

When being Weak is Brave: Privacy Issues in Recommender Systems

Naren Ramakrishnan, Benjamin J. Keller, and Batul J. Mirza
Department of Computer Science
Virginia Tech, VA 24061
Email: {naren,keller,bmirza}@cs.vt.edu

Ananth Y. Grama
Department of Computer Sciences
Purdue University, IN 47907
Email: ayg@cs.purdue.edu

George Karypis
Department of Computer Science
University of Minnesota, MN 55455
Email: karypis@cs.umn.edu

May 18, 2001

Abstract

We explore the conflict between personalization and privacy that arises from the existence of weak ties. A weak tie is an unexpected connection that provides serendipitous recommendations. However, information about weak ties could be used in conjunction with other sources of data to uncover identities and reveal other personal information. In this article, we use a graph-theoretic model to study the benefit and risk from weak ties.

1 Introduction

Privacy in Internet services is typically thought of in terms of protecting attributes of users (and can thus be related to solutions in database security). However, information provided by a recommender system can also allow the privacy of some of its users to be compromised, when used in conjunction with other information. For example, consider a system that recommends books by finding correlations between a user's ratings and those of other participants for the same books (a nearest neighbor algorithm [3]). Suppose as a user, person X rates only books on computer networking, but has an interest in Indian classical music. We could imagine the following dialogue between person X and the recommender system:

person X: Would I like "Evolution of Indian Classical Music"?

recommender: Yes.

person X: Why? (surprised)

recommender: People who liked the books you have liked also liked this book.

This is an example of *serendipity* in recommendation: person X does not expect to receive a recommendation for a book on a topic outside of those that he/she has rated. But, in fact, it is not by luck that the recommendation was provided — it means that at least one person has rated the book in addition to the books that person X has rated (and possibly others). The explanation conveys this, but also indicates that the algorithm used is a nearest neighbor algorithm. The serendipity, if person X is malicious, is that he/she has found a candidate *weak tie*.

In social network theory, a tie is a relationship between people. The strength of a tie is measured in terms of the number of shared associations. A family is an example: a child knows her mother and her father, and her parents know each other. Weak ties, on the other hand, form bridges between two groups of people who would not otherwise interact. Most importantly for us, weak ties allow us to deduce identities. For instance, one may know that a particular computer science department has a networking professor who was an interest in Indian classical music. This professor thus provides a tie between university faculty and the Indian classical music aficionados. Consequently, upon meeting the *only* Indian networking professor in that department, you can safely identify this person as the Indian classical music enthusiast.

A similar situation occurs in recommendation where weak ties provide the opportunity for serendipitous recommendations. In our example, a candidate weak tie has been found between networking books and books on Indian classical music. We don't know if this truly is a weak tie, but can test it by varying what we rate (the process would involve masquerading as different users and rating different sets of books to probe around the original ratings). Through this process, if the query fails for most variations on the ratings, then we can be more confident that we have a weak tie. And, in fact, the more restrictive the set of ratings that yield the positive recommendation, the more confidence we have that there is one person who has rated the books.

Our ability to find such a minimum set of ratings is where the risk lies — we can use this rating set to determine what the other person has rated using queries, and perhaps fit this information to our knowledge of people who use the system. In the end, it is conceivable that we could identify our hypothetical Indian music enthusiast/networking professor in the recommendation system, and determine what he may have rated. If the system allows us to vary the ratings, then we might be able to estimate the person's ratings as well.

In this example, a person actually forms a tie between books through their ratings: a book relates to another if they have both been rated by someone. The alternate view that we will take is that there is a tie between two people if they have rated some number of books in common. So, in the example, the Indian music enthusiast/networking professor has ties to people who rated networking books and people who rate books on Indian classical music. The ties to people who rate Indian classical music books are weak only if there are relatively few people who have the same tastes.

Clearly, the task of identifying someone through their ratings is difficult, and is even harder if more people have similar rating patterns. But the above example illustrates the risk that may exist even in simple recommendation. The risk is really to people who would participate in weak ties, because they are the people who could most easily be identified. In some application domains (voting preferences and membership on boards), even knowing that a weak tie exists constitutes a breach of privacy.

