Group title	Paper BIB	Year and Venue	Abstract
1. Titel: “Vanilla Matrix Factorization”			
1	Koren, Yehuda, Robert Bell, and Chris Volinsky. "Matrix factorization techniques for recommender systems." Computer 42.8 (2009): 30-37.	Computer 42.8 (2009)	As the Netflix Prize competition has demonstrated, matrix factorization models are superior to classic nearest-neighbor techniques for producing product recommendations, allowing the incorporation of additional information such as implicit feedback, temporal effects, and confidence levels
2	Hoyer, Patrik O. "Non-negative matrix factorization with sparseness constraints." Journal of machine learning research 5.Nov (2004): 1457-1469.	Journal of machine learning research 5.Nov (2004)	Non-negative matrix factorization (NMF) is a recently developed technique for finding parts-based, linear representations of non-negative data. Although it has successfully been applied in several applications, it does not always result in parts-based representations. In this paper, we show how explicitly incorporating the notion of 'sparseness' improves the found decompositions. Additionally, we provide complete MATLAB code both for standard NMF and for our extension. Our hope is that this will further the application of these methods to solving novel data-analysis problems.
3	Salakhutdinov, Ruslan, and Andriy Mnih. "Probabilistic matrix factorization." NIPS. Vol. 20. 2011.	NIPS. Vol. 20. 2011.	Many existing approaches to collaborative filtering can neither handle very large datasets nor easily deal with users who have very few ratings. In this paper we present the Probabilistic Matrix Factorization (PMF) model which scales linearly with the number of observations and, more importantly, performs well on the large, sparse, and very imbalanced Netflix dataset. We further extend the PMF model to include an adaptive prior on the model parameters and show how the model capacity can be controlled automatically. Finally, we introduce a constrained version of the PMF model that is based on the assumption that users who have rated similar sets of movies are likely to have similar preferences. The resulting model is able to generalize considerably better for users with very few ratings. When the predictions of multiple PMF models are linearly combined with the predictions of Restricted Boltzmann Machines models, we achieve an error rate of 0.8861, that is nearly 7% better than the score of Netflix’s own system.
2. Titel: “Cold Start”			
4	Pilászy, I. and Tikk, D. (2009). Recommending new movies: Even a few ratings are more valuab than metadata. In RecSys. 34, 49, 58	RecSys 2009	The Netflix Prize (NP) competition gave much attention to collaborative filtering (CF) approaches. Matrix factorization (MF) based CF approaches assign low dimensional feature vectors to users and items. We link CF and content-based filtering (CBF) by finding a linear transformation that transforms user or item descriptions so that they are as close as possible to the feature vectors generated by MF for CF. We propose methods for explicit feedback that are able to handle 140,000 features when feature vectors are very sparse. With movie metadata collected for the NP movies we show that the prediction performance of the methods is comparable to that of CF, and can be used to predict user preferences on new movies. We also investigate the value of movie metadata compared to movie ratings in regards of predictive power. We compare our solely CBF approach with a simple baseline rating-based predictor. We show that even 10 ratings of a new movie are more valuable than its metadata for predicting user ratings.
5	Park, Seung-Taek, and Wei Chu. "Pairwise preference regression for cold-start recommendation." Proceedings of the third ACM conference on Recommender systems. ACM, 2009.	Recsys 2009	Recommender systems are widely used in online e-commerce applications to improve user engagement and ther to increase revenue. A key challenge for recommender systems is providing high quality recommendation to use in ``cold-start" situations. We consider three types of cold-start problems: 1) recommendation on existing items for new users; 2) recommendation on new items for existing users; 3) recommendation on new items for new users. We propose predictive feature-based regression models that leverage all available information of users and items, such as user demographic information and item content features, to tackle cold-start problems. The resulting algorithms scale efficiently as a linear function of the number of observations. We verify the usefulness of our approach in three cold-start settings on the MovieLens and EachMovie datasets, by comparing with five alternatives including random, most popular, segmented most popular, and two variations of Vibes affinity algorithm widely used at Yahoo! for recommendation.
