Carnegie Mellon University
Tepper School of Business

Improving Click-through Rate on Mobile Lock Screen Advertisements

MS Business Analytics Program

Hugo Camou, MSBA
Johnathan Khodr, MSBA
Jamie Khoo, MSBA
Andrew Lee, MSBA

Carnegie Mellon Tepper School of Business

ABSTRACT

Glance delivers its users interactive content (cards) to their phone's lock screens. "Clicks" on sponsored cards is a measure of success for a campaign. Dividing clicks by number of impressions gives click through rate (CTR), a key metric when comparing success across campaigns.

Glance already utilizes a sophisticated recommendation system to distribute cards to its users. (Oli et al., 2020) However, to maximize revenue, it is important to understand what motivates a user to interact with a sponsored card. Thus, our goal was to create a modeling methodology that **improves upon the existing recommendation system**, makes **consistent predictions** across multiple card types, and offers **interpretable results**.

By predicting clicks on a target glance at a user level, we were able to develop a modeling methodology that achieves each of these goals. We feature engineered user data and past click history to make our predictions and found that past click history was most predictive of future clicks.

INTRODUCTION

Glance's end goal is to improve the performance of their ad campaigns. Our approach was to enhance Glance's ability to assign a card to the user most likely to click it. To that end, Glance desires **interpretable assignments**: to know not only when a user is likely to click on a certain card, but why they do so.

We leveraged three main data sources: **user, card and impression data**. User data contains features such as the user's phone, subscribed categories, phone apps, region, and language. Card data contains information related to the card's content and category. Impression data records the user ID and card ID of the impression and whether it was a click. Clicks can occur in the lock screen or in the Glance app. For this project we focused only on lock screen clicks.

METHODS AND MATERIALS

We began by using lasso regularization and elastic nets to remove irrelevant features. This filtered dataset became the basis for predicting whether a user could click on a specific target sponsored card.

The biggest challenge we faced with our data was a very large class imbalance, as most users do not click on the cards (>90%). **Lift** was the preferred KPI to demonstrate how our models add value to the existing recommendation system. Lift is a metric that compares how much better our model is at correctly identifying clicking users vs randomly picking users. All our models were tuned with cross validation, with lift in the top 20 percentile as KPI.

Feature engineering was also important. Data was gathered on users and their mobile devices: model, price, tier, installed apps, and user self-elected subscribed categories of cards. We supplemented these features with cards users have clicked on.

We used these features to construct a user-ID indexed sparse matrix, which can be seen in Figure 1. We also attempted to use our sparse matrix to create topic models (Latent-Dirichlet Allocation), however this did not improve lift and severely compromised interpretability, so this effort was suspended.

Figure 1: Representation of the user-ID indexed sparse matrix.

User ID	Card X	Card Y	App X	App Y	Category X	Category Y	Phone X	Target Card
123	0	1	1	1	1	0	1	1
456	1	0	0	1	0	0	1	0

Whether the user clicked on cards x and y

Whether the user has an app on their phone, is subscribed to a particular category, or has a specific phone model.

Whether the user clicked on the target card.

Interaction Features

User Features

The sparse matrix was generated for roughly 47,000 users highly-active users (10+ clicks). The matrix included 4092 unique cards, 627 apps, 31 subscribed categories, and 68 phone models. Due to the large number of "Interaction Features", when predicting on a specific target card, only interactions within a two week period prior to the target card's expiration date were considered.

Multinomial Naive Bayes (MNB) and Decision Tree (DT) models were trained using an 80-20 train-test split along the user ID index.

RESULTS

We ran the DT and MNB models on 8 cards from each category - sports, finance, and arts & entertainment. The table below shows the average lift generated by targeting the top 20% of users in each category for both models. The MNB model outperforms the DT model in each category.

Model Type	Sports	Finance	Arts & Entertainment
MNB	**1.82**	**1.79**	**1.83**
DT	1.10	0.95	0.97

Table 1: Lift in top 20% generated across 3 card categories by MNB and DT models.

The high lift generated from the MNB model across each of the 3 categories shows we are able to improve upon Glance's existing recommendation system in a consistent manner, satisfying two of our goals.

The last step was to validate the interpretability offered by our model by analyzing the top 10 **important features** the MNB model used to predict a user click. Feature buckets for the top 10 features were created in order to simplify this process.

- All apps and subscribed categories were placed under "User Features"
- All phone models and tiers were placed under "Phone Features"
- All card interactions were bucketed by their category. (Ex: if two features were interactions with Finance cards, then placed under "Finance Cards".)

Figure 2 shows the results of running the MNB model on 8 sports cards. We can see that predictions derived from a user's previous clicks provided us with the most predictive value, while user and phone features are also useful.

