A General Model for Out-of-town Region Recommendation (Pham et al., 2017)

1. The author seems only output the only one region, do you think is it realistic? Also usually users decide the region first and decide the POI, do you think the recommend system is practical or attractive to users?
2. Sometimes when people go out-of town travel, people usually are not alone. Do you think if we can add the variable to determine the group to travel, is it better? I think if people travel with their partners is totally different from traveling with their friends or family.
3. If in the real recommended system, users would like to set some filters. For example, users like to go the beach or not hot regions. In this case, is it possible to return more candidates regions and filter them or we have to adjust the model and retrain them?
4. The author used the conditional probability to set the inference between two POIs. Do you think the parameters are too few, because it only depends on user not time not the people users hang out with. And also the conditional probability is calculated by the number of users, it seems not personalize, right?

   (Kuang Hsuan Lee)

1. In addition to just predicting the number of POIs that are located within a region, it seems that the quality and variety of the POIs also play a big role in whether someone will visit a new area. How can these additional sources of information be taken into account?
2. Is it possible to run a live or A/B test of an algorithm like this? It seems that just separating the observed users into training/test sets and evaluating might not be a realistic scenario in terms of helping users decide where to travel.
3. Are there additional characteristics of POIs and co-occurrences that could be taken into account? For example, it could be useful if some POIs should only be visited together, or a certain set of places should be visited in a sequence.

Rajiv Pasricha

1. How does this model account for the temporal aspects that needs to be modeled in this problem, say seasonal effects that can influence users in visiting certain places. What would happen if a POI is no more entertaining, even though it was popular once.
2. How is the satisfaction of an user for a particular region computed? For example, if an user tends to visit a place repeatedly, will that have an effect on the value?
3. Is this method biased towards the interaction between the POIs? What would be the performance of the model if a person needs to visit a POI which is not popular or among other POIs?

4. Just modelling the user satisfaction score is pretty vague, as each user might rate a POI differently even though they might like it. Can we leverage some other data, say the rext reviews?
   (By Vignesh Gokul)

1. Shouldn't they compare their model with Individual POI recommendations with a region dxd centered around the POI? This would had been a better evaluation of their model instead of their baselines.
2. People would generally visit a popular POI out-of-town. Would seeding the searching algorithm with most famous POIs be more efficient?
3. Does the training set also has users who have at least two POI covered by dxd region? Otherwise, the second term in Eq 5 would be zero. Fig. 2 shows #POIs in regions in mostly 1. So, the first term in Eq 5 would dominate?
4. They have not accounted for the categories of the POIs in a region, but this can be important when two POIs are complementary/alternatives or other case might be, one POI is a tourist place and other POIs are secondary POIs like hotels, restaurants etc. In the later case, more weight should be put on tourist place.

- **Kulshreshth Dhiman**

A Probabilistic Model for Using Social Networks in Personalized Item Recommendation (Chaney et al., 2015)

1. How does NCRR compare to other ranking-specific evaluation metrics such as AUC? How feasible is it to define a pairwise version of the algorithm similar to BPR?
2. What are the advantages of using a probabilistic approach to add social connections into a factorization model. Could it also be effective to add additional social features to standard matrix factorization models which are directly trained using gradient ascent?
3. Would a social network joined with a purchase or rating dataset also be effective? For example if we joined a user's Facebook connections with their Flixster or Etsy purchases.

(Rajiv Pasricha)

1. This model treats all friends equally while calculating influence. How can we incorporate strength of social link (trusted friend vs not) in the model? Weighting the social component of a friend by strength factor work?
2. Friends can have different influence for different items (or a group of items) like different genres of movies etc. How can this be accommodated in the model?
   -Kulshreshth Dhiman

Online popularity and topical interests through the lens of instagram (Ferrara et al., 2014)

1. The authors said “users who already uploaded a large number of media are more likely to do so,”. Do you think if the author used the x is the number of posts last year and y is the number of this year, is it better to measure the dynamics?
2. Do you think if we split the data to different group and run the tests again, is it more interesting?
3. The author said “the unpopular users tend to be more focused in their interests with respect to more popular users”. Do you think whether is it possible that the users posting more various topics easily attract more other users attention? **(Kuang Hsuan Lee)**