6	Li, B., Zhu, X., Li, R., Zhang, C., Xue, X., and Wu, X. (2011). Cross-domain collaborative Filtering over time. In Proceedings of the Twenty-Second international joint conference on Artificial Intelligence	RecSys 2011	Collaborative filtering (CF) techniques recommend items to users based on their historical ratings. In real-world scenarios, user interests may drift over time since they are affected by moods, contexts, and pop culture trends. This leads to the fact that a user's historical ratings comprise many aspects of user interests spanning a long time period. However, at a certain time slice, one user's interest may only focus on one or a couple of aspects. Thus, CF techniques based on the entire historical ratings may recommend inappropriate items. In this paper, we consider modeling user-interest drift over time based on the assumption that each user has multiple counterparts over temporal domains and successive counterparts are closely related. We adopt the cross-domain CF framework to share the static group-level rating matrix across temporal domains, and let user-interes distribution over item groups drift slightly between successive temporal domains. The derived method is based o a Bayesian latent factor model which can be inferred using Gibbs sampling. Our experimental results show that our method can achieve state-of-the-art recommendation performance as well as explicitly track and visualize user-interest drift over time.
3. Titel: “Online Update”			

Group title	Paper BIB	Year and Venue	Abstract
7	Sarwar, Badrul, et al. "Incremental singular value decomposition algorithms for highly scalable recommender systems." Fifth International Conference on Computer and Information Science. 2002.	CCIS 2002	We investigate the use of dimensionality reduction to improve the performance for a new class of data analysis software called "recommender systems". Recommender systems apply knowledge discovery techniques to the problem of making personalized product recommendations during a live customer interaction. The tremendous growth of customers and products in recent years poses some key challenges for recommender systems. These are:producing high quality recommendations and performing many recommendations per second for millions of customers and products. Singular Value Decomposition(SVD)-based recommendation algorithms can quickly produce high quality recommendations, but has to undergo very expensive matrix factorization steps. In this paper, we propose and experimentally validate a technique that has the potential to incrementally build SVD-based models and promises to make the recommender systems highly scalable
8	Vinagre, João, Alípio Mário Jorge, and João Gama. "Fast incremental matrix factorization for recommendation with positive-only feedback. "International Conference on User Modeling, Adaptation, and Personalization. Springer International Publishing, 2014.	UMAP 2014	Traditional Collaborative Filtering algorithms for recommendation are designed for stationary data. Likewise, conventional evaluation methodologies are only applicable in offline experiments, where data and models are static. However, in real world systems, user feedback is continuously being generated, at unpredictable rates. One way to deal with this data stream is to perform online model updates as new data points become available. This requires algorithms able to process data at least as fast as it is generated. One other issue is how to evaluate algorithms in such a streaming data environment. In this paper we introduce a simple but fast incremental Matrix Factorization algorithm for positive-only feedback. We also contribute with a prequential evaluation protocol for recommender systems, suitable for streaming data environments. Using this evaluation methodology, we compare our algorithm with other state-of-the-art proposals. Our experiments reveal that despite its simplicity, our algorithm has competitive accuracy, while being significantly faster.
9	Matuszyk, Pawel, et al. "Forgetting methods for incremental matrix factorization in recommender systems." Proceedings of the 30th Annual ACM Symposium on Applied Computing. ACM, 2015.	Symposium on Applied Computing. ACM, 2015	Numerous stream mining algorithms are equipped with forgetting mechanisms, such as sliding windows or fading factors, to make them adaptive to changes. In recommender systems those techniques have not been investigated thoroughly despite the very volatile nature of users' preferences that they deal with. We developed five new forgetting techniques for incremental matrix factorization in recommender systems. We show on eight datasets that our techniques improve the predictive power of recommender systems. Experiments with both explicit rating feedback and positive-only feedback confirm our findings showing that forgetting information is beneficial despite the extreme data sparsity that recommender systems struggle with. Improvement through forgetting also proves that users' preferences are subject to concept drift.
