Comealong: Empowering Experience-Sharing through Social Networks

Beatrice Valeri, Marcos Baez, Fabio Casati
Department of Information Engineering and Computer Science
University of Trento, Italy
Via Sommarive, 5 I-38123 Trento
{valeri, baez, casati}@disi.unitn.it

Abstract—In this paper we present a social platform that empowers real-life social experiences by using the crowd to help people discover interesting events around a particular location and make the best of the actual experience. We demonstrate how current systems cover only partially the social and individual motivations on deciding and teaming up for events, and propose i) a system of implicit and explicit incentives that rewards intention and experience sharing within circles based on friendship or interests and ii) a discovery algorithm that combines the importance of coming along with friends with the personal interests to recommend events. The resulting platform has been applied to generic events as well as domain-specific ones such as shows in movie theaters.

Index Terms—Experience sharing, social networks, recommender systems

I. INTRODUCTION

The motivation of this paper - and this line of work - comes from the perception that, despite the growing number of activities and events planned around us and despite the abundant sources of information to find out about such events, we are missing out on an overwhelming number of opportunities to have fun, keep fit, learn, or socialize - not because we don't want to, but because we are not aware of the possibility.

The reason, paradoxically, can be related to the overwhelming number of options we are exposed to, made available by a growing number of services on the Internet that often narrow our choices [1]. This is not surprising as it is now ascertained that more options and more information leads to a *narrowing* of the choices rather than to a *broadening*. As a result, not only individuals looking for activities see their ability to choose reduced but also organisations find it more difficult to disseminate the activities they are promoting.

A plethora of web and social applications have grown over the past few years to increase the flexibility, efficiency, and effectiveness of communication beyond what was traditionally provided by newspapers and posters. In most cities, the local tourist agency maintains more or less updated and usable web pages listing events. Travel and hotels/restaurants guides such as *LonelyPlanet*, *Michelin*, etc provides information on the main points of interest and places to eat or sleep, often with some ratings. Businesses, such as movie theaters, also maintain their own web sites listing events at their premises. While all these sources do help, they are incomplete (for example, some show only events, other show only places, etc) and, as we

will see, they don't always include (or rate) the restaurant or business that fits our liking and taste. Finally, they typically lack a social and crowdsourcing dimension.

Social networking sites, microblogging platforms and mobile computing have made possible to move an important part of our conversations online, providing richer and real-time social footprints that were not available before. Prominent examples are partially crowdsourced travel guides such as *TripAdvisor* and *Yelp*, or real-time location-sharing and meetup services such as *Foursquare*. TripAdvisor works on the assumption that crowdsourced ratings are better (that is, closer to the average customer's taste) than expert reviews. It shares word of mouth and works thanks to the trust that is established between users and reviewers [2]. Giving completely freedom to users brings also the problem of fake reviews added by some businesses for increasing their ranking or decreasing the one of their competitors [3]. Foursquare[1] allows users to discover places nearby, add new ones and *check in* to places in order to collect points and possibly get discounts from participating businesses. It is also a social application, where you can receive notification from your friends and know where they are. Foursquare was designed for discovering places, but it is used in many different ways, as the interview and surveys in [4] show. For example, check-ins are used for bookmarking places. Some people always check in at home to let their friends know about their availability, while others consider it as spamming with boring information. The study shows also that the gaming aspect given by badges motivate people in using Foursquare at the beginning while later they tend to use it because of their social network connected to it.

For private events, online invitation systems such as *Evite*, *Eventbrite*, or *Facebook* let us create events and invite friends, also managing attendance and interaction among potential participants. Last but not least, general purpose social networking sites such as Facebook and *Twitter* let us advertise and share event information among our friends and followers.

These systems provide us with information as well as elements that support our decisions around activities - including possibly information on *who else* is attending an event which, as we will see, is an important aspect in the decision-making process. Again, the problems of inaccuracy of information (for

---

[1]https://foursquare.com/

our personal needs and taste) or of too much information mean that we are still left with the perception that we miss out on many opportunities we would like to capture.

In the line of work, studies, and tools described in this paper, that we refer to as *comealong*, we aim at verifying if this perception is true and, if so, at understudying the causes of the problem and how IT can lead us to a possible solution.