1. The observations were made on data collected over a very short period of time(1 month) this doesn’t capture any dynamic behavior. For example some tags like #christmas are popular only during particular times of the year. Is it better to collect information spanning longer duration and will it change these observations significantly ? **(Sai Chaitanya Kolasani)**
2. The collected media and tags are from photographic contest. Don’t this fact biases the results ? For eg: Users already winning these contests generally posts more as there is an incentive but it may not be general human behaviour on this network.**(Sai Chaitanya Kolasani)**
3. The results show users with very broad interests and specialized users are equally likely to be popular. What might be the explanation behind this ? Is it just because users with broad interests tend to post more and thereby may become popular ? **(Sai Chaitanya Kolasani)**

(Sai Chaitanya Kolasani)

1. Was there any way in which they could remove the bias in the data or crawled data in a different way? The chose to crawl the posts of an Instagram contest as their original clustering coefficient was low..but maybe taking all these users did not give correct trends since media and hashtags other than #whp might also not tell much about users themselves and were instead specific to the contest.

2. They explain that individual user vocabulary size is limited and majority of them use a few tags-Is this a correct finding since they used data of only a specific timeframe? I feel like the user’s vocabulary increases over time since they follow on more trends of using ‘tags’ and they should have studied longer temporal trends .

3. So according to them the very popular users have either very specific topical interests or very broad ones. Is this a valid conclusion taking into account only these many users.? There can be very popular users who would like only a few topics. Or is popularity correlated to how many and how different topics a user is interested in?

-Akanksha Grover

1. **The authors seem to want to catch the dynamics in social tagging and user popularity with only one-month data, do you think that's efficient? Because different time in a year, people tend to focus on different things, like during the election or summer holiday time.**
2. **Do you think "unpopular users tend to be more focused in their interests with respect to more popular user" and "larger sets of tags assigned to the same media are increasingly more unlikely to be observed" contract each other?**
3. **Do you think that the finding that "a large majority of users adopting only few tags" is because of the small time difference in the data they used?**

(Yiwen Gong)

Recommendations in Signed Social Networks (Tang et al., 2016)

1. **How do the authors decide the Wij and whether we can set the parameters on f(x) because maybe the rank is not follow this function? How do they decide this f(x)?**
2. **The author used Uip, Uin and max it to get the result, what is the difference between Pearson correlation between two user latent vectors?**
3. **Do you think some recently good friend make more impacts on the user? If you agree with that, how to modify this model?**
4. **Do you think we should include the social relation strength on the second terms?**

   (Kuang Hsuan Lee)

1. **In the paper, there is a statement that RecSNN provides a unified recommendation framework with unsigned and signed social network. Do you think that the previous SocialMF model is a sort of special case of this model?**
2. **Could the model be modified to consider the weighted relations?**

3. **How can we extend this model with some other user attributes? For example, the model just regard foes as the same but if we can get more attributes could we have a better way to define circles?**
   (Wen Liang)

Additive Co-Clustering with Social Influence for Recommendation (Du et al., 2017)

1. **It seems the author does not consider the negative connection. Do you think this model can be used in negative connection?(foe)**
2. **Do you think if the graph is the two rings where k-means usually does not work, this model still works well?**
3. **Do you think if there are complementary friends who have different interests, whether we can split the social group to two parts and modify the mode?**

(Kuang Hsuan Lee)

1. The authors state that for the purpose of the paper, they have treated all graphs as undirected. What change do you think would be seen if the graphs were to be directed?
2. The improvement in performance over ACCAMS is not very pronounced. Do you think this could be because of the sizes of the datasets used?

   Aditi Ashutosh Mavalankar

1. How can graphical structural information be considered between items in this SIACC setup? This might be useful to alleviate cold start for items
2. Social links are just one kind of contextual information that can be used to alleviate cold start issues. What other contextual information can be expressed in terms of inter-item and inter-user links to use a framework such as SIACC?
3. Can an evaluation metric that is more suitable to reflect the fact that recommendations are tailored to social relations emphasize the improvements that SIACC presents? How can such an evaluation metric be constructed?
   Siddharth Dinesh

1. The authors mention that they use only direct relationships as of now. I was just wondering how indirect relations and other graph parameters like degrees/centrality could be exploited here?