4. Titel: "Implicit feedback"			
10	Pilászy, István, Dávid Zibriczky, and Domonkos Tikk. "Fast als-based matrix factorization for explicit and implicit feedback datasets." Proceedings of the fourth ACM conference on Recommender systems. ACM, 2010.	Recsys 2010	Alternating least squares (ALS) is a powerful matrix factorization (MF) algorithm for both explicit and implicit feedback based recommender systems. As shown in many articles, increasing the number of latent factors (denoted by K) boosts the prediction accuracy of MF based recommender systems, including ALS as well. The price of the better accuracy is paid by the increased running time: the running time of the original version of ALS is proportional to K3. Yet, the running time of model building can be important in recommendation systems; if th model cannot keep up with the changing item portfolio and/or user profile, the prediction accuracy can be degraded. In this paper we present novel and fast ALS variants both for the implicit and explicit feedback datasets, which offers better trade-off between running time and accuracy. Due to the significantly lower computational complexity of the algorithm - linear in terms of K - the model being generated under the same amount of time is more accurate, since the faster training enables to build model with more latent factors. We demonstrate the efficiency of our ALS variants on two datasets using two performance measures, RMSE and average relative position (ARP), and show that either a significantly more accurate model can be generated under the same amount of time or a model with similar prediction accuracy can be created faster; for explicit feedback the speed-up factor can be even 5-10.
11	Fang, Yi, and Luo Si. "A latent pairwise preference learning approach for recommendation from implicit feedback." Proceedings of the 21st ACM international conference on Information and knowledge management. ACM, 2012.	CIKM 2012	Most of the current recommender systems heavily rely on explicit user feedback such as ratings on items to model users' interests. However, in many applications, it is very hard to collect the explicit feedback, while implicit feedback such as user clicks may be more available. Furthermore, it is often more suitable for many recommender systems to address a ranking problem than a rating predicting problem. This paper proposes a latent pairwise preference learning (LPPL) approach for recommendation with implicit feedback. LPPL directly models user preferences with respect to a set of items rather than the rating scores on individual items, which are modeled with a set of features by analyzing clickthrough data available in many real-world recommender systems. The LPPL approach models both the latent variables of group structure of users and the pairwise preferences simultaneously. We conduct experiments on the testbed from a real-world recommender system and demonstrate that the proposed approach can effectively improve the recommendation performance against several baseline algorithms.

Group title	Paper BIB	Year and Venue	Abstract
12	Hu, Yifan, Yehuda Koren, and Chris Volinsky. "Collaborative filtering for implicit feedback datasets." 2008 Eighth IEEE International Conference on Data Mining. Ieee, 2008.	ICDM 2008	A common task of recommender systems is to improve customer experience through personalized recommendations based on prior implicit feedback. These systems passively track different sorts of user behavior, such as purchase history, watching habits and browsing activity, in order to model user preferences. Unlike the much more extensively researched explicit feedback, we do not have any direct input from the users regarding their preferences. In particular, we lack substantial evidence on which products consumer dislike. In this work we identify unique properties of implicit feedback datasets. We propose treating the data as indication of positive and negative preference associated with vastly varying confidence levels. This leads to a factor model which is especially tailored for implicit feedback recommenders. We also suggest a scalable optimization procedure, which scales linearly with the data size. The algorithm is used successfully within a recommender system for television shows. It compares favorably with well tuned implementations of other known methods. In addition, we offer a novel way to give explanations to recommendations given by this factor model.
