

Part 6: LightFM

Robert Kwieciński

OLX Group and Adam Mickiewicz University

June 5, 2020

Motivation for using hybrid recommender systems

Problems with collaborative models (e.g. classical Matrix Factorization):

- cold-start problem for users and items,
- ignores user and item features.

Motivation for using hybrid recommender systems

Problems with collaborative models (e.g. classical Matrix Factorization):

- cold-start problem for users and items,
- ignores user and item features.

Problems with content based approach:

- hard to catch dependencies which cannot be deduced from content features (maybe the same people buy bikes and chocolate cakes),
- usually worse recommendations quality.

Motivation for using hybrid recommender systems

Problems with collaborative models (e.g. classical Matrix Factorization):

- cold-start problem for users and items,
- ignores user and item features.

Problems with content based approach:

- hard to catch dependencies which cannot be deduced from content features (maybe the same people buy bikes and chocolate cakes),
- usually worse recommendations quality.

Solution

Hybrid recommender system like LightFM, which utilizes both collaborative and content data.

The full description you can read in the original paper [1].
In short words prediction is given by:

$$\hat{r}_{ui} = \mathbf{q}_u \cdot \mathbf{p}_i + b_u + b_i,$$

where \mathbf{q}_u is the user latent vector, \mathbf{p}_i is the item latent vector b_u is the user bias, b_i is the item bias.

Each user latent vector is a (weighted) sum of feature embeddings of each of his features.

The same rule applies to items and biases.

For example embedding of the item "white rabbit" might be a sum of the embedding of white and the embedding of rabbit.

Example-number of parameters

For our Movielens dataset we will use the following item features:

- genre (19 possibilities),
- data of release (240 possibilities).

As user features we'll take:

- age (61 possibilities),
- sex (2 possibilities),
- profession (21 possibilities),
- zip-code (795 possibilities).

Example-number of parameters

For our Movielens dataset we will use the following item features:

- genre (19 possibilities),
- data of release (240 possibilities).

As user features we'll take:

- age (61 possibilities),
- sex (2 possibilities),
- profession (21 possibilities),
- zip-code (795 possibilities).

If we take 10 latent factors for each user/item feature our model will have:

- $10 \cdot (19 + 240) = 2590$ parameters needed for item latent vectors,
- number of item features (259) parameters for item biases,
- $10 \cdot (61 + 2 + 21 + 795) = 8790$ parameters needed for user latent vectors,
- number of user features (879) parameters for user biases.

LightFM as a generalization of MF

We can treat as feature being a given user or an item. In this case we will have:

- number of user features equals number of users,
- number of item features equals number of items.

Then LightFM model (possibly assuming biases equals 0) reduces to MF model.

Note that we can combine both types of discussed features.

In the paper BPR: Bayesian Personalized Ranking from Implicit Feedback [2] novel loss function was introduced. It could be applied for various existing models, including Matrix Factorization.

Denote by D the set of all triples (u, i, j) , where u interacted with i , but not interacted with j .

$$\text{BPR-OPT} = \sum_{(u,i,j) \in D} \ln \sigma(\hat{r}_{ui} - \hat{r}_{uj}) - \lambda \|\Theta\|^2,$$

where \hat{r}_{ui} is estimated by a given model and Θ represents the parameter vector.

For optimizing BPR-OPT we use the stochastic gradient ascent algorithm based on bootstrap sampling of training triples.

In simple words:

- we choose triple $(u, i, j) \in D$ randomly with equal probability for each triple,
- we compute a gradient of BPR-OPT with respect to our model parameters and update the parameters (gradient ascent),
- we repeat these steps until convergence.

For optimizing BPR-OPT we use the stochastic gradient ascent algorithm based on bootstrap sampling of training triples.

In simple words:

- we choose triple $(u, i, j) \in D$ randomly with equal probability for each triple,
- we compute a gradient of BPR-OPT with respect to our model parameters and update the parameters (gradient ascent),
- we repeat these steps until convergence.

In the paper it was shown that optimizing BPR-OPT is similar to optimizing AUC.

WARP loss (Weighted Approximately Ranked Pairwise Ranking Loss) is also designed for ranking optimization by focusing on triples (u, i, j) , where the item i is a positive example for the user u and the item j is a negative example - for example user u interacted with i , but has not interacted with j .

The general idea is as follows:

- pick a random pair (u, i) where i is a positive example for u ,
- sample negative items as long as score of sampled negative item j exceeds \hat{r}_{ui} ,
- update model parameters based on a loss function

$$(\hat{r}_{uj} - \hat{r}_{ui}) \ln\left(\frac{|I| - 1}{N}\right),$$

where I is a set of all items and N is a number of samples needed to find j .

WARP loss (Weighted Approximately Ranked Pairwise Ranking Loss) is also designed for ranking optimization by focusing on triples (u, i, j) , where the item i is a positive example for the user u and the item j is a negative example - for example user u interacted with i , but has not interacted with j .

The general idea is as follows:

- pick a random pair (u, i) where i is a positive example for u ,
- sample negative items as long as score of sampled negative item j exceeds \hat{r}_{ui} ,
- update model parameters based on a loss function

$$(\hat{r}_{uj} - \hat{r}_{ui}) \ln\left(\frac{|I| - 1}{N}\right),$$

where I is a set of all items and N is a number of samples needed to find j .

Read some details and properties of WARP implementation in LightFM [3].

To do (especially for absent students):

- Go through - *P6. LightFM* notebook to:
 - check implementation of LightFM model using LightFM library - especially user and item feature preparations
 - compare the results using different loss functions
 - observe evaluation measures for LightFM model with different features