2. What is the main reason why using local and global influence in additive co-clustering gives better results as compared to techniques that use both in matrix factorization model?
3. How does the fuzzy factor actually help in deciding which cluster the user belongs to as it is mentioned that a user can belong to multiple clusters at the same time.

- Akanksha Grover

Growing Wikipedia Across Languages via Recommendation (Wulczyn et al., 2016)

1. The authors discuss the fact that they took into account that the emails themselves may have influenced content creation by sending emails with random recommendations as well. Do you think that the results could have been biased by the fact that the contracted editors knew that they were part of a study?
2. To aggregate edit histories into interest vectors for editor interest modeling, the authors take into account the number of bytes that were added to the article. Would this not skew computations to favour adding images over adding text?
3. What are the possible approaches to extending this method to induce new editors to translate articles, ie. what are the possible ways in which the cold start problem be solved in this setting?
   (Shreyas Udupa)

1. While ranking missing articles, authors consider number of links as a feature. Can we use PageRank algorithm kind of approach to determine the importance of an article, as different links should be given different importance?

2. Wikipedia editors often have a user page describing themselves and their interests. Is it possible to use that information to improve Editor interest modelling?

3. The authors reduce the size of interest vectors by considering only a number of the most recent articles edited by an editor. How will our results get impacted, if we rather reduce the size of interest vectors by only considering the top articles edited by an editor (maybe by considering the amount of editing done) because a lot of addition to an article might be more important than a few words changed in an article.

By Nitin Kalra

1. The authors mention that, there is a need to rank the missing documents because of varying cultural contexts and other reasons. So would the number of page views and wikidata count help to model this relationship?
2. In modelling just the topic interests of the editors, is the model failing or ignoring to capture the language interests a particular editor might have?
3. How would the model capture the change in interests of an editor, say the topic or the language over time?(By Vignesh Gokul)

1. The experiments are completed on Source language English and Target language French. These two languages have many similarities in terms of Vocabulary and Grammar. And both languages are mainly used by people who share the similar culture, i.e. "western culture". Will these facts create a bias for the results? In other words, will the method also work well from Source language English to a Target language that has a larger difference such as Chinese or Japanese?
2. The method considers recommending missing pages for editors. In Wikipedia, there exists large differences in quality of articles of different languages. An article well written in English might not have the same quality in French. Could it be the future work to recommend low quality pages for editors.

- Siyu Jiang

1. The authors first say that first say that articles in one language may not be as important in another, but then they go on to use the rank in the source as the ground truth for learning ranking criterion ? Isn't that a bit contradictory ?
2. Instead of trying to find people proficient in both source and target languages, wouldn't it be better to first convert the article to the target language and then ask the target editors to make suggestions/corrections/improvements to that article ?
3. Wouldn't it be more useful to also include the most similar articles in the editor's history while calculating features ? Just because they are working with a different topic recently doesn't mean that they do not want to work with topic related to earlier edits ? Wouldn't they be more knowledgeable and qualified for the expansion process ?

- Ajitesh Gupta

The Effect of Recommendations on Network Structure (Su et al., 2016)

1. Would we expect to have the same findings if the network had weighted edges or symmetric relationships / bidirectional edges?

2. Are the authors justified in using random walks of length two to simulate friends of friends algorithm? Wouldn't a user be influenced by which of his friends are following the recommended user?
3. How could the findings of this study be used to improve design of friend suggestions / recommendations? Is it better to use a content-based system than a graph-based system? Should the ranking be normalized somehow to neutralize popularity measures? Or is the fact that in such a network, the 'rich get richer but a rising tide lifts all boats' not such a bad thing after all?
(Shreyas Udupa)

1. Twitter uses a social network model where there is a possibility of following celebrities and more popular users in a unidirectional manner. It should be convincing that in such a model more popular users would be followed by less popular users. Do you think this might show similar results in case of a Facebook type of model, where celebrities have pages, and a connection is complete only when both users accept the connection?
2. It is claimed through an analysis where they divided users in buckets based on degree of connections in the network, that more connected users saw greater increase ind the number of connections after the recommendations were shown. Shouldn't the study be more conclusive if they had analyzed the number of connections added where the user was less popular than himself vs when the user was more popular?
3. In the daily number of edges added plot, we see the daily edges portraying a saw tooth like curve. Can we discuss why this kindof shape might be observed irrespective of the fact the Who To Follow was introduced or not.