5. Titel: "Ranking and Recommendation"			
13	Shi, Yue, et al. "CLiMF: learning to maximize reciprocal rank with collaborative less-is-more filtering." Proceedings of the sixth ACM conference on Recommender systems. ACM, 2012.	RecSys 2012	In this paper we tackle the problem of recommendation in the scenarios with binary relevance data, when only a few (k) items are recommended to individual users. Past work on Collaborative Filtering (CF) has either not addressed the ranking problem for binary relevance datasets, or not specifically focused on improving top-k recommendations. To solve the problem we propose a new CF approach, Collaborative Less-is-More Filtering (CLiMF). In CLiMF the model parameters are learned by directly maximizing the Mean Reciprocal Rank (MRR), which is a well-known information retrieval metric for measuring the performance of top-k recommendations. We achieve linear computational complexity by introducing a lower bound of the smoothed reciprocal rank metric. Experiments on two social network datasets demonstrate the effectiveness and the scalability of CLiMF, and show that CLiMF significantly outperforms a naive baseline and two state-of-the-art CF methods.
14	Shi, Yue, Martha Larson, and Alan Hanjalic. "List-wise learning to rank with matrix factorization for collaborative filtering." Proceedings of the fourth ACM conference on Recommender systems. ACM, 2010.	Recsys 2010	A ranking approach, ListRank-MF, is proposed for collaborative filtering that combines a list-wise learning-to-ran algorithm with matrix factorization (MF). A ranked list of items is obtained by minimizing a loss function that represents the uncertainty between training lists and output lists produced by a MF ranking model. ListRank-MF enjoys the advantage of low complexity and is analytically shown to be linear with the number of observed ratings for a given user-item matrix. We also experimentally demonstrate the effectiveness of ListRank-MF by comparing its performance with that of item-based collaborative recommendation and a related state-of-the-art collaborative ranking approach (CoFiRank).
15	Lee, Guang-He, and Shou-De Lin. "LambdaMF: Learning Nonsmooth Ranking Functions in Matrix Factorization Using Lambda." Data Mining (ICDM), 2015 IEEE International Conference on. IEEE, 2015.	ICDM 2015	This paper emphasizes optimizing ranking measures in a recommendation problem. Since ranking measures are non-differentiable, previous works have been proposed to deal with this problem via approximations or lower/upper bounding of the loss. However, such mismatch between ranking measures and approximations/bounds can lead to non-optimal ranking results. To solve this problem, we propose to model the gradient of non-differentiable ranking measure based on the idea of virtual gradient, which is called lambda in learning to rank. In addition, noticing the difference between learning to rank and recommendation models, we prove that under certain circumstance the existence of popular items can lead to unlimited norm growing of the latent factors in a matrix factorization model. We further create a novel regularization term to remedy such concern. Finally, we demonstrate that our model, LambdaMF, outperforms several state-of-the-art methods. We further show in experiments that in all cases our model achieves global optimum of normalized discount cumulative gain during training. Detailed implementation and supplementary material can be found at (http://www.csie.ntu.edu.tw/~b00902055/).
6. Titel: "Collective Matrix Factorization"			
16	Bouchard, Guillaume, Dawei Yin, and Shengbo Guo. "Convex Collective Matrix Factorization." AISTATS. Vol. 13. 2013.	AISTATS 2013	In many applications, multiple interlinked sources of data are available and they cannot be represented by a single adjacency matrix, to which large scale factorization method could be applied. Collective matrix factorization is a simple yet powerful approach to jointly factorize multiple matrices, each of which represents a relation between two entity types. Existing algorithms to estimate parameters of collective matrix factorization models are based on non-convex formulations of the problem; in this paper, a convex formulation of this approach is proposed. This enables the derivation of large scale algorithms to estimate the parameters, including an iterative eigenvalue thresholding algorithm. Numerical experiments illustrate the benefits of this new approach.

