Low-rank models for recommender systems with limited preference information.

Significance. The thesis is devoted to the development of new low-rank models for recommender systems. Its main focus are matrix- and tensor-based factorization techniques widely used in the industry. Matrices and tensors naturally arise in the collaborative filtering approach, where information about collective human behavior is used to build a latent factor-based recommendation model. In many practical cases a simple and well-known SVD-based model, called *PureSVD*, outperforms other state-of-the-art approaches. However, it also has certain limitations and drawbacks, related to its inability to reliably work with very sparse data induced by insufficient information about user interests. *The main goal of the thesis work* is to develop efficient low-rank factorization methods, which inherit the key benefits of PureSVD approach and do not suffer from its major limitations.

Novelty. Two new methods based on matrix and tensor factorizations are proposed in this work. The first method treats user feedback as an ordinal concept in contrast to commonly used cardinal representation. In this method the (*user, item, feedback*) triplets are encoded into a third order tensor which is factorized with the help of Tucker decomposition. The model allows to create a better representation of user preferences, which reduces the risk of generating irrelevant recommendations, especially when not much is known about a user. The second proposed method extends PureSVD model with side information and allows to incorporate user attributes and item features into the model. While there is a number of factorization techniques with similar functionality (the so called hybrid approach), none of them are based on PureSVD and therefore do not provide the same set of benefits. In contrast, the proposed approach allows to stay within the SVD-based computation paradigm using its generalized formulation.

Practicality. Inability to work reliably in the limited preferences setting leads to the so-called “trust busters” effect, when too many irrelevant recommendations make users feel the service is not good enough. Disappointed users may stop using the service and never want to return to it again. Both proposed approaches tackle this problem from different angles by lowering the number of irrelevant recommendations and employing additional knowledge for better predictions on relevant items.