by these models as relevance scores. We analyze trade-off between conversion and different share of voices. Particularly we analyze non-Prime and TV SoVs as complement of them will be Prime and Movie SoVs respectively.

Experimental Setting: We perform all experiments on US market place. The content agnostic and content based stages are logged separately. We sample from large number (Millions) pages over a month for both content agnostic/based. The number of carousels in a page for content agnostic logs is few hundreds to a thousand. The number of carousels in a page for content based logs is 15-20 after filtering out the pinned carousels.

Model	MAPE (%)
MLP	30.28
LSTM	55.42
Transformer	70.28

Table 1: Performance improvement of different ML models compared to baseline (affinity) for predicting propensity w.r.t. MAPE metric

4.1 Experiments - Content Agnostic Page Composition

4.1.1 Evaluation of Propensity Models. The main component of proposed CAPC is propensity model. The propensity model helps to understand customer’s future interests at offer level. As we have posed the propensity model as a time series forecasting model, we make use of the metrics that are used in time series forecasting. We make use of Mean Absolute Percentage Error (MAPE) for measuring the performance of different propensity models. Since in our case we perform multivariate time series (w.r.t. different offers), we extend it to account for this aspect.

$$MAPE = \frac{100}{UK} \sum_u \sum_k \frac{|y_{uk} - \hat{y}_{uk}|}{y_{uk}} \quad (5)$$

Where U is number of customers, K is number of distinct offer/content types, y_{uk} and \hat{y}_{uk} are actual and predicted offer/content type aggregated streams respectively. As discussed previously, we predict the customer’s next one month offer propensity using past 12 months behavior. The performance of prediction is measured using MAPE, which provides scale free intuitive metric as scales of different customer’s offer consumption can vary significantly.

We experiment with three different ML models for modeling customers long term offer propensity. The first one is feed forward neural network implemented using Multi-Layer Perceptron (MLP), the second one is a Recurrent Neural Network (RNN) implemented using Long and Short Term Memory (LSTM) and finally Transformer with regression loss function as discussed in proposed methods section.

We compare the ML propensity models against naive affinity (aggregation past streams) based method. It can be observed that in Table 1, MLP has 30.28 % lesser MAPE as compared to affinity baseline. LSTM has 55.42% lesser MAPE as compared to affinity baseline. Further, Transformer has 70.28% lesser MAPE as compared to affinity baseline. The proposed transformer model transformer model turns out to be best at predicting customer propensity. The propensity estimated by model is raw count. We normalize the propensity by sum of all offer/content type propensities at customer level.

4.1.2 Balancing of Offer/Content Type. Next, we combine the long term propensity scores generated at customer/offer level with the relevance score at carousel level as per equation 1. It can be seen that for different values of λ we can achieve different levels of non-Prime share of voice (Figure 3(a)). Also, we can see that there is trade-off between the non-Prime SoV and Inverse Propensity Score (IPS) normalized conversion rate [23].

When we introduce exploration with $c = 0.01$ in equation (3), the SoV changes more aggressively as expected as can be seen in Figure 3 (b). As an initial A/B experiment, we plan to focus on balancing between SoV and conversion without performing exploration. In subsequent A/B experiments we plan to introduce exploration. Therefore, we’ll focus on non-exploration based methods for further offline experiments in this paper.

There is inherent trade-off between conversion and non-Prime SoV as typically customers stream the content that is already available with Prime. This trade-off can be depicted using Pareto Frontier [9]. The overall trade-off between conversion and non-Prime SoV is represented using Pareto Frontier in Figure 4(a). Though we see direct trade-off between non-Prime SoV and conversion, in practice (through A/B experiment) we can study how much of real trade-off is there as we influence non-Prime SoV through customer’s long term propensity towards non-Prime offers. From these plots we can infer that carousel level instantaneous conversion is not in synchronization/aligned with long term customer-offer level propensity.

Similar analysis is performed at content type level to balance between the TV SoV and IPS normalized conversion. Due to space limitations, we do not include these results.

In practice we’re not just interested in offer or content type trade-off with conversion separately as analyzed till now. The actual use case that we’re interested in the combined effect of offer and content type trade-off with conversion. To account for offer and content type together, we apply equation 2 to balance 3-way between carousel relevance, customer’s long term offer propensity and customer’s long term content type propensity. We vary the λ_1 and λ_2 parameter in steps forming a grid. Then, we obtain similar SoV curves and Pareto Front. In this case Pareto front is 3-D surface which becomes difficult to visualize. Also, picking appropriate points on Pareto surface becomes tricky.

