Visualization and exploration for recommender systems in enterprise organization

Z. Karni[a], L. Shapira[b]
[a]HP Lab Israel, Technion City, Haifa, Israel; [b]Microsoft Research, Redmond, WA USA

ABSTRACT

Recommender systems seek to predict the interest a user would find in an item, person or social element they had not yet considered, based upon the properties of the item, the user's past experience and similar users. However, recommended items are often presented to the user with no context and no ability to influence the results. We present a novel visualization technique for recommender systems in which, a user can see the items recommended for him, and understand why they were recommended. Focusing on a user, we render a planar visualization listing a set of recommended items. The items are organized such that similar items reside nearby on the screen, centered around real-time generated categories. We use a combination of iconography, text and tag clouds, with maximal use of screen real estate, and keep items from overlapping to produce our results. We apply our visualization to expert relevance maps in the enterprise and a book recommendation system for consumers. The latter is based on Shelfari, a social network for reading and books.

Keywords: Graph, Social Network, Expert Relevance, Visual Layout, Similarity Measure, Recommender systems

1. INTRODUCTION

Recommender systems utilize information analytics techniques to predict the interests of its users based on the experience that has been acquired directly from them or indirectly by the system's activity. Typically, when a user logs in, the system analyzes the interactions and builds a representative profile. This profile is compared against the rest of the profiles in the system, and similarity is used for recommendations.

The majority of recommender systems use a ranking of interests sorted by their relevance. For example, NetFlix, a popular website for movie rentals, presents a list of recommended movies, sorted by "importance"'. However, even though the list may reflect the general interest of the user in movies, additional information would improve the recommendations. For example, depending on the mood, the user may be interested in a comedy. However, it may be that no comedies were ranked as relevant and therefore none will appear in the recommendation list.

The visualization challenge increases when considering social networks. Social networks usually consist of many layers of information. A similarity measure may be defined between entities of the same layer, and connections may exist between entities from different layers. The connections in social networks are usually represented in a graph and although many algorithms in graph drawing exist, general algorithms do not capture the special structure in social networks. Moreover, recently there has been a growing interest in enterprise social graphs, connecting employees and experts within the organization, in effect mapping the organizational knowledge. Here, it is even more important for a user, to lay out the structure of his social graph, illustrating the context of each connection.

In this paper, we present an interactive visualization for recommender systems. It enables a user of the system to easily locate other users sharing similar interests. The system consists of user entities, interest entities and auxiliary entities. Interest entities are connected using a similarity measure. Users are connected to interest entities and other entities via secondary connections. The auxiliary entities are usually added to enhance the context of given parts of the graph. Given a user entity in the network on which to focus, the system draws a graphical layout with users and interests that reflect relevance and similarity to the central user entity (Figure 1). Utilizing the graphical layout, the user can easily deduce who are the most relevant users in the network and what is the basis for their relevance. The main contributions in this paper are:

- A novel visualization technique which puts recommendations in context.
- An algorithm to lay out dense adjacency matrices with correlation to the actual distances.

Updated 1 March 2012

Figure 1: Visualization for recommender systems : (left) A profile at an enterprise social network and (right) A profile in the Shelfari social network

2. RELATED WORK

Recommender systems are a tool within the information analytics portfolio that come to assist in handling with the proliferation of digital data and information from the all variety of resources, types and content. Many algorithms have been developed especially in the last decade to try to face this challenge and they are mostly divided into content-based, collaborative and hybrid recommendation. Although recommendation algorithms and applications are beyond the scope of our paper, the interested reader is referred to this representative surveys and books [1][2][3].

From the surveys, we learn that many of the recommender algorithms use scalar ranking. The ranking is related to the entity in focus and common visualization techniques and UIs use sorted lists to show the recommendation results. This visualization is simple and provides immediate conclusion drawing. However, due to its simplicity, the underlying structure of the system and the reason for the conclusion remain unclear to the user and the amount of control a user has to refine the results is limited. Over the years, several ideas were suggested to improve the original linear sorting and provide a better user experience and informative views. For example, using color-coding [4], dividing the visual real-estate to several different views [5][6] and using dynamic gestures and animation [7][8].