Our Approach

Our goal is to model the benefit from and risk to users who participate in weak ties. In particular, we would like to characterize benefits and risks on an algorithm-independent basis. To achieve this, we use the model of ‘jumping connections’ [6] that casts recommendation as making a series of jumps between people, based on common ratings (in nearest neighbor recommendation, there is only one jump). We describe how this model has relevance to conventionally accepted metrics of evaluation (Section 2). Using this model we can then identify causes for weak ties in terms of the rating patterns of people (Section 3). Furthermore, the model allows us to qualify benefits in terms of reachability in a graph, and risk in terms of weak ties (Section 4). Finally (Section 5), we look at how policies and social structures can be designed that can support and enable recommender systems.

2 Recommendation: The Jumping Connections Model

Recommendation algorithms work in a wide variety of ways, from forms of graph search to learning. This variety presents a difficulty when attempting to study the risks of recommendation in general. However, we can think of recommendation as making connections between people who have rated artifacts in common. We represent people, the artifacts they rated (e.g., movies), and their ratings as a bipartite graph (Fig. 1 (a)). An algorithm can then *jump* over the common artifacts to form a connection between two people. Fig. 1 illustrates a skip jump where two people are brought together if they rate at least one movie in common.

A jump induces a *social network graph* (Fig. 1 (b)), which includes only people and edges between them (in social network theory, such a graph shows direct relationships, but here two connected people need not know one another and their connection depends on the jump). The *recommender graph* (Fig. 1 (c)) orients the edges in the social network graph and adds back the movies. An algorithm can then find paths from a person making a query to a person who has rated the movie of interest. Note that Fig. 1 illustrates only one way of jumping — other jumps are identified by Mirza [6].

In this article, we restrict our attention to *hammock* jumps. A hammock jump of width w connects two people if they have rated at least w movies in common (a skip is a hammock jump of width one). A hammock path of length l is a sequence of l hammocks, as illustrated in Fig. 2. Our hypothesis is that hammock jumps underlie most recommendation approaches, and at the very least can be used as the basis to design metrics for studying privacy issues. Note that nearest neighbor algorithms (e.g., GroupLens [4], LikeMinds, and Firefly) use an implicit hammock sequence of length 1. The ‘horting’ algorithm of Aggarwal et al. [1] uses sequences of explicit hammock-like jumps.

Our model completely ignores accuracy of predicted values of ratings, and instead focuses on the parameters of hammock width and path length. This is because if recommendation is truly a matter of making (the right) connections, then a recommendation of a particular movie for a given person can be characterized by values for the hammock width w and the hammock path length l . Notice that we do not emphasize how individual ratings for the nodes (movies) spanning a hammock are transformed into a prediction. In addition, it is highly likely that there are multiple paths between the same set of nodes, with various constraints on w and l . Intuitively, since considering more common ratings can be beneficial (see [3] for approaches) having a wider hammock could be better (this has to be carefully done when correlations between ratings are considered [3]). But if we insist on a wide hammock, we might have to traverse longer paths to reach a particular movie from a given person [6]. However, recommendations

