MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Multidimensional Visualization for Collaborative Filtering Recommender Systems

Frederick J. Igo Jr., Matthew Brand, Kent Wittenburg, David Wong and Shinsuke Azuma

TR2002-39 September 2002

Abstract

This paper describes a movie recommender system that demonstrates both an incremental SVD prediction algorithm as well as multidimensional visualization of the prediction results. The recommender system makes accurate recommendations in real-time and supports incremental updating, also in real time. In addition the movie recommendation system demonstrates the advantages of multidimensional visualization of the recommender system's results. Recommendations for all movies, along with several attributes of the movies being recommended, are viewed in parallel as bargrams in the multidimensional visualization tool. This tool allows the user to browse the recommendations and filter the results using multiple movie attributes at once. The user can dynamically navigate to a short list of recommendations with only a few clicks. By adding recommendations for other users as additional dimensions we also demonstrate how the system can recommend for groups just as easily as for individuals.

This extended abstract has been submitted as a poster to InfoVis 2002.

Submitted August 2002.

Multidimensional Visualization for Collaborative Filtering Recommender Systems (Extended Abstract)

Frederick J. Igo, Jr., Matthew Brand,
Kent Wittenburg and David Wong
Mitsubishi Electric Research Laboratories
{igo,brand,wittenburg,wong}@merl.com

Shinsuke Azuma
Mitsubishi Electric Corp.
azuma@isl.melco.co.jp

Abstract

This paper describes a movie recommender system that demonstrates both an incremental SVD prediction algorithm [1] as well as multidimensional visualization of the prediction results. The recommender system makes accurate recommendations in real-time and supports incremental updating, also in real-time. In addition the movie recommendation system demonstrates the advantages of multidimensional visualization of the recommender system's results. Recommendations for all movies, along with several attributes of the movies being recommended, are viewed in parallel as bargrams [2] in the multidimensional visualization tool. This tool allows the user to browse the recommendations and filter the results using multiple movie attributes at once. The user can dynamically navigate to a short list of recommendations with only a few clicks. By adding recommendations for other users as additional dimensions we also demonstrate how the system can recommend for groups just as easily as for individuals.

1. Introduction

Most online recommender systems take ratings from a user and return the top few (three to fifteen) recommendations for that user. Unfortunately these top recommendations can fail to match the desires of the user if the recommendations have not been narrowed down to the user's interest at the moment. Some online recommender systems have started to address this problem by allowing the user to focus the recommendations by specifying a category or sub-category, such as movie genre. Our movie recommender system takes that idea and expands upon it. Instead of letting the user see only the top recommendations within one category, we return recommendations for all movies in the system, from worst to best, in real-time, and allow the user to browse and filter the recommendations interactively with a multidimensional visualization tool which includes several movie attributes, such as genre, year released, runtime, language and MPAA rating. Now the user becomes a part of the recommending system and can dynamically guide the system to the most appropriate result.

By adding recommendations for other user's as additional dimensions, the visualizer can quickly transform from being a recommender system for individuals to a recommender system for groups of users. The visualization tool displays recommendations for multiple users along side movie attributes as parallel dimensions. The users are able to filter and select using both movie attributes and recommendations for users, quickly narrowing down a long list to a handful of accurate recommendations that meet the interest of the entire group.

2. IISVD: Incremental Imputative SVD

Our movie recommender system uses MERL's novel Incremental Imputative Singular Value Decomposition (IISVD)[1] as an engine for analyzing ratings and making predictions. A "thin" SVD is a compact representation of how ratings are correlated across people and across movies. It contains information that could be characterized as "people who liked movie X and disliked movie Y tend to like movie Z more than average." IISVD computes an SVD incrementally from partial data in linear time--thereby making SVDs practical for large online consumer data applications. This allows our system to update its model of people's tastes in movies in real-time without compromising the accuracy of the system. Ratings from new users (and of new movies) are immediately reflected in future recommendations.

