Major Book Publisher

INDUSTRY

BOOK PUBLISHING

WEBSITE

MAJOR BOOK PUBLISHER

PRODUCT

MACHINE LEARNING EMBEDDINGS

DEEP RECOMMENDATION

RESULTS

- Increase **Conversion** (Sales)
- Increase **Engagement**
- Reduce Unsubscribe

In this specific success story, the goal of a Major Book Publisher was to optimize the conversion after sending an email to their user, i.e. having a user sufficiently intrigued or interested by the recommendation to first click on an item, and then purchase the item. Crossing Minds unique embedding of the users and the items recommended allowing MBP to tackle the cold start problem and recommend more accurately new books or new users.

Goal

Increase user engagement, increase book sales, and create more personalized recommendations through email.

Challenges

Our client, referred to in this document as Major Book Publisher (MBP), is one the largest paperback publisher in the world. They operate with an annual revenue of $3B+ and are a division of Bertelsmann. With one of the largest direct competitors, Amazon, plus a large inventory, MBP needs to keep its users engaged and aware of the latest book releases. To do so, the company sends on a bi-

embedding of the users and the items recommended allowing MBP to tackle the cold start problem and recommend more accurately new books or new users.

Goal

Increase user engagement, increase book sales, and create more personalized recommendations through email.

Challenges

Our client, referred to in this document as Major Book Publisher (MBP), is one the largest paperback publisher in the world. They operate with an annual revenue of $3B+ and are a division of Bertelsmann. With one of the largest direct competitors, Amazon, plus a large inventory, MBP needs to keep its users engaged and aware of the latest book releases. To do so, the company sends on a bi-

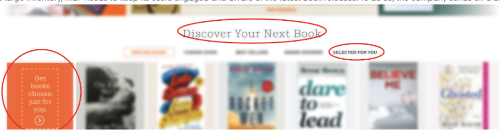

Deliverables

Crossing Minds analyzes the MBP dataset through the Crossing Minds Engine and provide reports to MBP with the results of such analysis. The deliverable for this use case are:

- 3 extensive Crossing Minds reports
- Analysis, insights, and suggestions about the dataset handled by MBP
- **Book recommendations for email marketing**

Work Process

In this very specific case, the goal of MBP is to **optimize the conversion** after sending an email to their user, i.e. having a user sufficiently intrigued or interested by the recommendation to first click on an item, and then purchase the item.

To that extent, the dataset considered contains clicks from users (or the absence) to book links, which are considered as implicit feedback. Our first contribution was to provide MBP with clearer insights on the data they are handling with respect to recommendation ability; and suggest the best addition to this dataset to improve the recommendations. Being the core of any supervised recommender system, it is crucial to get accurate insights into the user-item interaction graph. We presented our density analysis of the adjacency matrix, and then further details on the connectivity of the graph to find “information bottlenecks”. Our next step was to pinpoint which **additional data** about the items and the users would have the highest impact.

In our user-facing application Hai, we confirmed that deep learning allows a significant breakthrough for recommendation accuracy. Although MBP data only contains sparse implicit feedback for the users, we knew using deep learning would allow correlating user tastes with all additional data sources. The first improvement we achieved is the **aggregation of multiple MBP datasets** as part of the training procedure: users’ features, items features, and additional interactions, from both their current system and their older “pre-merging” Penguin database and Random House database. We combined two different learning approaches, supervised (using interactions) and unsupervised (without interactions), to extract the most of these databases.

The second major improvement was **leveraging Hai’s dataset** to improve the general quality of the recommendations. Thanks to our B2C application Hai, we are collecting the best possible dataset to train recommender systems. This is because Hai users have full control of their data, and therefore explicitly train their own AI to get the best recommendations. As a consequence of our gamified experience, the average number of ratings per user in Hai’s dataset is 165, compared to only 3 for the MBP dataset. Merging the two datasets enabled denser and finer learning of each user’s tastes and preferences.

	AUC (higher is better)	FNR (lower is better)
LightFM	71.1%	26.5%
Crossing Minds	**92.4%**	**24.7%**

Comparison of Recommendation Accuracy Scores

Why Crossing Minds

- Recommendations are **more accurate** and lead to **higher conversion**.
- **Crossing Minds unique embedding** of the users and the items recommended allowing MBP to tackle the **cold start problem** and recommend more accurately new books or new users.
- Crossing Minds leverage the weight and algorithms trained on the **iOS app Hai** in order to provide faster and better recommendations.
- All user and item data are **fully de-identified and secure**.