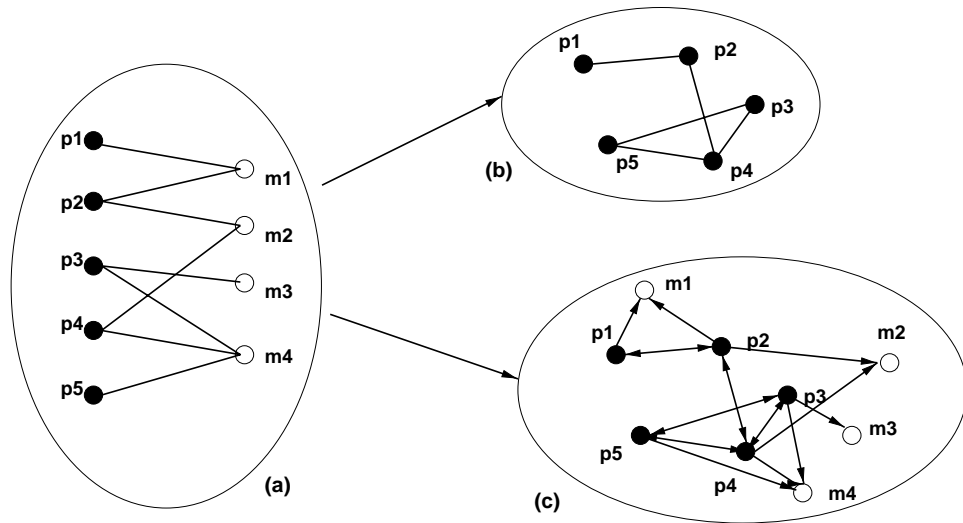


Figure 1: Illustration of the *skip jump*. (a) bipartite graph of people and movies. (b) Social network graph, and (c) recommender graph.




Figure 2: A path of hammock jumps, with a hammock width $w = 4$.




Figure 3: (left) Influence of hammock width w on quality of recommendation. The annotations denote the fraction of people and movies reachable for different values of the hammock width. (right) Influence of hammock path length l on quality of recommendations. The annotations denote the number of recommendations possible for each value of l .

involving shorter path lengths are preferred, for reasons of explainability, over longer paths. From a graph-theoretic point of view, w and l thus qualify the reachability of different movies from a given person, and indirectly provide a measure of the expected quality of predictions.

Preliminary analysis of the relationship between w , l , and predictive accuracy supports this intuition. Fig. 3 (left) shows a plot of the average discrepancy between predicted and actual ratings for each hammock width when using the LikeMinds algorithm (as described in [1]). These results were determined by a leave-one-out study, where an available rating was masked, and a prediction was made for that rating (using the remaining data). The number of common ratings between the given person and the person with the highest agreement scalar (and who contributed to the recommendation) was used as the hammock width. The results indicate that (for the LikeMinds algorithm), wider hammocks contribute to better ratings. Notice that LikeMinds’s hammocks do not just model commonality, they also posit agreement between the rating values spanning a hammock. While it is certainly true that we can get a poor quality recommendation even with a wide hammock (involving perhaps noisy ratings or a faulty aggregation procedure), overall quality of predictions is influenced by greater hammock widths. However, increasing the hammock width results in a progressive disconnection of the social network graph into many components. As a result, fewer and fewer connections can be made — Fig. 3 (left) also lists the fraction of people and movies reachable for various levels of hammock width. A w of 53 for instance reaches only 48% of the people and 93% of the movies. By the time the abrupt improvement in agreement values is observed (after $w \approx 110$), less than 25% of the people and only about 86% of the movies are reachable.

Fig. 3 (right) describes the results of an experiment where a minimum hammock width constraint was set at $w = 113$ (according to the LikeMinds definition) and the resulting recommender graph was analyzed for paths of varying lengths from people to movies. We used the transformation technique described in [1] to make predictions of ratings from others’ ratings, once again using the leave-one-out method. Paths in the recommender graph involve 1, 2, or 3 hops to the person providing a recommendation and a final hop to the movie being recommended (hence the bucketing of values into 2, 3, and 4 in Fig. 3, right). As can be seen, greater lengths (for the same w) cause a faster-than-linear decay in the quality of predictions. We should caution that horting [1] may exhibit different behavior, though we still expect longer paths to be of lower quality.

These results support the intuition that wider hammocks and shorter paths provide better ratings. Hammock widths are determined by rating patterns that ensure significant overlap. We see this in the MovieLens dataset for which each participant rates a minimum of 20 movies and which has a connected social network graph for all $w \leq 17$ [6].

The primary cause of shorter paths is having more connections in the graph. In the MovieLens dataset, a recommendation is almost always possible using a path of length no longer than 3. This is due to the power-law degree distribution of the rating patterns. Other graphs, such as *small-world networks* [8], have small clusters of vertices that are connected by relatively few edges. ‘Weak ties’ are important in both these situations because they make some recommendations possible, and provide others with shorter paths. Therefore, weak ties are very important to recommendation.