To address this challenge, we ran Skyline algorithm [24] to figure out points that are on the Pareto front. With that we are left with around 35 points out of 120 points generated by varying λ_1 and λ_2 . As we can’t perform A/B experiment on 35 treatments, we filter further based on threshold on IPS normalized conversion minimum of 0.20 as we don’t want to compromise on customer relevance. Then we pick final 6 treatments so that we cover broad range of TV / non-Prime SoV. This is achieved by picking points around min, max and median of TV / non-Prime SoV. These points are shown on the 3-D plot as blue points (Conversion > 0.20, TV SoV: (min: 0.12, median: 0.21, max: 0.45), non-Prime SoV: (min: 0.17, median: 0.39, max: 0.46)) and will be picked as final treatments for A/B experiment.

4.2 Experiments - Content Based Page Composition

As can be seen in Figure 1, once the candidates are chosen in the PC layer 1, the carousels are re-ranked by content based relevance ranker. These relevance scores are used in final layer 2 of page composition. The objective of the layer 2 of page composition is to diversify the carousels so that neighboring carousels are not from same offer / content type. Toward this we apply equation 4 as we construct page by adding one carousel at a time to the page.

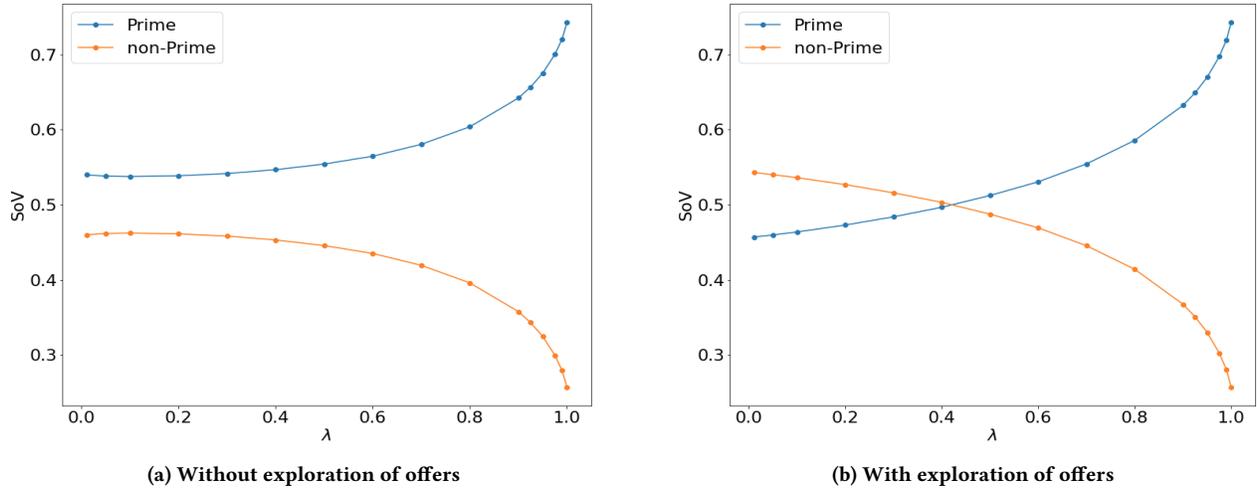


Figure 3: Prime and non-Prime SoV for different values of λ

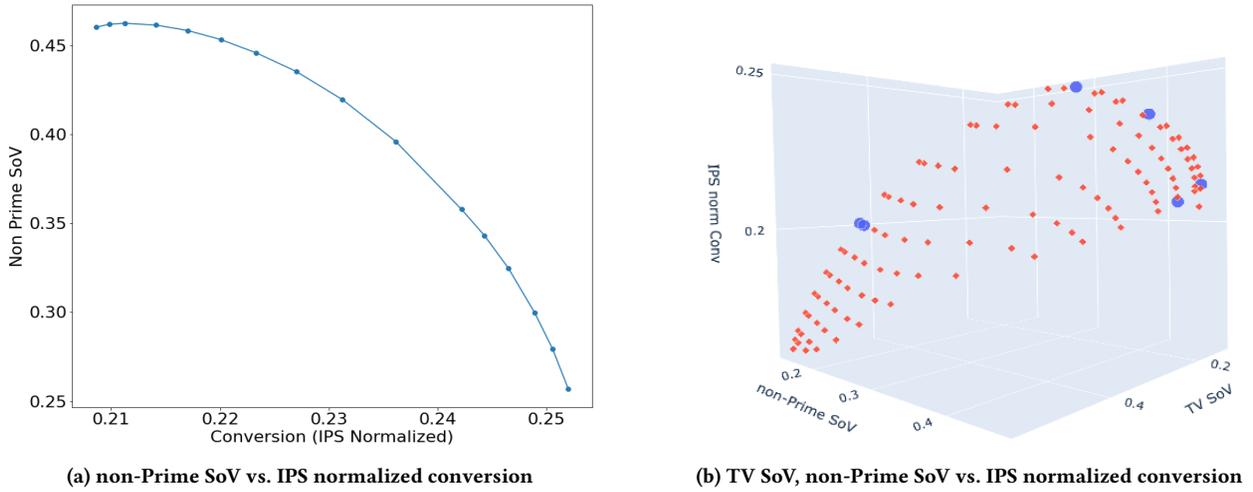


Figure 4: Pareto front for different SoV vs. IPS normalized conversion with CAPC

We define a new metric to measure the neighborhood diversity in terms of offer / content type:

$$ISL - Div = \frac{\sum_k (\# \text{ of distinct neighbors at rank } k)}{\sum_k (\max(\# \text{ of distinct neighbors at rank } k))} \quad (6)$$

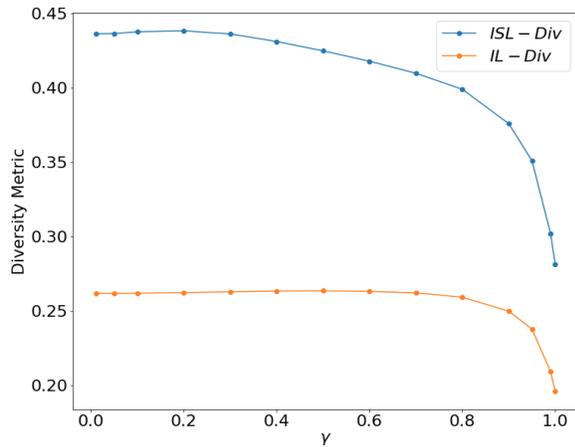
The traditional intra-list diversity ($IL - Div$) [25] and the $ISL - Div$ is shown in Figure 5(a) for different values of γ . It can be seen that as γ increases $IL - Div$ remains constant whereas $ISL - Div$ changes gradually suggesting change in offer/content type in the neighborhood.

The Figure 5 (b) shows the Pareto front for Conversion (IPS normalized) vs. $ISL - Div$. It can be seen that as diversity increases, conversion also increase slightly (right bottom side). This can be due to diverse carousels in fact lead to higher conversion. Overall, Conversion reduces as diversity increases as expected. Therefore, for achieving right level of $ISL - Div$ we recommend to have 3

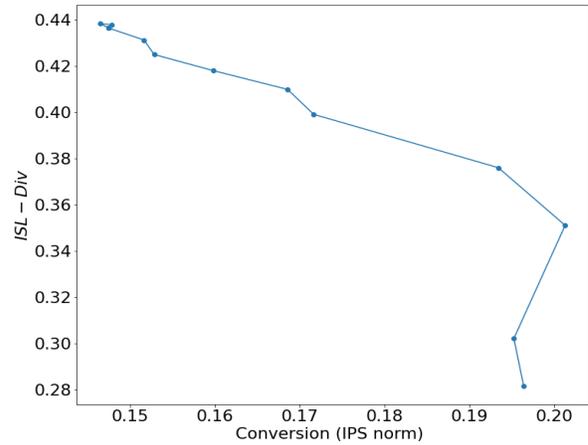
conservative treatments with $\gamma = 0.8, 0.9, 0.95$ for A/B experiment of CBPC.

4.3 Online A/B Experiments

Having established the utility of long term propensities through offline experiments, we next perform an online A/B experiment to see actual benefits of the proposed long term propensities on real customers. As its challenging to test out multiple dimensions simultaneously, we first focus on improving content type engagement of customers while keeping movie/offer engagement constant. Towards this we first estimate customers' content propensity for next month using method suggested using transformer model. For this we just aggregate customer's TV predicted propensity across offers. Based on the different λ s and Pareto analysis as shown in Figure 3 and 4, we select 6 treatments that was tested online during subsequent month.