In many cases the underlying data structure of recommender systems is a graph or a tree. It is of course possible to traverse the graph in a meaningful order and to display the recommendation results linearly as a list. However, the graph structure includes important knowledge such a topology and connections, which may be vital to the user to see and understand. Graph drawing and visualization is a well-established field in computer science. Automatic graph drawing is commonly an optimization problem that tries to position the graph nodes in the 2D plane and to connect them with the graph's edges in a way that meets an aesthetic criteria and rules [7][9]. Due to the complex structure of the graph, edge intersection and node overlap cannot be avoided. This may render the visualization useless since it affects the graph's readability and understanding of its structure. Beside aesthetic considerations, graph layouts are usually hard to control. Graphs are commonly modeled using masses and springs and unless proper constrains are posed on the model springs tends to shrink pulling entire nodes after them, leaving a significant part of the rendering plane empty.

Social networks place a modern twist on interaction between individuals, small businesses and even enterprises. For the latter two analyzing the network is critical since it can guide marketing and productization. The underlying graph of social networks is usually characterized by global sparseness and local dense regions. Visualizing such data using graph visualization techniques does not utilize space efficiently, and typically obfuscates information within local regions. Therefore dedicated algorithms and approaches are required. In [10] given a selected item (usually a user), the focus is on a small region most relevant to that item. Moreover, the rendering is preformed according to predefined preferences and not dynamically adaptive.

3. OUR APPROCH

Given a user and a set of recommendations, our objective is to present the information in a way that is informative, aesthetic and interactive. Our approach attempts to put each recommendation in context (informative), maximize usage of screen real estate while avoiding clutter (aesthetic) and empower the user to request additional information and update the visualization accordingly (interactive). Our solution consists of the following steps:

Figure 2: (Left) Layers, from back to front: anchors, auxiliary, users and interests. (Right) composed visualization.

1) Receive a relevance graph for the selected user

2) Enhance the graph with auxiliary information, emphasizing the context of each recommendation

3) Layout the relevance graph on the available screen real estate.

For a selected user *u* we receive a relevance graph $G_u=(V,E_w)$. Each node in the graph $v \in V$ represents an entity in the system of a certain type. The edges of the graph E_w represent symbolic relationships such as manager-employee, reader-book, author-paper, or weighted similarity between entities. We usually augment the relevance graph with additional information, adding context to the recommendations. We add both nodes representing contextual information such as categories or tags, and new edges which define the context.

Our layout algorithm separates the different entity types within the relevance graph into layers, and performs the layout one layer at a time (Figure 2-Left). The base layer is termed the anchor layer. This layer contains entities for which we have a similarity (or distance) measure. Using dimensionality reduction measures such as multi-dimensional scaling (MDS) we embed these anchor entities in 2D. The embedding serves as seed coordinates for each anchor entity. We then traverse the relevance graph from each anchor entity, calculating minimum-spanning tree (MST) to the rest of the nodes. The seed coordinates for all the remaining entities are a weighted combination of the reachable anchor entities. Finally, we implement a packing algorithm that ensures that there is no overlap between entities, and that the distance between each entity and its seed location, and hence its related entities, is preserved (Figure 2-Right).

4. APPLICATIONS

We apply our visualization algorithm to two very different recommender systems: The first is an expert relevance mapping for the enterprise, created for a HP internal conference (Figure 3). We used the complete archives of the conference to calculate, for a conference attendee, an expert relevance graph. The graph is based on the user's own work (submitted and presented papers) and contains papers and experts of interest. The abstracts are categorized into roughly 70 subjects, which have been refined over the years. Our goal was to create an application accessible by all attendees, which would analyze a user's past contributions and suggest interesting people (experts) to meet and papers to watch for. The application was distributed to 750 attendees of the conference and was well received. We received many positive feedbacks from people who have experienced the visualization. The visualization exposed them to people of interest they were not aware of, and papers they would otherwise have missed.