The main intuition behind the work is that by motivating people to share their *intentions* in terms of event participation - if properly communicated - can lead to a crowdsosurced way to inform about events and also to participations to events along with people we like. Indeed, the need of social interaction is still present and the participation of social events together with friends bring entertainment and happiness in people's life. Attending an event alone or with friends, to which we can talk about it and share immediate thoughts and comments, are two completely different experiences. For this reason, people tend to avoid attending events alone and they often give up to the event if they are not able to find someone that comes along. But what happens if social networks start to be used for supporting real social life? They could be a mean for finding friends that are interested in attending the same event, for example. Moreover, the activity of sharing the experience, living it together, showing pictures or telling about it, is very important for completing the entertaining activity. Starting from these considerations, we explore how social networks and new technologies can be used to facilitate the experience sharing, enriching the experience gained in social and cultural events.

The second intuition is that planned or dynamic group discounts - or similar benefits - can motivate intention sharing, thereby creating benefits for event participants (more event awareness, better prices) and for businesses (increased attendance). The third intuition (or, rather, "assumption" at this point) is that recommendations are important to separate signal from noise, and that for recommendations to be effective we need to collect a large body of unbiased, trusted ratings and opinions, and to use them to identify similarity in taste. Recommendations are usually done taking in consideration the location where the users are, but proximity (both in time and location) is not the only important aspect people take into account when deciding about events to attend. E. M. Daly and W. Geyer in [5] show that people can attend far events if they know that their friends are going there, so the social aspect of the event is really important. In the platforms described previously only Foursquare has strong support to the social aspect of meeting up, letting the users know where their friends are checking in.

Finally, the fourth intuition is that we need to be able to combine private (e.g., parties) and public (e.g., theater) events in a single platform, because they cater to a same need - finding interesting things to do in one's free time.

If we are successful in exploiting these intuitions and assumptions, then we can transition to a situation where people can find events they might be interested in (mainly thanks to implicit and explicit recommendations from friend), and naturally and dynamically organize to enjoy it along with friends, all with minimal or no search effort.

The challenges lie in verifying if the intuitions make sense and in transforming them to a socio-technical system that achieves the goal. We believe this systems blends in, in the appropriate way and taking the appropriate elements, the characteristic of systems like TripAdvisor, Foursquare, Facebook, Groupon, and Evite. In this paper we describe the system and the initial learnings we take home after a few seed experiments. From a research standpoint, we aim at studying the *comealong* effect, that is, if (and how much) attendance intentions and group attendance benefits affect people's participation and satisfaction. We also aim at studying effective techniques for motivating people to rate a business or event once they have attended, thereby providing a large dataset for recommendations. The specific subgoal of this paper, besides introducing the camealong concepts and platform, are to understand the usefulness of ratings and the effectiveness of recommendations when in presence of a large corpus of unbiased ratings. The interested reader can peek the system and learn about up to date results at http://comealong.me.

In the following, we first presents the results of a study that motivates the need for comealong. Then, we briefly discuss the comealong features, before diving into the experiments related to the effectiveness and efficiency of ratings. The key results presented in this paper are i) the demonstrated need for an event management platform, ii) the inaccuracy of conventional and web rating systems for recommending places to visit, iii) the importance of personalized, similarity-based (Amazon-style) recommendations, and iv) the importance of considering your friends network when computing recommendations (with a caveat, discussed later, related to the hypothesis that considering friends preferences for providing better recommendations is effective because friendship is a proxy for other aspects such as being in the same age group.)

II. ASSESSING THE NEED FOR A BETTER EVENT DISCOVERY

To assess whether people are comfortable with the way they find out about interesting events in their town today, as well as to understand to which extent social motivations lead event attendance, we run a survey of nearly 150 persons [2], where just over a third where in their college years and the remainder were adults up to 60 years old.