3 Ties: Strong, Weak, and Brave

In contrast to weak ties¹, a strong tie connects two people who share many associations (like in a family or some other close-knit group). We can think of weak ties as forming bridges between groups of people who would otherwise not interact. Of course, strength and weakness are relative, and there is no agreed definition of what a strong tie is in terms of the number of shared associations.

¹It is important to note that there is nothing fundamentally feeble or fragile about a weak tie; a weak tie creates a powerful and robust link between nodes from different neighborhoods.




Figure 4: Strong ties in a social network graph (right) induced by a hammock jump on a recommendation dataset (left) with $w = 2$.

In a graph, strong ties are characterized by a triangle of vertices (a *triad* in the social network literature). Fig. 4 illustrates how these triads can occur in a social network graph induced by a hammock jump. In this case, the width of the hammock jump is 2, and what looks like two relationships becomes three in the social network graph. Notice that, in this example, if the hammock jump width were three, then the resulting social network would not have a triad and so neither edge would represent a strong tie. It is a classical argument in social network theory that no strong tie can be a bridge and that two strong ties would imply a third tie [2].

Weak ties are of most interest to us, because they are the foundation for our notion of risk. As discussed earlier, a weak tie in a social setting allows people to identify someone with other information that they've been given. Weak ties occur simply because someone knows someone else outside of their usual circle of friends; or perhaps there is a person (an 'outsider') who is friends with a few people who each have strong(er) ties to people in different groups.

In recommendation, weak ties originate from the rating patterns of the participants, but the jump process also plays a crucial role. We hypothesize two fundamental rating patterns. One can be observed in the public movie recommendation datasets (MovieLens and EachMovie), and the other is what we would assume for a domain where people have stronger bias in their tastes (such as books or music).

The movie datasets exhibit a power-law degree distribution as illustrated in Fig. 5 (top, left). The power-law rating pattern comes from preferential attachment; for example, some movies (the *hits*) are rated by almost everyone, and some people (the *buffs*) rate almost all movies. Weak ties are rare in this setting but might occur when a person shows no strong allegiance to any genre and rates relatively few movies in each (he/she would not be a buff). The real risk to these people is that they might not have enough ratings in common with anyone so they can be given recommendations.

The second rating pattern would occur where most people exhibit a preference for a particular kind of artifact. This is illustrated in Fig. 5 (bottom, left) where there are three subgraphs with power-law structures, connected by a relatively small number of ratings. This diagram illustrates one source of weak tie in this setting, which is when someone who ordinarily only rates artifacts in one domain (e.g., networking books) rates an artifact in another domain (e.g., Indian classical music books). Another possibility is someone with more eclectic tastes who rates artifacts across many domains, and unlike in the power-law graph is truly a weak tie. The risk with weak ties in this rating pattern is that they may allow us to identify a person whose ratings can get us from one domain to another.

The jump process, described in Section 2, can also create weak ties when using common ratings as the basis for making connections between people. Many people might have rated across several domains, but only a few have enough ratings to satisfy the jump being used. A final reason relates to merging of data collected from different settings. For instance, the recent purchase of eToys consumer data by another retail giant signals the possibility of the creation of a social network graph with weak ties.

The risk in a weak tie really comes from being the only person with a peculiar rating pattern — there is safety in numbers, or at least in homogeneous tastes (as in power-law graphs). The more people who rate the same kinds of things, the less likely that any one of them will be identifiable as participating in a weak tie. But notice that if the jump definition weeds some of those people out, the risk is still there (although it is less likely that any additional