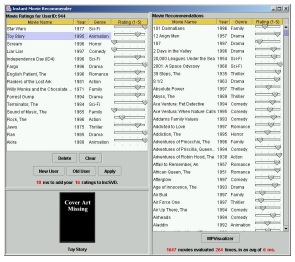

Figure 1. IISVD Movie Recommender

In addition to being incrementally updatable, IISVD is also fast and accurate when recommending. Our movie recommender system re-computes the recommendations for all movies in the system (~1600) every time the user moves one movie's rating just one rating point. IISVD can compute recommendations for all the movies in less than 10 ms, which allows the system's user-interface to react in real-time as the user adjusts the ratings. Using a standard benchmarking movie database, the IISVD was found to accurately predict a user's rating (within 1 point on a 1-5 scale) 82% of the time.

3. MPV: Multi-Parametric Visualization

Dealing with a complete list of recommendations can be clumsy and unwieldy without the proper tools. We use a tool we call Multi-Parametric Visualization (MPV), based on Verizon's EZChooser[tm] [2], to browse the complete list of movie recommendations. This tool displays in parallel the system's recommendations for the user as well as various attributes of the movies, all as individual dimensions. Each dimension is split into segments and displayed as bargrams, which are individually selectable by the user during navigation.

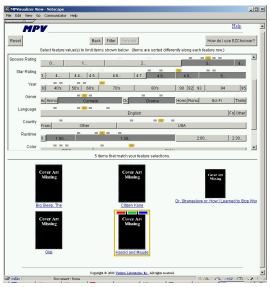

Figure 2. Multi-Parametric Visualization

Unlike a query interface where the user enters queries in a trial and error sequence until a satisfactory result set is discovered, MPV is used to dynamically build a 'query' by clicking on the bargrams. For example, the user could select 'drama' from the genre dimension, 'English' from the language dimension and '5 stars' from the recommendation dimension to discover a list of recommended dramas in English. If the list is still too large the user can further restrict the query by selecting other attributes, such as runtime or MPAA rating. If the result set is too small or empty, the user can relax the query by adding '4 star' recommendations.

This tool can also be used to visually discover relationships previously unknown to the user. For example, selecting highly rated movies may show that a user prefers movies from the 1980's since the majority of the user's highest recommendations fall within the 1980's bargram. The ability to navigate, browse and discover interactively allows the user to search without knowing in advance exactly which attributes will be important.

Another concept we demonstrate with our movie recommendation system is that the system can be used to make recommendations for groups just as easily as it can for individuals. If two or more users individually rate at least 15-20 movies, the recommendations from IISVD for each user can be viewed in parallel along with the movie attributes. MPV allows the users to select which ratings and attributes are important for making a decision for the group. The result is a set of recommendations that meet the needs of all users in the group.

4. Future Work

Our plan for future work includes integrating new features into the recommending system, such as downdating, and improving the multidimensional visualization by adding support for dimensions with sets or hierarchy.

5. Conclusion

Our movie recommender system demonstrates how the addition of multidimensional visualization to recommender systems can improve the overall system for individuals and as well as groups. Movies have been used in this demonstration since movie recommending can be understood more easily by a broad audience, however we hope to apply this combination of visualization and recommending to other domains such as risk management and marketing.

6. Availability

http://brisa.merl.com/mpv

7. Acknowledgements

Thanks to the GroupLens Research Project for making the movie data available.

8. References

[1] Brand, M.E., "Incremental Singular Value Decomposition of Uncertain Data with Missing Values", European Conference on Computer Vision (ECCV), Vol 2350, pps. 707-720, May 2002 (Lecture Notes in Computer Science)

[2] Wittenburg, et al, "Parallel Bargrams for Consumer-based Information Exploration and Choice". In Proceedings of the ACM Symposium on User Interface Systems and Technology (UIST '01), Orlando, Florida, November 2001, pp. 51-60.