The first and, for us, most important result from the survey is that there is a widely recognized need for a better way to be informed about events. In the information age, this is quite surprising, especially given that we are asking people about events in their own town, so respondents might indeed be aware of where to look for information (which newspapers, which web site, etc). Only 10% of the respondents said that they have satisfactory information about what to do in the place where they live. The overwhelming majority feel either

---

[2]You can read more at http://survey.comealong.me

non-informed or are unsure (Figure 1). The results are near identical for youngsters and adults.
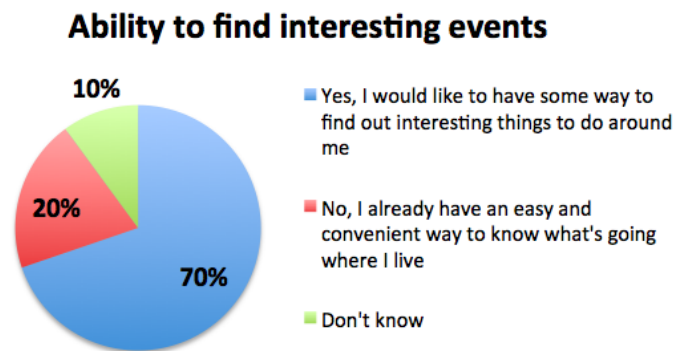
Ability to find interesting events

Fig. 1. Results from the survey on availability and sources of event information in one's own town: Ability to find interesting events.

After this response, we looked at *where* do people get their event information from.
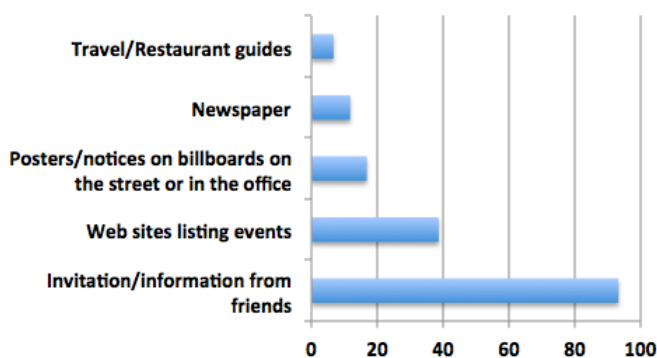
Sources of event information

Fig. 2. Sources of event information

The result, shown in Figure 2, tells us that, in the majority of the cases, event discovery and participation is a social phenomenon: we find out about things to do mainly through invitations from friends, with a whopping 93% of the people citing this as one of their top two sources of information. Nearly 40% also looks at web event-listing web sites, while the rest is rather negligible with, perhaps surprisingly, only 11% of the respondents looks at newspapers, on average. Again the results are similar across age groups, with only one major difference: *less than one percent of the surveyed sample under the age of 25 learns about events from newspapers!*. This number is just short of 16% for adults.

Another data that surprised us is the correlation between sources and ability to find events. As Figure 3 shows, all respondents declare the same need for better event information regardless of where they obtain information from today. Interestingly, the people for which this need is more widely felt are the ones having event web sites as their sources.

Again, we recall that the study focused on event in one's own town. This means that people looking regularly at event
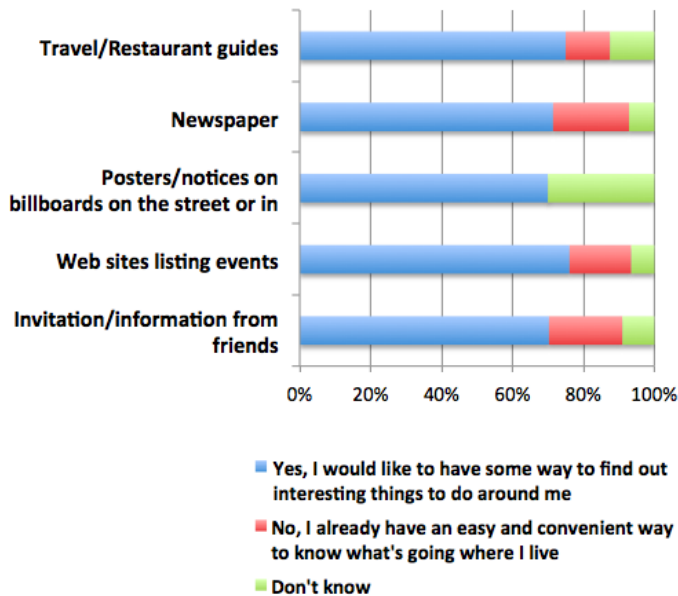
Sources vs ability to find events

Fig. 3. Correlation between sources and ability to find events.

web sites for information do know where to look. The conclusion we draw from the survey is that either these web sites are not sufficiently maintained, or have poor usability, or that interesting events are lost in the noise and are not easily identifiable.