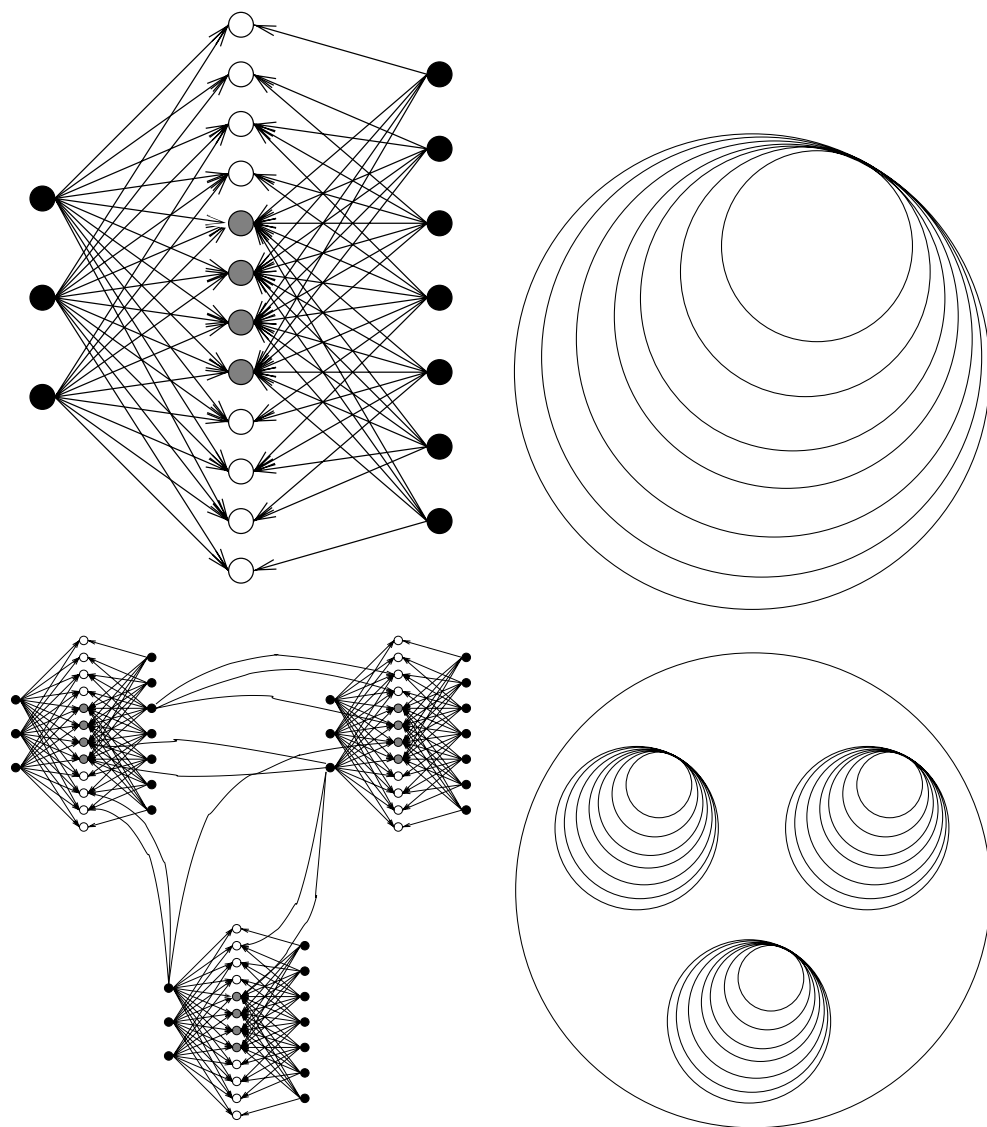


Figure 5: Two different types of induced social networks that can exhibit weak ties. (top left) A dataset with a power-law induces a low-risk social network (top right) where increasing hammock widths cause a ‘nested clam shells’ picture. Each circle in the social network picture denotes a group of people brought together. Increasing hammock widths cause the circles to get progressively smaller. (bottom left) A dataset with power-laws in only subgraphs and a few weak ties induces a high-risk social network (bottom right) characterized by the breakdown of a connected network into disconnected networks. Some experimental data supporting the diagrams above can be found in [6].

Table 1: Movies used in analyzing the benefits of ratings on personalization. ‘Star Wars’ and ‘Scream of Stone’ had the highest and lowest number of ratings, respectively.

<i>Movie Name</i>	<i>Number of Ratings</i>
Star Wars	583
Tomorrow Never Dies	180
Robin Hood: Men in Tights	56
Scream of Stone	1

information could be used to identify a single person).

4 The Benefits and Perils of Personalization

Intuitively, a user desires the most benefit from a recommendation that is based on wider hammocks and shorter path lengths. Of course, to get these qualities we have to provide more ratings, and the risk is that we might introduce a weak tie. The problem then is can we relate how much we rate to the benefit and risk inherent in recommendation?

