Capstone: Generating Product Recommendations for small businesses at scale

Subhashree Rengarajan

Yingtian Yang

Advisors and Mentors: Neel Shivdasani, Rahul Mazumder, Hussein Hazimeh

What is MailChimp?

Mailchimp is the world's leading marketing automation platform for small businesses. To this end, the platform offers services including marketing automation, landing pages, email templates and product recommendations (affectionately known as P-REX).

MailChimp's goals are to publish the right content to the right person at the right place at the right time.

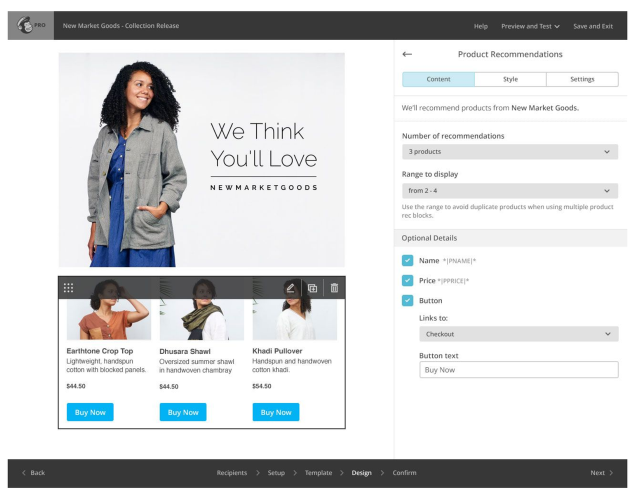

What are Personalized Product Recommendations?

Using the purchase history of each customer to make smart, data-driven predictions about what they'll want to buy in the future.

Our 1st few weeks were reviewing customer feedback about the existing system, understanding pain points, and seeing if there were ways we could improve the existing P-REX system.

Central Business Question: Can we improve the relevance of P-REX for consumers who are the recipients of Product Recommendations from MailChimp customers?

Datasets

Tens of thousands of customers use product recommendations each month.

Raw Data:

- Sample of ~1,000 stores
- Historical transactions for 3 years
- Product details, including text descriptions

Store ID	Customer ID	Product ID	Ordered At
59165197	1358525	1795	2017-02-21
59165197	1274065	1802	2017-02-01
56892273	NaN	115	2017-08-20
56892273	1432704	112	2017-09-07

Product ID	Title	Description
7523	Raven Disco Jumper	With a **high neck**, **low back**, this **dress** was designed for you in all the right places.

Cleaning and Processing:

- Removed NA's, aggregated sales for the same customer, and same products
- Transformed the datasets into user * product matrices

Purchases	Hats	Socks	Hoodies
Ben		✅	
Dan			✅
Neel		✅	✅
Mitch			✅
Jasmine	✅		

Phase 1: *Review customer feedback*

Phase 2: *Data Processing and Metrics*

Phase 3: *Choosing the right Algorithm:* ***Soft-Impute***

Phase 4: *Results and Recommendations*

Recommender Systems: Solving the problem of missing data

After exploring multiple machine learning algorithms used in recommender systems (BM25, Weighted Alternating Least Squares), we settled on Soft-Impute as it requires the fewest parameters to tune.

Soft-Impute: The idea is to impute the missing values where people have not bought anything with educated guesses while also minimizing the error on the observed values.

$$\underset{Z}{\text{minimize}} \sum_{\text{Observed}(i,j)} (X_{ij} - Z_{ij})^2 \quad \text{subject to } ||Z||_* \leq \tau$$

$$||Z||_* = \sum_j \lambda_j(Z)$$

The algorithm is based on singular value decomposition, the breakdown of a matrix into 3 submatrices, which reduces the dimensionality as well as providing some interpretability to the system.