Group title	Paper BIB	Year and Venue	Abstract
17	Nickel, Maximilian, Volker Tresp, and Hans-Peter Kriegel. "A three-way model for collective learning on multi-relational data." Proceedings of the 28th international conference on machine learning (ICML-11). 2011	ICML 2011	Relational learning is becoming increasingly important in many areas of application. Here, we present a novel approach to relational learning based on the factorization of a three-way tensor. We show that unlike other tensor approaches, our method is able to perform collective learning via the latent components of the model and provide an efficient algorithm to compute the factorization. We substantiate our theoretical considerations regarding the collective learning capabilities of our model by the means of experiments on both a new dataset and a dataset commonly used in entity resolution. Furthermore, we show on common benchmark datasets that our approach achieves better or on-par results, if compared to current state-of-the-art relational learning solutions, while it is significantly faster to compute
18	Singh, Ajit P., and Geoffrey J. Gordon. "Relational learning via collective matrix factorization."Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, 2008.	SIGKDD 2008	Relational learning is concerned with predicting unknown values of a relation, given a database of entities and observed relations among entities. An example of relational learning is movie rating prediction, where entities could include users, movies, genres, and actors. Relations encode users' ratings of movies, movies' genres, and actors' roles in movies. A common prediction technique given one pairwise relation, for example a #users x #movies ratings matrix, is low-rank matrix factorization. In domains with multiple relations, represented as multiple matrices, we may improve predictive accuracy by exploiting information from one relation while predicting another. To this end, we propose a collective matrix factorization model: we simultaneously factor several matrices, sharing parameters among factors when an entity participates in multiple relations. Each relation can have a different value type and error distribution; so, we allow nonlinear relationships between the parameters and outputs, using Bregman divergences to measure error. We extend standard alternating projection algorithms to our model, and derive an efficient Newton update for the projection. Furthermore, we propose stochastic optimization methods to deal with large, sparse matrices. Our model generalizes several existing matrix factorization methods, and therefore yields new large-scale optimization algorithms for these problems. Our model can handle any pairwise relational schema and a wide variety of error models. We demonstrate its efficiency, as well as the benefit of sharing parameters among relations.
7. Titel: “Time aware”			
19	Koren, Yehuda. "Collaborative filtering with temporal dynamics." Communications of the ACM53.4 (2010): 89-97.	Communications of the ACM53.4 (2010)	Customer preferences for products are drifting over time. Product perception and popularity are constantly changing as new selection emerges. Similarly, customer inclinations are evolving, leading them to ever redefine their taste. Thus, modeling temporal dynamics is essential for designing recommender systems or general customer preference models. However, this raises unique challenges. Within the ecosystem intersecting multiple products and customers, many different characteristics are shifting simultaneously, while many of them influenc each other and often those shifts are delicate and associated with a few data instances. This distinguishes the problem from concept drift explorations, where mostly a single concept is tracked. Classical time-window or instance decay approaches cannot work, as they lose too many signals when discarding data instances. A more sensitive approach is required, which can make better distinctions between transient effects and long-term patterns. We show how to model the time changing behavior throughout the life span of the data. Such a model allows us to exploit the relevant components of all data instances, while discarding only what is modeled as being irrelevant. Accordingly, we revamp two leading collaborative filtering recommendation approaches. Evaluation is made on a large movie-rating dataset underlying the Netflix Prize contest. Results are encouraging and better than those previously reported on this dataset. In particular, methods described in this paper play a significant role in the solution that won the Netflix contest.