When there are multiple recommendation paths between a given combination of person and movie, we would like a benefit formula that captures our preference for wider hammocks and shorter path lengths. By defining the benefit of a recommendation as:

$$\text{benefit} = \frac{w}{l^2}$$

we can give more weight to improvements in path length from 2 to 1 than, say, from 3 to 2. This non-linear dependence of quality of interaction on the length is supported by research in diffusion processes [8], social networks [2] and also our own experiments (see Fig. 3, right).

We can explore benefit in terms of the number of artifacts that are rated. Typically, recommender systems require that users rate a minimum number of artifacts before they can make queries, and so we look at the incremental benefit received by providing additional ratings. For this purpose, we analyze the MovieLens dataset where it is required that a user rate 20 movies, and add a new person. The MovieLens dataset consists of 943 people, 1682 movies, and is connected as a graph.

For the experiment, we introduce a 944th person and incrementally add ratings from the new person to movies (so that movies with a higher rating frequency were more likely to be rated). After each rating was added the path lengths l to particular movies (see Table 1) were computed for each hammock width w . Twenty repetitions were performed for each additional rating.

The benefit from additional ratings for the movies in Table 1 is shown in Fig. 6. Each colored cell indicates that the particular benefit is possible for the corresponding number of ratings. The feasible benefit regions are actually monotonically increasing by popularity of the movie — with more possibilities for ‘Star Wars’ than for ‘Tomorrow Never Dies.’ The plot shows that if you want a good recommendation for a less popular movie, you need to provide more ratings, but can receive good recommendations for popular movies with fewer ratings. In particular, requesting an improvement in benefit for a ‘Star Wars’ recommendation from 5 to 14 requires no extra ratings!

The danger involved in recommendation relates to the probability that a weak connection is exposed; unfortunately, this is not a static property of a recommendation path and can only be studied in reference to the social network graph *in the absence* of the considered connection. This means that we need a more complete understanding of the dynamics by which weak ties are introduced, modeled, and employed in a social network. Such an understanding could take the form of a graph-generation model. Here we use the model of Watts and Strogatz [8] as a basis for our study of risk.

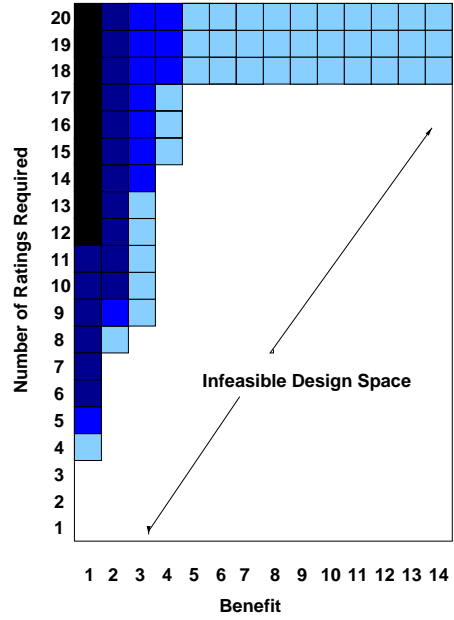


Figure 6: Benefit vs. number of additional ratings required, for various choices of movie destination nodes. The cells are colored with greater intensities corresponding to movies with fewer ratings.

