KDD2018 Deep Learning Day

SCHIBSTED MEDIA GROUP

Five lessons from building a deep neural network recommender for marketplaces

Simen Eide (simen.eide@schibsted.com), Audun M. Øygard (amoygard@gmail.com), Ning Zhou (ning@duoja.com)

Live demo available! Ask us!

What is an online classified marketplace?

- An internet platform where users can buy and sell new or second-hand items

What is special about building recommenders for marketplaces?

- User-generated content, often less structured and incomplete
- Freshness and location proximity are very important
- High amount of semi-unique items, i.e. high item volatility comparing to user
- Noisy transaction signals due to payments happening outside the platform

How do we recommend? - Find similar items to the one users are looking at

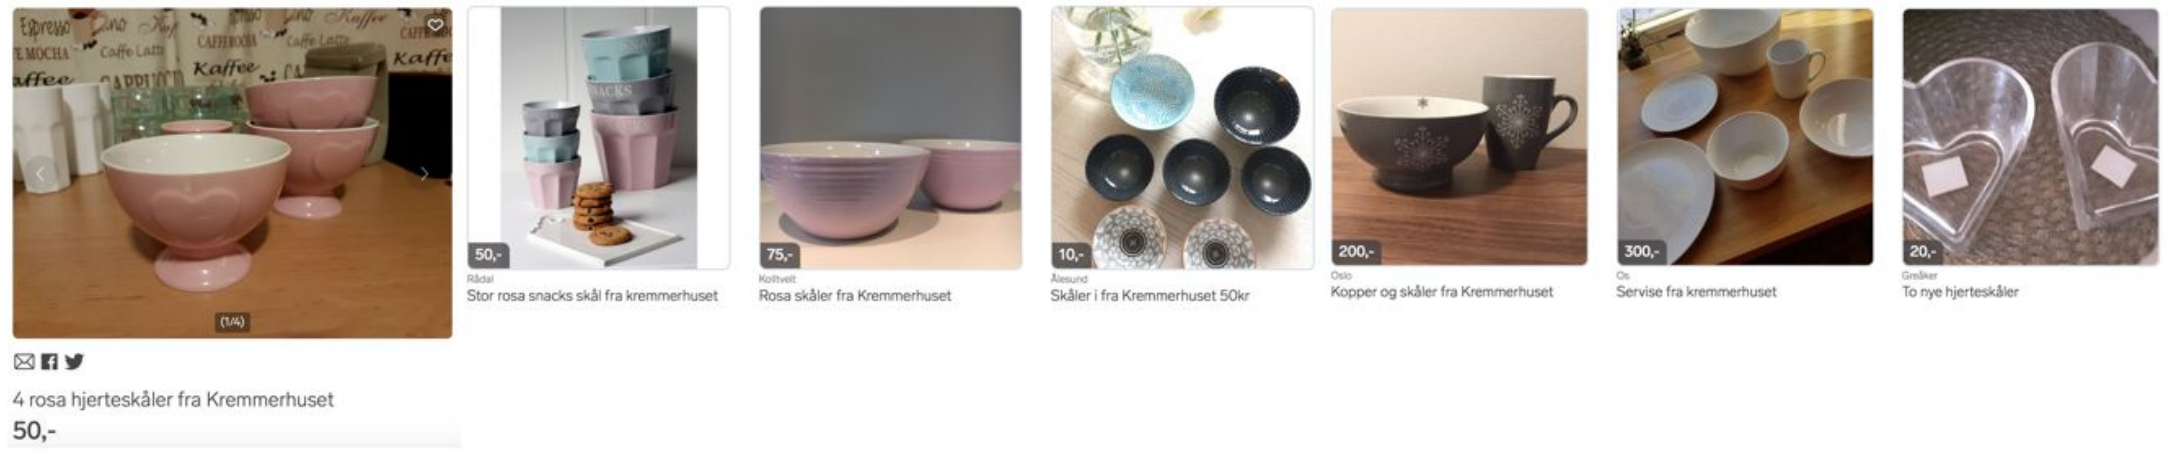

User looking at — Similar items recommended

Tactic 1: Use rich behavior signals in matrix factorization

Correct inclusion of various signals along the conversion funnel into a matrix factorization model makes a big difference

- Adjust importance of conversion signals. The platform lacks accurate transaction signals, however user behavior signals are very rich: check seller's phone number, contact seller by message or phone call, save to favorite, bookmark, follow the seller, etc.
- Use a higher vector dimension to capture the increased richness of more signals.