(Dhruv Sharma)

A Location-Sentiment-Aware Recommender System for Both Home-Town and Out-of-Town Mobile Users (Wang et al., 2017)

1. The authors fix some hyper-parameters (e.g. alpha, beta) for simplicity in the model. How do the authors choose these fixed hyper parameters? Is there any way to learn these parameters?
2. JIM has the 2nd best performance in the evaluation. How does JIM take "Home-town"/"Out-of-town" feature into consideration? I think if JIM takes this feature into 1 model instead of 2, JIM is more robust than LSARS model.
3. The dataset only contains the active users. Could this model also work on the cold start problem?

(Zeng Fan)

1. How does LSARS model the overall sentiment of a point of interest based on a small sample of reviews ? **(Sai Chaitanya Kolasani)**
2. The amount of user drift could be a function of the distance between the POI's can this model be improved by including the distance into the model rather than having a fixed threshold of 100kms ? **(Sai Chaitanya Kolasani)**
3. How can LSARS be extended to handle the exploration exploitation dilemma ? It is tries to recommend places only with positive sentiment. New places which the user may sometimes like are never ranked in the top k ? **(Sai Chaitanya Kolasani)**

- **Sai Chaitanya Kolasani**

Exploiting Socio-Economic Models for Lodging Recommendation in the Sharing Economy (Vazquez et al. 2017)

1. For Perceived Risk (PR) of an Airbnb listing, authors consider a few relevant numerical features from the dataset. Can we also incorporate additional features using Perceived Risk specific lexicon, as done in case of Perceived Authenticity to improve performance? When should such an approach (lexicon based) be used, in general?

2. The authors derive features from Airbnb dataset, for various aspects, to recommend the lodgings. Can we further improve the results by incorporating features derived from other sources. For example, crime index data for various neighbourhoods might be a good feature while considering perceived risk, as users often do some sort of quick search about an area/neighborhood on google or elsewhere before making a booking.

3. How will new listings be affected in such an approach, as many features depend on data not available for new listings. Would it be a good idea to give some boost to new listings, to provide them visibility, if such a recommendation system is to be actually used?

By Nitin Kalra

1. Electronic word-of-mouth methods may not be completely reliable, and may contain spam/deliberate positive/negative advertising. How do you think the current algorithm can handle that (if possible)?
2. Instead of using popularity (just the count of the bookings of a particular lodging), wouldn't a scaled popularity be more indicative of a hotel's reputation? Scaling could be done along various lines, like the bookings at hotels near that particular location, or hotels that offer the same price range, etc.

3. What other metrics, besides MRR, would you recommend for statistically analyzing the current recommendations?

   Aditi Ashutosh Mavalankar

1. It would be interesting to see how often Superhosts come up on a higher rank with the proposed model. It would also be interesting to see how the prediction might be affected if whether the host is a Superhost or not is added as a feature. One might think that characteristically this is covered by other features like ratings and reviews, however, Superhosts vary on all the other aspects apart from the PR which might make it an important addition. What do you think?
2. For a city like London, distance is perceived more by how long it would take on the tube to get to the venue and the proximity of the lodging to a tube station/what zone the lodging is in. Do you think these factors are minimized by merely normalizing the distance between the lodging and venue?
3. A user's price sensitivity may change over time and may also vary with their intent behind the trip. In this model, the PS features are primarily location and lodging focused. Do you think the model will be able to account for the user's current needs?

- *Sejal Shah*

Social Collaborative Viewpoint Regression with Explainable Recommendations (Ren et al. 2017)

1. The SCVR model seems to have a lot of moving parts and distributions. Is there a way to approximate the various parts and distributions in the probabilistic model so that inference can be done even when the dataset is smaller in size or for sparser datasets?
2. The authors assume that each user review would contain only one viewpoint. A viewpoint is defined as a combination of a concept, topic and a sentiment. Considering that there is previous work on aspect based sentiment analysis for reviews, is it a valid assumption that there would only be one viewpoint in each review?
3. Can other graphical relationships between users or items be modeled in a similar fashion such as how the authors have modeled the influence on trust networks on a user's ratings?