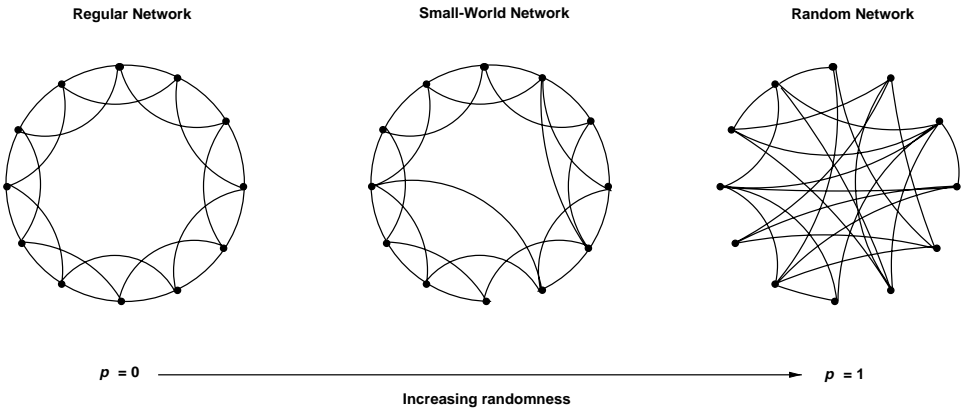


Figure 7: Random rewiring, starting from a regular wreath network, introduces weak ties that help model small-world graphs. Figure adapted from [8].




Figure 8: (left) Average path length and clustering coefficient versus the rewiring probability p (from [8]). All measurements are scaled w.r.t. the values at $p = 0$. (right) Quantifying the risk as a function of rewiring probability p .

The intuition is that risk occurs when we have edges that are weak ties between subgraphs that are cliques (or at least nearly so), and the risk decreases as more of these edges are added. In particular, the risk is highest when a new weak tie occurs and the lengths between people decrease dramatically. As more weak ties are added, the risk decreases.

This idea of risk can be explored in the Watts-Strogatz model for small-world networks. They show how to generate a spectrum of graphs from a regular wreath graph by adjusting the probability p of rewiring an edge (Fig. 7). When p is zero, we have the wreath; but when p is one, we have a random graph. The risk from weak ties is low in both the wreath and random graphs, but increases as the average path length drops but the vertices are still clustered. Fig. 8 (left) illustrates the relationship between length and clustering (see [8] for details of the definitions). When p is between 0 and 0.1, the graph is a small-world network, and poses the most risk from weak ties.

We can express the risk of weak ties in terms of p : the risk in having ratings that form weak ties can be quantified as the rate at which l reduces, as a function of p :

$$\text{risk} = -\frac{\partial l}{\partial p}$$

The risk for the dynamics described in Fig. 7 is given in Fig. 8, right (the length values are scaled with respect to the length at $p = 0$ before calculating the risk). Notice that the risk increases rapidly (as weak ties are introduced) and drops down gradually (as more weak ties share responsibility for length reduction). This captures our intuition pertaining to disclosure of sensitive information by ferreting out weak ties. However, our jumping connections model is not directly parameterized by p .

To be useful, the above formula for risk must relate length reduction to a metric that could be used to balance personalization and privacy. We can illustrate the risk of becoming a weak tie by studying what happens as we decrease the hammock width w in Fig. 5 (bottom). Consider the situation when the social network graph is in three disconnected components. Decreasing the hammock width would introduce new edges that are weak ties, which would contribute to length reduction and thus, quantification of risk. However, recall that increasing the hammock width is desirable from the viewpoint of benefit. Taken another way, benefits improve monotonically with increasing width w but risk rises rapidly (as fewer weak ties share responsibility for length reduction) upto a point and then drops sharply.

To explore this setting, we created an artificial dataset that consists of three subgraphs with power-law degree distributions, each with 200 people and 75 artifact vertices. Each person node is linked to at most 15 artifact nodes




Figure 9: Risk as a function of hammock width w .

within the same subgraph. Specifically, the people and artifacts were ordered, and the b^{th} person rated the first $\lceil 75b^{-\epsilon} \rceil$ artifacts. The value of ϵ was calibrated to achieve a minimum rating of 15 artifacts. Then three extra people were added who rate (at most 15) artifacts in all three connected components, again with a ‘master’ power-law.

For a hammock width of 9, the social network of this graph consists of three disconnected components. By decreasing the hammock width, weak ties will be introduced into the social network, and the path lengths decrease. The results are plotted in Fig. 9 (lengths are scaled against the path length for $w = 8$). As could be expected, risk is highest when the graph is first connected.