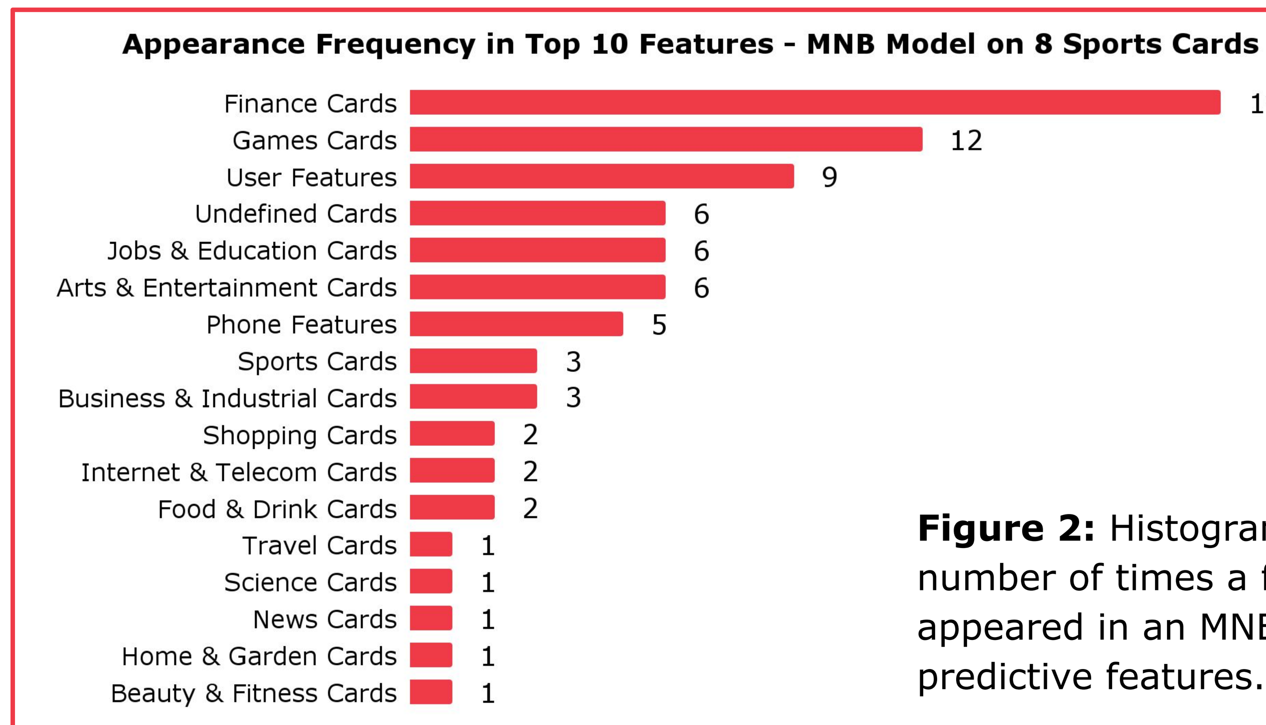

Figure 2: Histogram showing number of times a feature bucket appeared in an MNB model's top 10 predictive features.

DISCUSSION

Our MNB model is able to consistently identify a subset of users who are likely to interact with a sponsored glance card, and can do so in an interpretable way. With regards to Sports cards, we are able to identify a subset of 20% of the users that click 1.8x as often.

The MNB model determined that users who had previously clicked on Finance and Games cards were more likely to click on Sports cards. The model is likely picking up features related to

- **Age**: Many sports audiences are younger and are more interested in getting financial tips/news from their phones.
- **Interests**: Games are closely related to sports, likely with overlapping interests.

Even with cross validation efforts, the DT model was unable to offer similar performance to the MNB model likely due to its inability to handle both the extremely sparse feature set and the imbalanced class data. By treating the sparse matrix features as counts, MNB was able to provide better lift.

LOOKING AHEAD & CONCLUSION

One limitation with our analysis is scalability. The current solution achieves better CTR by identifying the best 20% of users to receive a card, potentially limiting the card's impression volume. This could pose an issue for Glance when needing to scale out an advertising campaign. To overcome this, we suggest reallocating impressions by applying this methodology on a sufficiently large set of users and sponsored cards.

For example, if a user is not within the top 20% for a specific card, our model may indicate they are within the top 20% for another card, at which point they may be reallocated to the latter. Running this a sufficient number of times for many users and many target cards, Glance will be able to retain the overall number of impressions while still improving CTR, resulting in an overall higher number of clicks.

In conclusion, we have created robust predictive models that can identify whether a set of users will click on a given glance across a variety of card types. By running these models over a sufficient number of users and cards, Glance will be able to better allocate sponsored campaign cards to their users, thereby improving CTR.

CONTACT

Hugo Camou: hcamouca@tepper.cmu.edu
Johnathan Khodr: jkhodr@tepper.cmu.edu
Jamie Khoo: jkhoo@tepper.cmu.edu
Andrew Lee: awl2@tepper.cmu.edu

REFERENCES

Oli, N., Patel, A., Sharma, V., Dacharaju, S.D., & Ikhar, S. (2020, August 22). *Personalizing Multi-Modal Content for a Diverse Audience: A Scalable Deep Learning Approach* [Paper Presentation]. IRS2020.