20	Karatzoglou, Alexandros, et al. "Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering." Proceedings of the fourth ACM conference on Recommender systems. ACM, 2010.	RecSys 2010	Context has been recognized as an important factor to consider in personalized Recommender Systems. However, most model-based Collaborative Filtering approaches such as Matrix Factorization do not provide a straightforward way of integrating context information into the model. In this work, we introduce a Collaborative Filtering method based on Tensor Factorization, a generalization of Matrix Factorization that allows for a flexible and generic integration of contextual information by modeling the data as a User-Item-Context N-dimensional tensor instead of the traditional 2D User-Item matrix. In the proposed model, called Multiverse Recommendation different types of context are considered as additional dimensions in the representation of the data as a tensor. The factorization of this tensor leads to a compact model of the data which can be used to provide context-aware recommendations.We provide an algorithm to address the N-dimensional factorization, and show that the Multiverse Recommendation improves upon non-contextual Matrix Factorization up to 30% in terms of the Mean Absolute Error (MAE). We also compare to two state-of-the-art context-aware methods and show that Tensor Factorization consistently outperforms them both in semi-synthetic and real-world data - improvements range from 2.5% to more than 12% depending on the data. Noticeably, our approach outperforms other methods by a wider margin whenever more contextual information is available.

Group title	Paper BIB	Year and Venue	Abstract
21	Pragarauskas, H., and Oliver Gross. "Temporal collaborative filtering with bayesian probabilistic tensor factorization." SIAM (2010).	SIAM 2010	Real-world relational data are seldom stationary, yet traditional collaborative filtering algorithms generally rely o this assumption. Motivated by our sales prediction problem, we propose a factor-based algorithm that is able to take time into account. By introducing additional factors for time, we formalize this problem as a tensor factorization with a special constraint on the time dimension. Further, we provide a fully Bayesian treatment to avoid tuning parameters and achieve automatic model complexity control. To learn the model we develop an efficient sampling procedure that is capable of analyzing large-scale data sets. This new algorithm, called Bayesian Probabilistic Tensor Factorization (BPTF), is evaluated on several real-world problems including sales prediction and movie recommendation. Empirical results demonstrate the superiority of our temporal model.
8. Titel: "Social Regularization"			
22	Ma, Hao, et al. "Sorec: social recommendation using probabilistic matrix factorization." Proceedings of the 17th ACM conference on Information and knowledge management. ACM, 2008.	CIKM 2008	Data sparsity, scalability and prediction quality have been recognized as the three most crucial challenges that every collaborative filtering algorithm or recommender system confronts. Many existing approaches to recommender systems can neither handle very large datasets nor easily deal with users who have made very few ratings or even none at all. Moreover, traditional recommender systems assume that all the users are independent and identically distributed; this assumption ignores the social interactions or connections among users. In view of the exponential growth of information generated by online social networks, social network analysis is becoming important for many Web applications. Following the intuition that a person's social network will affect personal behaviors on the Web, this paper proposes a factor analysis approach based on probabilistic matrix factorization to solve the data sparsity and poor prediction accuracy problems by employing both users' social network information and rating records. The complexity analysis indicates that our approach can be applied to very large datasets since it scales linearly with the number of observations, while the experimental results shows that our method performs much better than the state-of-the-art approaches, especially in the circumstance that users have made few or no ratings.
23	Jamali, Mohsen, and Martin Ester. "A matrix factorization technique with trust propagation for recommendation in social networks." Proceedings of the fourth ACM conference on Recommender systems. ACM, 2010.	RecSys 2010	Recommender systems are becoming tools of choice to select the online information relevant to a given user. Collaborative filtering is the most popular approach to building recommender systems and has been successfully employed in many applications. With the advent of online social networks, the social network based approach to recommendation has emerged. This approach assumes a social network among users and makes recommendations for a user based on the ratings of the users that have direct or indirect social relations with the given user. As one of their major benefits, social network based approaches have been shown to reduce the problems with cold start users. In this paper, we explore a model-based approach for recommendation in social networks, employing matrix factorization techniques. Advancing previous work, we incorporate the mechanism o trust propagation into the model. Trust propagation has been shown to be a crucial phenomenon in the social sciences, in social network analysis and in trust-based recommendation. We have conducted experiments on two real life data sets, the public domain Epinions.com dataset and a much larger dataset that we have recently crawled from Flixster.com. Our experiments demonstrate that modeling trust propagation leads to a substantial increase in recommendation accuracy, in particular for cold start users
24	Ma, Hao, et al. "Recommender systems with social regularization." Proceedings of the fourth ACM international conference on Web search and data mining. ACM, 2011.	WSDM 2011	Although Recommender Systems have been comprehensively analyzed in the past decade, the study of social-based recommender systems just started. In this paper, aiming at providing a general method for improving recommender systems by incorporating social network information, we propose a matrix factorization framework with social regularization. The contributions of this paper are four-fold: (1) We elaborate how social network information can benefit recommender systems; (2) We interpret the differences between social-based recommender systems and trust-aware recommender systems; (3) We coin the term Social Regularization to represent the social constraints on recommender systems, and we systematically illustrate how to design a matrix factorization objective function with social regularization; and (4) The proposed method is quite general, which can be easily extended to incorporate other contextual information, like social tags, etc. The empirical analysis on two large datasets demonstrates that our approaches outperform other state-of-the-art methods.