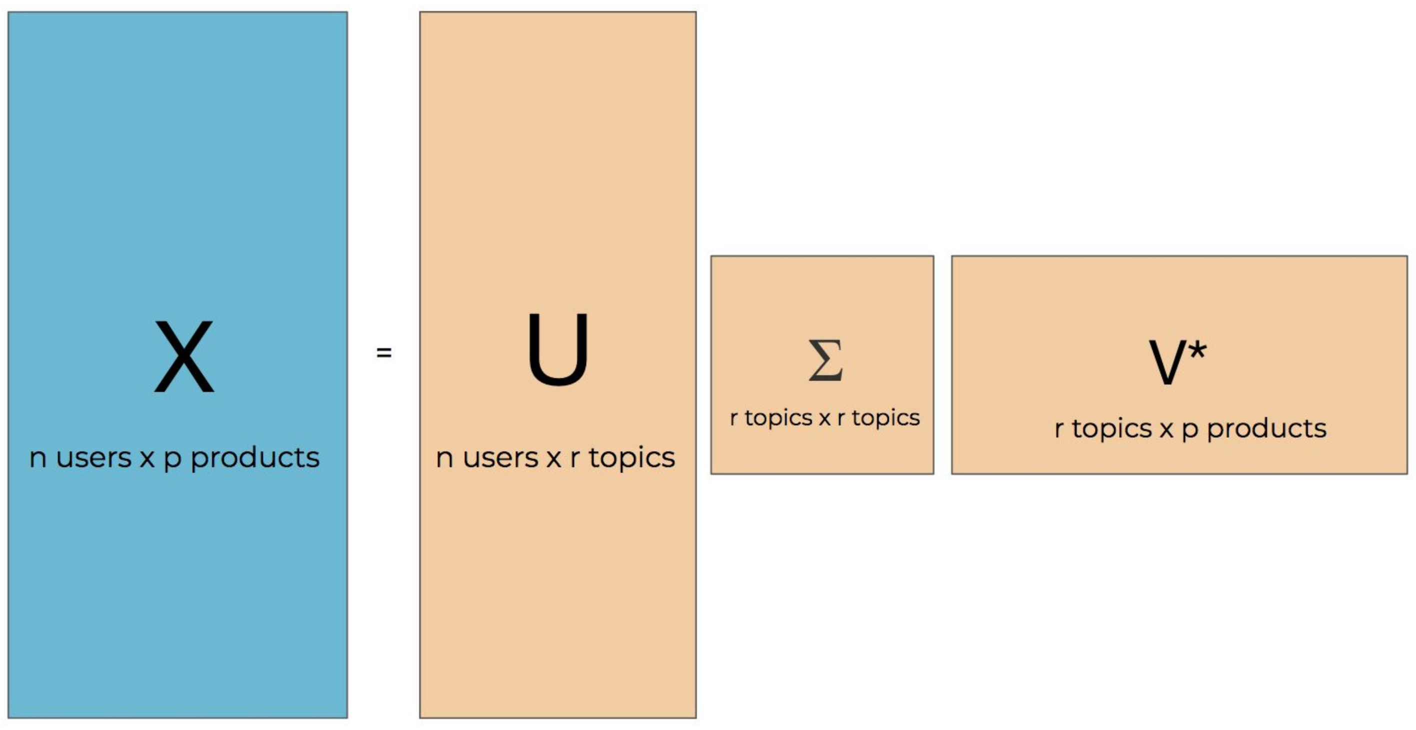

The only tuning parameter: λ, as a penalization coefficient. Similar to the penalization parameter in LASSO, here λ is a penalty on the nuclear norm $||Z||^*$. Once we generate our approximation Z, we're able to make estimations on what people will like and dislike.

Purchases	Hats	Socks	Hoodies
Ben	💖	✅	👎
Dan	👎	💖	✅
Neel	👎	✅	✅
Mitch	👎	👎	✅
Jasmine	✅	💖	💖

Testing and Results:

We tested cosine similarity, which is MailChimp's current method, and SoftImpute on small and medium stores*. The metrics we used to train the model and tune λ is NMSE:

$$Normalized\ Mean\ Square\ Error\ (NMSE) = \frac{mean\,(\sum Model\ Predictions - Observed\ Ratings)}{(Observed\ Ratings)^2}$$

With the optimal λ, we masked 20% of the purchase matrix and tested recommendations using Hit Rate @ 3: how many items the model can detect as being purchased i.e. the top 3 items likely to be purchased by the user.