It is not possible to provide a traditional benefit-risk profile, as is customary in analysis. This is because recommender systems aggregate the ratings of many participants when computing a recommendation. A user’s benefits comes from ‘plugging into’ the social network by providing a sufficient number of ratings, but a user’s risk depends not only on what is rated but also on what other people rate. Ultimately, the difficulty comes from the fact that risk occurs even if recommendation queries are not made, but benefit requires that the user make queries.

The two qualitative conclusions from our studies are that (i) a few weak ties are more risky than a lot of weak ties, and (ii) more so, in some (induced) social networks than others.

5 Concluding Remarks

The very factors that make weak ties useful are the ones that raise the threat of privacy. We have demonstrated that under certain conditions, recommendations could involve weak ties and could potentially compromise the privacy of individuals. Like most problems in computer security, the ideal deterrents are better awareness of the issues and more openness in how recommender systems operate in the market place. In particular, policies and methodologies employed by an individual site should be made clear. Sites that involve multiple homogeneous networks have a crucial responsibility in clarifying the role of weak ties in their system designs and what forms of mechanisms are in place to thwart hackers.

Ideally, recommender systems should convey to the user both benefits and risks in an intuitive manner. One possibility is to present the user with plots of benefit and risk versus user-modifiable parameters — ratings, w , and l (if the algorithm allows their direct specification). Another possibility is to qualify the risks and benefits associated with rating each individual artifact (as a function of the previous ratings in the system). Providing a rating for ‘Scream of Stone’ for instance would provide dramatic improvements in benefit than providing a rating for ‘Star Wars.’ At the same time, the system should qualify the extent to which a user becomes a weak tie, by such a rating.

Singh and colleagues [7] make a provoking observation in drawing comparisons from community-based networks to recommender systems — namely, that people really want to control to whom they reveal their ratings but would like to know how recommendations are being made. In a distributed setting, one can imagine a scenario where people specify how data collected from their interactions should be modeled and used. Interfaces for privacy management are woefully inadequate and their role is only now being recognized [5]. Extending the results here to a distributed setting where people can set arbitrary constraints on their station in the social network graph (whether they are willing to participate in a path?; are there constraints on such participation?; would they provide ratings if they knew that it would contribute to a weak tie?) is a possible direction for future research.

One wonders if weak ties will happen at all, if concerns are raised about their compromise. Social network theory postulates that they are the primary mechanisms by which micro-level interactions can manifest at macro levels, and that such ties will be kindled whenever communities have to be mobilized for collective action. It remains to be seen if weak ties induced by jumps in a recommender system also conform to similar distributed organization.

References

- [1] C.C. Aggarwal, J.L. Wolf, K.-L. Wu, and P.S. Yu. Horting Hatches an Egg: A New Graph-Theoretic Approach to Collaborative Filtering. In *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 201–212. ACM Press, 1999.
- [2] M.S. Granovetter. The Strength of Weak Ties: A Network Theory Revisited. *Sociological Theory*, Vol. 1:pp. 203–233, 1983.
- [3] J.L. Herlocker, J.A. Konstan, A. Borchers, and J. Riedl. An Algorithmic Framework for Performing Collaborative Filtering. In *Proceedings of the Twenty Second Annual International ACM SIGIR Conference*, pages 230–237. ACM Press, 1999.
- [4] J.A. Konstan, B. Miller, D. Maltz, J.L. Herlocker, L. Gordan, and J. Riedl. GroupLens: Applying Collaborative Filtering to Usenet News. *Communications of the ACM*, Vol. 40(3):pp. 77–87, March 1997.
- [5] T. Lau, O. Etzioni, and D.S. Weld. Privacy Interfaces for Information Management. *Communications of the ACM*, Vol. 42(10):pp. 88–94, October 1999.
- [6] B.J. Mirza. Jumping Connections: A Graph-Theoretic Model for Recommender Systems. Master’s thesis, Department of Computer Science, Virginia Tech, February 2001. 64 pages.
- [7] M.P. Singh, B. Yu, and M. Venkataraman. Community-Based Service Location. *Communications of the ACM*, Vol. 44(4):pp. 49–54, April 2001.
- [8] D.J. Watts and S. Strogatz. Collective Dynamics of ‘Small-World’ Networks. *Nature*, Vol. 393(6):pp. 440–442, June 1998.