9. Titel: " Recommendation in signed social networks"			

Group title	Paper BIB	Year and Venue	Abstract
25	Tang, J., Aggarwal, C., & Liu, H. (2016, April). Recommendations in signed social networks. InProceedings of the 25th International Conference on World Wide Web (pp. 31-40). International World Wide Web Conferences Steering Committee.	www 2016	Recommender systems play a crucial role in mitigating the information overload problem in social media by suggesting relevant information to users. The popularity of pervasively available social activities for social media users has encouraged a large body of literature on exploiting social networks for recommendation. The vast majority of these systems focus on unsigned social networks (or social networks with only positive links), while little work exists for signed social networks (or social networks with positive and negative links). The availability of negative links in signed social networks presents both challenges and opportunities in the recommendation process. We provide a principled and mathematical approach to exploit signed social networks for recommendation, and propose a model, RecSSN, to leverage positive and negative links in signed social networks. Empirical results on real-world datasets demonstrate the effectiveness of the proposed framework. We also perform further experiments to explicitly understand the effect of signed networks in RecSSN.
26	Song, D., & Meyer, D. A. (2015, January). Recommending Positive Links in Signed Social Networks by Optimizing a Generalized AUC. InAAAI (pp. 290-296).	AAAI 2015	With the rapid development of signed social networks in which the relationships between two nodes can be either positive (indicating relations such as like) or negative (indicating relations such as dislike), producing a personalized ranking list with positive links on the top and negative links at the bottom is becoming an increasingly important task. To accomplish it, we propose a generalized AUC (GAUC) to quantify the ranking performance of potential links (including positive, negative, and unknown status links) in partially observed signed social networks. In addition, we develop a novel link recommendation algorithm by directly optimizing the GAUC loss.We conduct experimental studies based upon Wikipedia, MovieLens, and Slashdot; our results demonstrate the effectiveness and the efficiency of the proposed approach.