We have run the Soft-Impute methodology over 74 small stores and 112 medium stores using a stratified sample,*

For both the small and medium stores, Soft-Impute outperforms the Cosine Similarity recommender system, but the difference is only statistically significant for small stores.

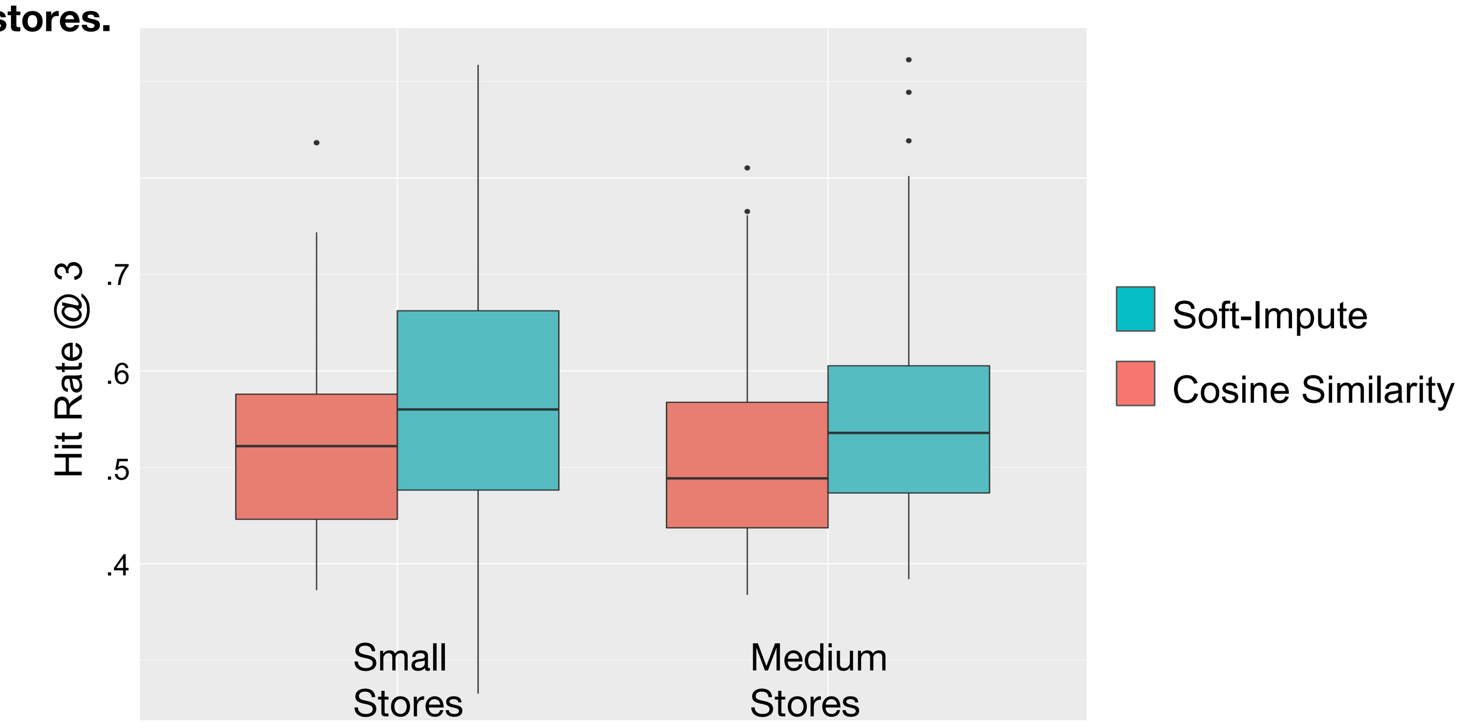

Recommendations:

Business Impact: Expansion of the P-REX feature will give MailChimp's customers a greater ability to grow their small business by using personalized e-commerce tailored to their consumers.

For MailChimp, we have observed that the most benefit would be applying Soft-Impute to the small-stores who are not already able to generate recommendations. We see a net benefit to expanding this feature to more small businesses who may not qualify for P-REX under the current schema.

Thanks for a wonderful summer in Atlanta, Georgia!

* Small stores contain 1-29 products (not including variants of size or color), medium stores contain 30 - 100 products