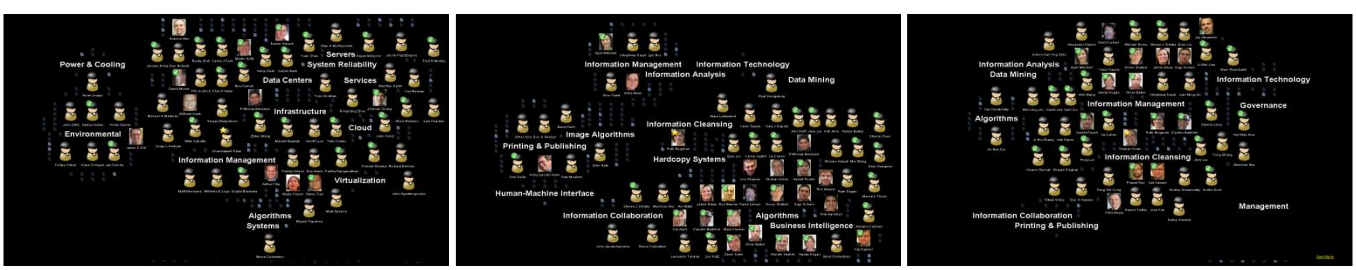

Figure 3: Examples of the expert relevance mapping for the enterprise application representing the visualization for three different users.

Figure 4: Examples out of the Shelfari application representing the visualization of two different users.

The second is a book recommendation system, based on collaborative filtering. We used Shelfari (www.shelfari.com), which is a book-oriented social network. Users can rate books they have read, write reviews, synopsis and character descriptions, and add tags to books. Users can form groups on various topics, and follow the reading of their friends. Currently, Shelfari only offers as recommendations books highly rated by your friends or books by authors you've already read.

The Shelfari application (Figure 4) was tested by generating reading recommendations for several long time Shelfari users. They were asked whether the visualization reflected their taste in reading, and would they read some of the recommended books. Their responses were positive, and they expressed interest in adding interactivity to the visualization, allowing one to focus on specific types of books. Specifically one noted that since he never marks books he dislikes in Shelfari, the application suggested some false positives i.e. authors he doesn't favor.

5. CONCLUSIONS

We presented visualization for recommender systems. Unlike conventional recommendations which are one-dimensional and usually lack in context, we allow a user to explore recommended items in context. Recommendations are laid out together with context enhancing information (such as tags and categories). Interactivity is added, allowing a user to influence the recommendations and receive more information. We demonstrated the visualization on two very different recommender systems, an enterprise expert relevance suggestion, and an application for book recommendations, both of which proved highly successful.

REFERENCES

[1] Adomavicius, G. and Tuzhilin, A., "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions", IEEE Transactions on Knowledge and Data Engineering, 17(6):734–749 (2005).

[2] Ricci, F., Rokach, L., Shapira, B., and Kantor, P. B., "Recommender Systems Handbook", Springer (2011).

[3] Jannach, D., Zanker, M., Felfernig, A., and Friedrich, G., "Recommender Systems: An Introduction", Cambridge University Press (2010).

[4] Crossen, A., Budzik, J., and Hammond, K. J., "Flytrap: intelligent group music recommendation", pages 184–185, ACM, (2002).

[5] Cosley, D., Konstan, J. A., and Riedl, J., "Polylens: A recommender system for groups of users", In Proceedings of the European Conference on Computer-Supported Cooperative Work, pages 199–218. Kluwer Academic, (2001).

[6] McCarthy, K., Salam, M., Coyle, L., McGinty, L., Smyth, B., and Nixon, P., "Cats: A synchronous approach to collaborative group recommendation", Proceedings of the Nineteenth International Florida Artificial Intelligence Research Society Conference Melbourne Beach FL, pages 86–91, (2006).

[7] Jameson, A., Baldes, S., and Kleinbauer, T. "Two methods for enhancing mutual awareness in a group recommender system", In Proceedings of the International Working Conference on Advanced Visual Interfaces (2004).

[8] Guy, I., Ronen, I., Wilcox, E., Carmel, M., Road, H., and Jose, S., "Do you know? Recommending people to invite into your social network". Proceedings of the 13th international conference on intelligent user interfaces, 236(3):77–86, (2009).

[9] Chen, Y., "Interface and interaction design for group and social recommender systems", In Proceedings of the fifth ACM conference on Recommender systems, RecSys '11, pages 363–366, New York, NY, USA, (2011).

[10] TouchGraph LLC, Navigator2, (2012). <http://www.touchgraph.com.>