27	Song, D., Meyer, D. A., & Tao, D. (2015, August). Efficient latent link recommendation in signed networks. In Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1105-1114). ACM.	SIGKDD2015	Signed networks, in which the relationship between two nodes can be either positive (indicating a relationship such as trust) or negative (indicating a relationship such as distrust), are becoming increasingly common. A plausible model for user behavior analytics in signed networks can be based upon the assumption that more extreme positive and negative relationships are explored and exploited before less extreme ones. Such a model implies that a personalized ranking list of latent links should place positive links on the top, negative links at the bottom, and unknown status links in between. Traditional ranking metrics, e.g., area under the receiver operating characteristic curve (AUC), are however not suitable for quantifying such a ranking list which includes positive, negative, and unknown status links. To address this issue, a generalized AUC (GAUC) which can measure both the head and tail of a ranking list has been introduced. Since GAUC weights each pairwise comparison equally and the calculation of GAUC requires quadratic time, we derive two lower bounds of GAUC which can be computed in linear time and put more emphasis on ranking positive links on the top and negative links at the bottom of a ranking list. Next, we develop two efficient latent link recommendation (ELLR) algorithms in order to recommend links by directly optimizing these two lower bounds, respectively. Finally, we compare these two ELLR algorithms with top-performing baseline methods over four benchmark datasets, among which the largest network has more than 100 thousand nodes and seven million entries. Thorough empirical studies demonstrate that the proposed ELLR algorithms outperform stateof-the-art approaches for link recommendation in signed networks at no cost in efficiency
10. Titel: "Graph knowledge"			
28	Menon, Aditya Krishna, and Charles Elkan. "Link prediction via matrix factorization." Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Springer Berlin Heidelberg, 2011.	ECML 2011	We propose to solve the link prediction problem in graphs using a supervised matrix factorization approach. The model learns latent features from the topological structure of a (possibly directed) graph, and is shown to make better predictions than popular unsupervised scores. We show how these latent features may be combined with optional explicit features for nodes or edges, which yields better performance than using either type of feature exclusively. Finally, we propose a novel approach to address the class imbalance problem which is common in lin prediction by directly optimizing for a ranking loss. Our model is optimized with stochastic gradient descent and scales to large graphs. Results on several datasets show the efficacy of our approach
29	Gu, Quanquan, Jie Zhou, and Chris HQ Ding. "Collaborative Filtering: Weighted Nonnegative Matrix Factorization Incorporating User and Item Graphs." SDM. 2010.	SDM 2010	Collaborative filtering is an important topic in data mining and has been widely used in recommendation system. In this paper, we proposed a unified model for collaborative filtering based on graph regularized weighted nonnegative matrix factorization. In our model, two graphs are constructed on users and items, which exploit the internal information (e.g. neighborhood information in th user-item rating matrix) and external information (e.g. content information such as user's occupation and item's genre, or other kind of knowledge such as social trust network). The proposed method not only inherits the advantages of model-based method, but also owns the merits of memory-based method which considers the neighborhood information. Moreover, it has the ability to make use of content information and any additional information regarding user-user such as social trust network. Due to the use these internal and external information, the proposed method is able to find more interpretable lowdimensional representations fo users and items, which is helpful for improving the recommendation accuracy. Experimental results on benchmark collaborative filtering data sets demonstrate that the proposed methods outperform the state of the art collaborative filtering methods a lot

Group title	Paper BIB	Year and Venue	Abstract
30	Nickel, Maximilian, Volker Tresp, and Hans-Peter Kriegel. "Factorizing yago: scalable machine learning for linked data." Proceedings of the 21st international conference on World Wide Web. ACM, 2012.	WWW2012	Vast amounts of structured information have been published in the Semantic Web's Linked Open Data (LOD) cloud and their size is still growing rapidly. Yet, access to this information via reasoning and querying is sometimes difficult, due to LOD's size, partial data inconsistencies and inherent noisiness. Machine Learning offers an alternative approach to exploiting LOD's data with the advantages that Machine Learning algorithms are typically robust to both noise and data inconsistencies and are able to efficiently utilize non-deterministic dependencies in the data. From a Machine Learning point of view, LOD is challenging due to its relational nature and its scale. Here, we present an efficient approach to relational learning on LOD data, based on the factorization of a sparse tensor that scales to data consisting of millions of entities, hundreds of relations and billions of known facts. Furthermore, we show how ontological knowledge can be incorporated in the factorization to improve learning results and how computation can be distributed across multiple nodes. We demonstrate that our approach is able to factorize the YAGO~2 core ontology and globally predict statements for this large knowledge base using a single dual-core desktop computer. Furthermore, we show experimentally that our approach achieves good results in several relational learning tasks that are relevant to Linked Data. Once a factorization has been computed, our model is able to predict efficiently, and without any additional training, the likelihood of any of the 4.3 · 1014 possible triples in the YAGO~2 core ontology.