Our final result (Figure 4) is about how important is it to know who else is going prior to deciding to attend an event. Here, more than 70% of the respondents stated that they often or sometimes do look at other attendees, while 27% typically doesn't. The result is analogous regardless of whether we learn about events from friends or from other sources, and is consistent across different age groups. The perhaps rather obvious learning is that the social dimension is very important when deciding which events to take part to.

Armed with this knowledge and convinced on the need for a better event management which is also characterized by a social and crowdsourcing dimension, we proceed to describe the model and system behind comealong.

III. A PLATFORM FOR EXPERIENCE SHARING

Comealong captures the void left by current tools by building on the four intuitions mentioned earlier in the paper. Its goal is to include both public events (e.g., a concert) and private events (e.g., a birthday party) as well as places (such as restaurants and points of interests). Private events are entered by the organizers, while public events and places are entered by the community in a fashion analogous to entries in wikipidia - hence, moderated by the community itself. The reason for having both public and private information is to provide a one-stop shop for event information. The classes of events supported by comealong is rather broad, going from movies to fitness classes to religious events.
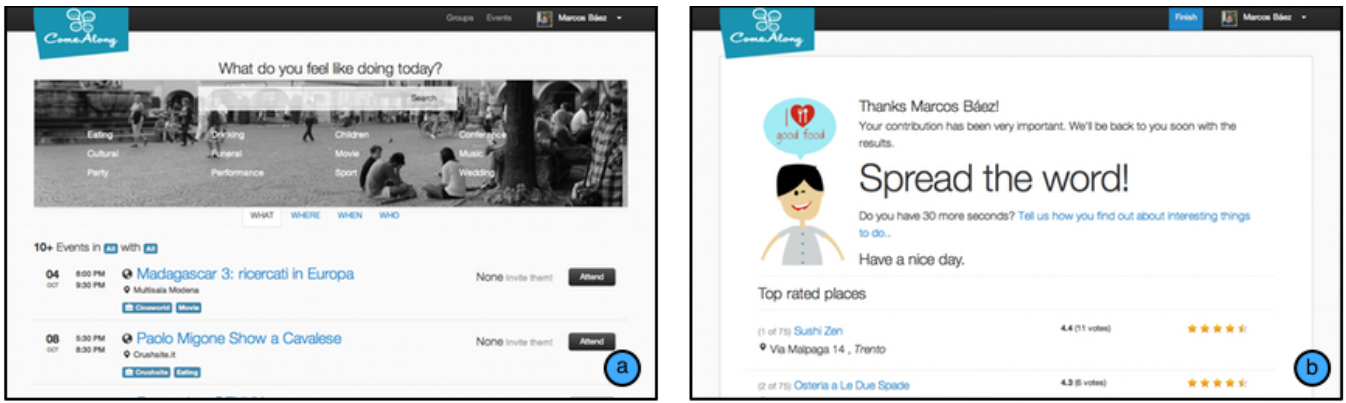
Fig. 5. Comealong UI. a) Interface for discovering events b) Interface for crowdsourcing ratings
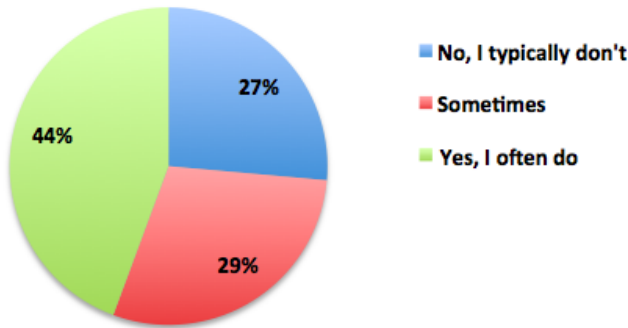

Fig. 4. Looking at what friends do before choosing event attendance.