Siddharth Dinesh

1. We have seen various papers that model short reviews called 'Tips'. Is a viewpoint different from Tips?
2. The feature detection system finds the concept of the review by finding the similarity between a concept and each word in the review. Using word embeddings is pretty vague because there could be a review "I love American food, but I love this place" for concept

“Italian food”. Wouldn’t a sentence embedding scheme such as skipthoughts perform better? (By Vignesh Gokul)

Exploiting Socio-Economic Models for Lodging Recommendation in the Sharing Economy (Vazquez et al. 2017)

1. Users may have different needs for different trips, such as number of persons will vary, and the most important holiday Christmas and New Year is not included in the experiment. This paper did not seem to capture user preference during their searching, online recommendations based on user input might be more accurate.
2. The experiments were conducted on 2 metropolis, New York and London. How do you think the training results may affect users’ preference while they go to rural area?
3. For section 4.1, “ we replicate each test case five times, varying the candidate set I each time within the original radius of 2 km”, I am not sure if I fully understand this sentence. Does that mean the target location would be changed 5 times, and how much overlap would this provide?

(Yiwen Gong)

STAR: Semiring trust inference for trust-aware social recommenders

1. The author states that “The trust value domain is application dependent” but doesn’t argue against using discrete vs continuous values? How do we decide among these?
2. In almost all social networks, the maximum distance between any two nodes is at most 8 nodes - So, how important is the varying path length characteristic?
3. Comparisons made on just one dataset with discrete definition of trust (although author claims it works equally well on a continuous trust definition) not very indicative of performance.

- Balasubramaniam Srinivasan

1. What are some drawbacks and future directions for this work?

(Chester Holtz)

1. Would partial reciprocity in discrete setting imply full reciprocity? In that case would it be better to convert continuous data to discrete form?
2. The baselines considered in the paper seem to be old. Are there any recent work in your knowledge that could be taken into consideration for comparison?
3. Could this method somehow be adapted to model generation of trust/distrust links between two members over time? For example, if two people were not friends before there would no link between them, and when they became friends, we get a trust link between them.

- By Rishabh Misra

Fairness-Aware Group Recommendation with Pareto-Efficiency

1. Do you think the authors could use more appropriate datasets to evaluate their proposed framework? MovieLens does not have defined groups and MoviePilot usually has around 2 users as mentioned in the paper.
2. Do you think this framework would be applicable while recommending products to a segment of customers (back-to-school shoppers, high spenders, etc.) on an ecommerce website like eBay. Basically, can a segment of customers be considered as a group and will considering fairness lead to better recommendations?
3. Greedy algorithms are vulnerable to local optimization, do you think global optimization algorithms might be a better fit for solving the Pareto optimization problem?

- *Sejal Shah*

1 - Even though the proposed method beats all the baselines, precision and recall still seem to be very low? Do you have any insight into why it could be happening?
2 - Average Ranking Algorithm baseline, which ranks just on the basis of average relevance, seem to be performing equally well with the proposed method having marginal improvements. Doesn’t it mean that Fairness might not that be helpful as it seems in theory? And is it worth the extra complexity?
3 - Would it be a good idea to have fairness as a regularizer in our optimization objective that maximizes the social welfare? The problem could be much easier to solve then.

-By Rishabh Misra

1. Would a fairness metric based on vote share be a metric in terms of making group recommendations, say for example 5/10 people may hold

80% of the votes in the group (it is more relevant to them) - so would treating all users unhappiness/discontent equally be fair?

- Balasubramaniam Srinivasan

1. There are instances where in a group dynamics certain members (let's call them alpha members) tend to drive the preferences of the group.
    a. Two questions here: Should we be modeling it if the objective is fairness
    b. If yes, do you think modeling it as aggregation of users will work here or such cases are more rare for ephemeral groups?
2. While the author proposes an integer programming solution (which could be solved using SAT solvers) and also a way to model the integer programming solution into a continuous function where X need not be 0/1, it is still not considered for evaluation. Do you think it would have shown similar results or just the computational reasons can be sighted for this decision? In other words, do you think the continuous version is scalable and effective?
3. Both Social welfare and Fairness are calculated using just relevance. Do you think there can be any other metric to model the same considering relevance can change because of many reasons (temporal, last item etc)?

(Dhruv Sharma)

-