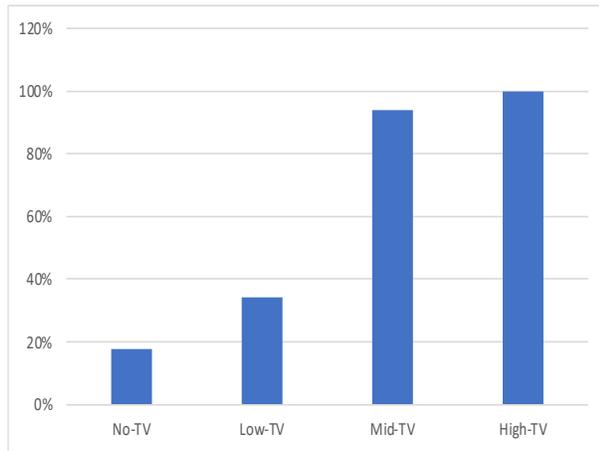


(a) Diversity of pages measured using *IL - Div* and *ISL - Div*

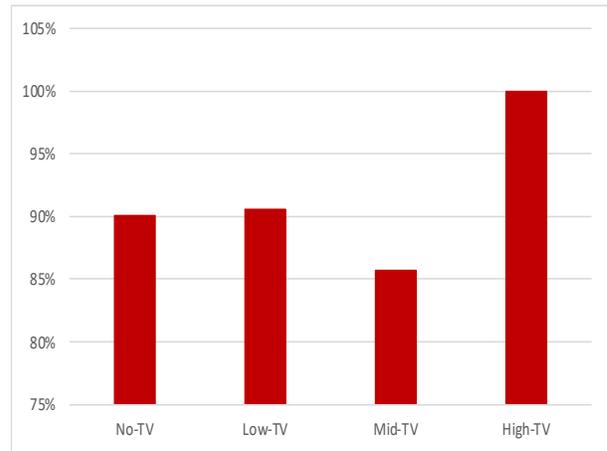


(b) Pareto front of *ISL - Div* vs. conversion (IPS normalized)

Figure 5: Diversity and Pareto front with CBPC



(a) For Treatment



(b) For Control

Figure 6: Number (normalized) of TV recommendations for different TV customer segments

To analyze how TV propensities influence the number of TV carousel recommendations in top 5, we group customers based on their propensity towards TV shows. For the dialed up treatment, the average TV share of voice for no-TV, low-TV, mid-TV, and High-TV customers is shown in Figure 6(a)¹, this is contrast to more or less constant number of TV recommendations for all segments in control (Figure 6 (b)). I.e. the proposed approach is able to consider long term customer propensity and tune number of TV recommendations whereas current production approach is not able to consider customer propensity, only the High-TV customers are recommended higher number of TV recommendations and no/low/mid TV customers are recommended more or less same number of recommendations. After the experiment, we have seen significant customer engagement with TV shows without drop in movie streaming or negative impact to offer balance. As a business

¹normalized by number of recommendations for High-TV customers

metric, customers engagement is measured in terms of time customers spent on the platform. During our online A/B tests in the US market, we observed a significant uplift of 4.6% in customer engagement metrics. Motivated by the impressive results from US market, we plan to conduct experiments in other major market places across world such as Europe, India.

5 CONCLUSION AND FUTURE WORK

In this paper, we proposed an approach for page composition driven by propensity of the customers towards offers/content types. We posed the long term propensity estimation as a time series forecasting problem, modeled using Transformers. The customer propensity is combined with relevance of carousels so that we select optimal number of carousels from different offer/content types. We also customized the MMR diversification approach so that neighboring

carousels in the page are not from same offer / content type. We evaluated the approach on PV homepage logs and selected the optimal points on the Pareto front for an A/B experiment. After conducting the online A/B experiment, we found out that proposed propensity based approach significantly improves customer engagement.

In the current work we have explored post-processing methods to combine propensity and relevance. Also, for diversification we used MMR based post-processing method. In future, we plan to use an ML model that can combine propensity/diversity aspects. Further, instead of looking at Prime/non-Prime, Movie/TV, we can extend the approach to handle other offer/content types and the exploration on offer/content types has to be studied more and needs a separate A/B experiment in the future to ascertain how much exploration we can introduce.