People can state their *intention* to attend an event. This has several implications: the first is that this is posted on Facebook and possibly twitter if the user allows, and this is likely to generate interest by other people to join. This effect is why we focus on intentions rather than presence, as Foursquare does. In this way, friends can more easily hook up, and business get more customers. The second is that by doing so the poster - and to some extent all followers as well - may get benefits from announcing their intention. These benefits may range from Foursquare-like points to actual discounts negotiated with businesses. For example, comealong users announcing their intention to attend a movie showing at least 2 hours before showtime enjoy a 20% discount at selected theaters. Furthermore, businesses can have a console at their disposal to monitor bookings and have a way to communicate with prospective event attendees. For example this is useful for fitness classes (that are indeed a social event), where a gym wants to communicate to members that a certain class has been canceled due to unavailability of the instructor.

From the platform side, intention declaration allows us to i) learn the preferences of the person, to then provide recommendations, and ii) obtain ratings for the event by pushing a request to the mobile app once the event has completed. This is crucial to the success of the platform as getting a large number of unbiased ratings is the essence of any recommendation system.

During the event, the mobile comealong app allows to take and share pics in real time with all attendees and with the event owner. This can be used to display right there and then on a large screen (this can be for example appealing for weddings or large parties, along the lines of what *weddingsnap* does) or for people to look at them *after* the event.

The system is expected to go into operation on October 31st following the initial comealong agreements with local businesses and limited to the city of Trento, Italy (the system works worldwide, but the business agreements for providing comealong discounts will only be in place for Trento, initially).

At this point we are not yet ready to study the comealong effects (increased participation and increased because of lack of a meaningful dataset. What we focused on in this stage in terms of study is to understand the power of collecting a large number of trusted and unbiased ratings from locals (as opposed to tourists) - which is one of the intended goals of comealong - and to understand how to provide recommendations based on this.

IV. EVALUATION: DO YOU LOVE FOOD?

A. Sample Population and comparisons with baselines

To study the effectiveness of ratings, we collected a rather large set of ratings about restaurants in Trento (694 ratings from 90 users, over a total of 75 restaurants we considered in the study, taken from TripAdvisor's top 75 ranked restaurants). Ratings where on 1-5 scale as TripAdvisor (TA) does. However, our population was people who lives in Trento and therefore knows the restaurant scenes rather well. Furthermore, each person rated on average just short of 8 places in the same town. We do not have in our survey ratings from occasional visitors who has only been in one or two places, and we do not have also ratings from restaurant owners or their competitors. The population in the study is composed of relatively young

adults that have a Facebook account (we required a Facebook login to perform analysis of the effectiveness of friends-based recommendations). This means that our population was mainly composed of young adults, but we don't consider this to be a bias because the sample is also the same as the population we target.

We take therefore the results of the ranking as a representation of the average opinion of locals.

As a first part of the analysis, we studied how well results from popular travel guides match our ground truth. We took as notable example the Lonely Planet 2012 (LP) guide for Italy[6]. The result is shown in Figure 6. The guide lists five restaurants. One of them is not among TripAdvisor's top 75 and therefore is not in our comealong test list. The other four are, interestingly, in very different position in the three rankings. LP considers them to be the four most interesting restaurants to visit, for comealong they are more or less evenly distributed among the top half but are quite far from the top 4 (with one exception), and the results are analogous for TA (the comealong and TA rankings are also very different, but we'll get back to this later). We can say that LP did not, in this case, guess the preferences of our test users population.

Lonely Planet	ComeAlong ranking (out of 75)	TripAdvisor ranking
Scrigno	5	10
Uva e menta	15	55
Tre Garofani	20	7
Pedavena	37	47

Fig. 6. LonelyPlanet recommended restaurants for Trento

As another baseline analysis, we consider TripAdvisor. Instead of using correlation coefficients we take a very pictorial way of performing the analysis: we consider the effectiveness of the top 10 rated restaurants on TA (assuming we would look at those to select where to go) and see how well we would do with that.

As Figure 7. shows, in three cases we would pick one of the comealong top 10 restaurants, while in the other cases we pick places that are rather far down the ranking. Indeed, given that we have a population of 75 restaurants, 6 out of these top 10 are around or below that midpoint!