+7.7%

Tactic 2: Feature embeddings trained with user generated content

Pure CF suffers from cold start. Some of our items only get 10-20 views. Combine matrix factorization Item embeddings with two models:

- A matrix factorization based on user-location instead of user-item
- Text features extracted by training a deep text CNN model

+11.9%

Tactic 3: Add Transfer learning for images

Images are ranked by users among the most important factors when looking for an item

- Leverage ample external labelled data from similar domains, e.g. ImageNet, and high-quality off-the-shelf models
- Use the project of item titles in the word embedding space as image labels

+2.1%

Our final solution – a hybrid recommender using pre-trained embeddings

- Pretrain item representations from different sources:
  - Collaborative filtering item representation
  - Textual embedding trained by predicting title and item description
  - Location embedding trained from user clicks
  - Image embedding trained with transfer learning
- Dataset consist of:
  - All item pairs where a user have converted on the same day (target = 1)
  - Negative samples of item pairs not converted (target = 0)
- During training, the model tries to predict whether a user will convert on both items or not

Recommender System Overview

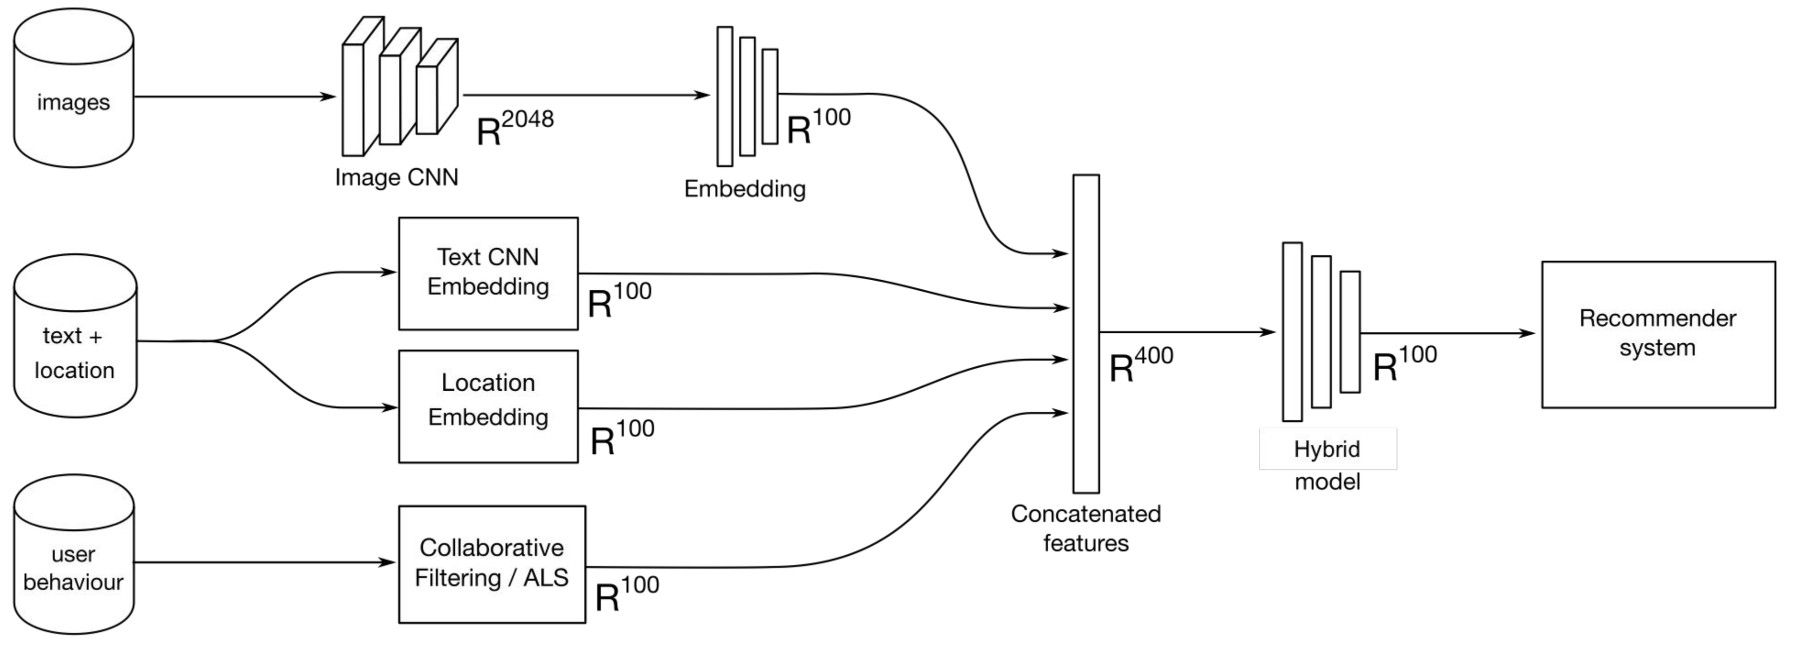

Tactic 4: Staged training strategy

Why we didn't choose end-to-end training:

- The relatively sparse and noisy user clicks are difficult to train end-to-end
- Staged training pipeline makes it easy to tune and debug module by module
- Faster training: With standalone modules, each can train in parallel.
- Practical to reuse the modules from other features such as text-based search and auto image categorization

Tactic 5: Attention mechanism

Use attention to select which features to focus on for each impression

- Focus on behavior-based features when the items already had sufficient clicks
- Focus on content-based features when it is a cold-start item
- An efficient focusing/re-weighting mechanism when some features are missing.

Marketplace Example

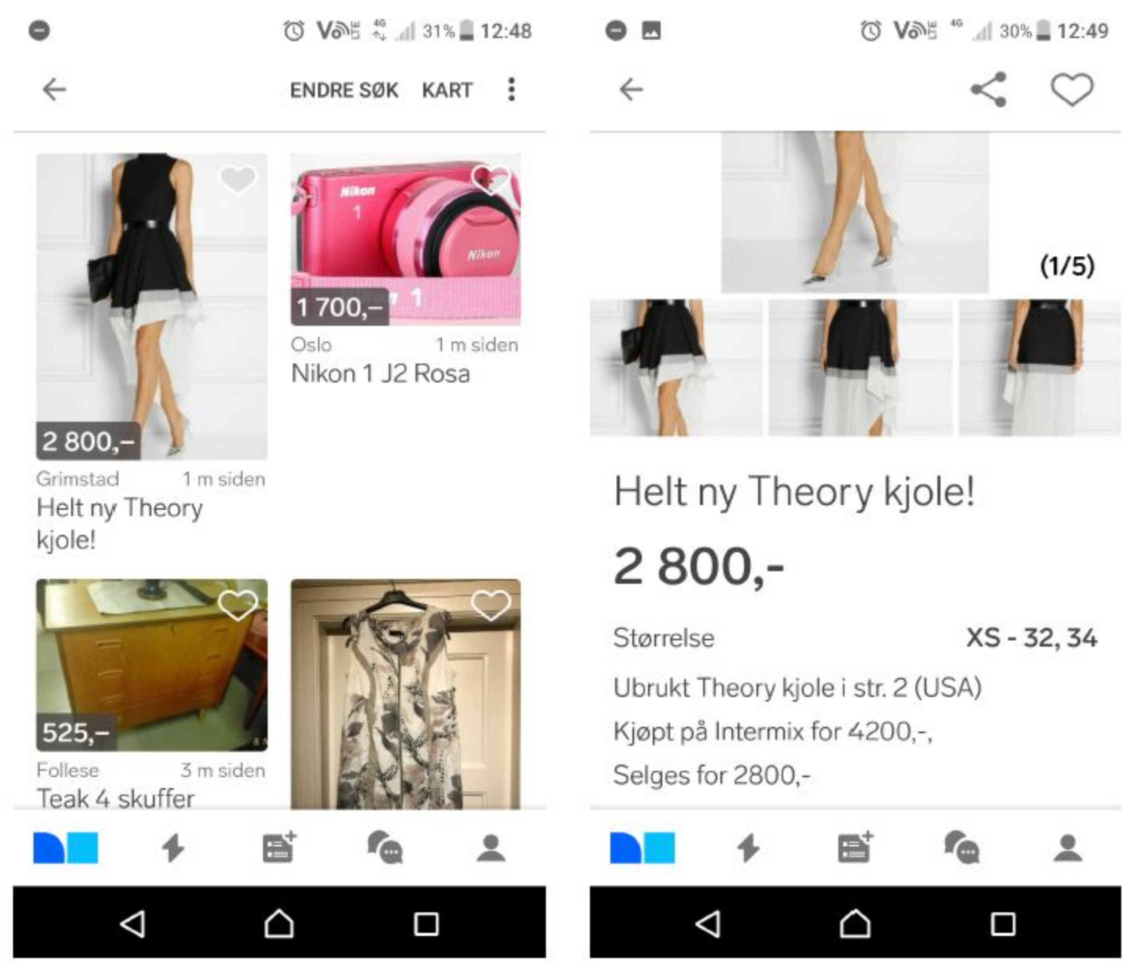

Sample Results