REFERENCES

- [1] Weicong Ding, Dinesh Govindaraj, and S. V. N. Vishwanathan. Whole page optimization with local and global constraints. In *KDD 2019*, 2019.
- [2] Deepak Agarwal, Shaunak Chatterjee, Yang Yang, and Liang Zhang. Constrained optimization for homepage relevance. pages 375–384, 05 2015.
- [3] Chris Alvino and Justin Basilico. Learning a personalized homepage. <https://netflixtechblog.com/learning-a-personalized-homepage-aa8ec670359a>. Accessed: 2015.
- [4] Zhiqiang Wang, Qingyun She, and Junlin Zhang. Masknet: Introducing feature-wise multiplication to ctr ranking models by instance-guided mask. *arXiv preprint arXiv:2102.07619*, 2021.
- [5] Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, pages 191–198, 2016.
- [6] Harald Steck. Calibrated recommendations. In *Proceedings of the 12th ACM conference on recommender systems*, pages 154–162, 2018.
- [7] Mark Wilhelm, Ajith Ramanathan, Alexander Bonomo, Sagar Jain, Ed H Chi, and Jennifer Gillenwater. Practical diversified recommendations on youtube with determinantal point processes. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 2165–2173, 2018.
- [8] Lequn Wang and Thorsten Joachims. User fairness, item fairness, and diversity for rankings in two-sided markets. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval*, pages 23–41, 2021.
- [9] Xiao Lin, Hongjie Chen, Changhua Pei, Fei Sun, Xuanji Xiao, Hanxiao Sun, Yongfeng Zhang, Wenwu Ou, and Peng Jiang. A pareto-efficient algorithm for multiple objective optimization in e-commerce recommendation. In *Proceedings of the 13th ACM Conference on recommender systems*, pages 20–28, 2019.
- [10] Wenjing Luan, GuanJun Liu, ChangJun Jiang, and MengChu Zhou. Mptr: A maximal-marginal-relevance-based personalized trip recommendation method. *IEEE Transactions on Intelligent Transportation Systems*, 19(11):3461–3474, 2018.
- [11] Yao Qin, Dongjin Song, Haifeng Chen, Wei Cheng, Guofei Jiang, and Garrison Cottrell. A dual-stage attention-based recurrent neural network for time series prediction, 2017.
- [12] Byeongjin Choe, Taegwan Kang, and Kyomin Jung. Recommendation system with hierarchical recurrent neural network for long-term time series. *IEEE Access*, 9:72033–72039, 2021.
- [13] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [14] Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. Transformers in time series: A survey, 2023.
- [15] Ruoxi Wang, Rakesh Shivanna, Derek Cheng, Sagar Jain, Dong Lin, Lichan Hong, and Ed Chi. Dcn v2: Improved deep & cross network and practical lessons for web-scale learning to rank systems. In *Proceedings of the web conference 2021*, pages 1785–1797, 2021.
- [16] Laming Chen, Guoxin Zhang, and Hanning Zhou. Improving the diversity of top-n recommendation via determinantal point process. In *Large Scale Recommendation Systems Workshop*, 2017.
- [17] Thore Graepel, Joaquin Quiñero Candela, Thomas Borchert, and Ralf Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *Proceedings of the 27th International Conference on Machine Learning ICML 2010, Invited Applications Track (unreviewed, to appear)*, June 2010. Invited Applications Track.
- [18] Jaime Carbonell and Jade Goldstein. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '98, page 335–336, New York, NY, USA, 1998. Association for Computing Machinery.
- [19] Wen Wu, Li Chen, and Yu Zhao. Personalizing recommendation diversity based on user personality. *User Modeling and User-Adapted Interaction*, 28(3):237–276, 2018.
- [20] Ananth Balashankar, Alyssa Lees, Chris Welty, and Lakshminarayanan Subramanian. What is fair? exploring pareto-efficiency for fairness constrained classifiers. *arXiv preprint arXiv:1910.14120*, 2019.
- [21] Richard S. Sutton. *Reinforcement learning: An introduction*. MIT, 2018.
- [22] Neo Wu, Bradley Green, Xue Ben, and Shawn O’Banion. Deep transformer models for time series forecasting: The influenza prevalence case, 2020.
- [23] Miroslav Dudik, John Langford, and Lihong Li. Doubly robust policy evaluation and learning, 2011.
- [24] S. Borzsony, Donald Kossmann, and Konrad Stocker. The skyline operator. pages 421 – 430, 02 2001.
- [25] Saúl Vargas and Pablo Castells. Rank and relevance in novelty and diversity metrics for recommender systems. In *Proceedings of the Fifth ACM Conference on Recommender Systems, RecSys '11*, page 109–116, New York, NY, USA, 2011. Association for Computing Machinery.