Restaurant	Tripadvisor	Comealong
Due spade	1	1
Duo	2	35
Vecchia Macina	3	35
Albert	4	35
Margon	5	63
Madruzzo	6	5
Tre Garofani	7	20
Gusto	8	54
Pineta	9	69
Scrigno	10	5

Fig. 7. TripAdvisor top 10 recommended restaurants for Trento

Given these results, we are confident that we need novel ways of collecting a large number of unbiased ratings as current technology and approaches seem to suffer from limitations, at least in our experience. This further motivates the need for a comealong-like platform, which has the potential of collecting ratings more than what current systems do thanks to the incentive and to the ability of asking feedback on the spot.

We now turn to look at how we can best use such ratings to perform accurate and personalized recommendations. Specifically, we first see how the overall comealong rating are accurate in giving recommendation to a user, and then compare this with recommendations given by considering only your friends' rating, by considering ratings from people with similar taste as yours (that rated similarly), and last by considering the ratings by similar friends.

B. Formal experiment definition

Let $\mathcal{U}$ be the set of all users and $\mathcal{P}$ the set of all the places. Liked represents the relation $(user, place) \in \mathcal{U} \times \mathcal{P}$ of places users rated positively, and Disliked the relation $(user, place) \in \mathcal{U} \times \mathcal{P} \setminus$ Liked of places users rated negatively. FriendOf denotes the relation $(user, user) \in \mathcal{P}^2$. Given these basic elements we define:

$$\mathsf{Known}(u) = \{p \in \mathcal{P} \mid \mathsf{Liked}(u, p) \vee \mathsf{Disliked}(u, p)\},$$

$$\mathsf{Unknown}(u) = \mathcal{P} \setminus \mathsf{Known}(u)$$

to denote the sets of places the user have rated and the ones yet to discover.

For the scoring function, we compute the average rating of the network in consideration:

$$\mathsf{score}(p, Net(u)) = \frac{\|\mathsf{Likes}(p, Net(u)) - \mathsf{Dislikes}(p, Net(u))\|}{\|Net(u)\|}$$

$$Net(u) = \{u' \in \mathcal{U} \mid \mathrm{sim}(u, u') > \delta\}.$$

$$\mathsf{Likes}(p, Net(u)) = \bigcup \{u' \in Net(u) \mid \mathsf{Liked}(u', p)\}$$

$$\mathsf{Dislikes}(p, Net(u)) = \bigcup \{u' \in Net(u) \mid \mathsf{Disliked}(u', p)\}$$

Finally, we define the recommendation of places p to a user u as:

1) $\mathsf{Rec}(u, k) \subseteq \mathsf{Unknown}(u)$,
2) $|\mathsf{Rec}(u, k)| = k$,
3) $\forall p \in \mathsf{Rec}(u, k) \ \forall p' \in (\mathsf{Unknown}(u) \setminus \mathsf{Rec}(u, k))$ $(\mathsf{score}(p) \geq \mathsf{score}(p')).$

C. Definition of recommendation strategies

We tested four different settings of the above definition to compare the effectiveness of the different networks to recommend places:

1) Recommendation based on overall popularity: In this setting we consider the ratings from all users to compute the score (score_o):

$$sim_o(u, u') = 1, \forall u, u \neq u'$$

TABLE I
EVALUATION OF THE DIFFERENT SCORING FUNCTIONS, K=5

Score function ($k = 5$)	Precision		Recall	
	Av.	sd	Av.	sd
score$_o$	0.09	0.15	0.08	0.12
score$_f$	0.16	0.21	0.12	0.17
score$_{rt}$	0.12	0.18	0.08	0.14
score$_{st}$	0.25	0.26	0.14	0.16

TABLE II
EVALUATION OF THE DIFFERENT SCORING FUNCTIONS, $k = 10$

Score function ($k = 10$)	Precision		Recall	
	Av.	sd	Av.	sd
score$_o$	0.11	0.12	0.16	0.16
score$_f$	0.17	0.16	0.22	0.20
score$_{rt}$	0.13	0.13	0.17	0.21
score$_{sf}$	0.22	0.21	0.25	0.28

2) Recommendation based on friends : This setting makes use of the network of friends to compute the score (score$_f$):

$$sim_f(u, u') = 1, \exists \mathsf{FriendOf}(u, u')$$

3) Recommendation based on similar users: This setting makes use of people with similar taste to compute the score (score$_{rt}$). The similarity function in this case is defined as:

$$sim_{rt}(u, u') = \frac{\|\mathsf{Co\_liked}(u, u') \cup \mathsf{Co\_disliked}(u, u')\|}{\|\mathsf{Known}(u) \cap \mathsf{Known}(u')\|}$$

$$\mathsf{Co\_liked}(u, u') = \bigcup\{p \in \mathcal{P} \mid \mathsf{Liked}(u, p) \wedge \mathsf{Liked}(u', p)\}$$

$$\mathsf{Co\_disliked}(u, u') = \bigcup\{p \in \mathcal{P} \mid \mathsf{Disliked}(u, p) \wedge \mathsf{Disliked}(u', p)\}$$

4) Recommendation based on similar friends: This setting makes use of friends with similar taste to compute the score (score$_{sf}$). The similarity function in this case is defined as:

$$sim_{sf}(u, u') = sim_{rt}(u, u'), \exists \mathsf{FriendOf}(u, u')$$

D. Evaluation of the different recommendations algorithms

We evaluated each recommendation strategy by computing all variations of the *score* function for all users u on every place $p \in \mathcal{P}$, dropping the place from the set of known places $Known(u)$ as the initial condition. Then we took the top k places and compared this list with how many places the user actually liked computing *precision* and *recall*. The summary of the results can be seen in Table I and Table II.

In the evaluation we use the following definition of *precision* and *recall*:

$$precision = \frac{\|\mathsf{Tp}(u)\|}{\|\mathsf{Tp}(u)\| + \|\mathsf{Fp}(u)\|}$$

$$recall = \frac{\|\mathsf{Tp}(u)\|}{\|\mathsf{Fp}(u)\| + \|\mathsf{Fn}(u)\|}$$

$$\mathsf{Tp}(u) = \bigcup\{p \in \mathsf{Rec}(u) \mid \mathsf{Liked}(u, p)\}$$
$$\mathsf{Fp}(u) = \bigcup\{p \in \mathsf{Unknown}(u) \mid \mathsf{Disliked}(u, p)\}$$
$$\mathsf{Fn}(u) = \bigcup\{p \in \mathsf{Known}(u) \setminus \mathsf{Rec}(u) \mid \mathsf{Liked}(u, p)\}$$

The results suggest, not surprisingly, that recommendations coming from similar users (based on rating behaviour) performs better than overall popularity. More interestingly, recommendations coming from friends, and specially if reduced to similar friends, outperforms considerably the other algorithms in both precision and recall. These results are promising and points to the potential of balancing personal tastes (in this case captured by rating behaviour) with the real social context (friends).

We should interpret the higher precision as higher chances of finding interesting events, and the higher recall as the reduced feeling of missing out something. When put in real context, these results indicate that events in newspaper and tabloids have less chances of providing useful suggestions than a comealong-like application.

V. RESULTS

The lesson we learn from the analysis is that ratings from travel guides and web sites do not satisfactorily predict our taste and in general are only partially adequate to recommend us good places to eat (and, likely, the same applies to other categories of places and events). Specifically, an interesting take home point is that not only unbiased ratings from locals differ from those of travel guides, but also that considering friendship and similarity have a profound effect on the accuracy of the recommendations. We assume that the friendship effect is not so much related to being friends but rather to being in the same age group, but we did not have the chance to verify this. Recommendations based on similar friends provide by far the best result (almost three times better than overall comealong ratings), combining (we believe) age with taste similarity.

We are currently with a beta version of the system, working as a crowsourcing platform for gathering feedback about local places in the Trento, and going through the alpha version of the full system. More information about the release dates can be found at http://www.comealong.me.

REFERENCES

[1] J. Sandweiss, "The future of scienti!c publishing," *Physical Rev. Letters*, vol. 102, no. 19, pp. 106–113, 2009.
[2] I. Jeacle and C. Carter, "In tripadvisor we trust: Rankings, calculative regimes and abstract systems," *Accounting, Organizations and Society*, vol. 36, no. 4–5, pp. 293 – 309, 2011.
[3] A. Fong, "The influence of online reviews," *Journal of Digital Research & Publishing Semester 1 2010 (7pm class)*, pp. 106–113, 2010.
[4] J. Lindqvist, J. Cranshaw, J. Wiese, J. Hong, and J. Zimmerman, "I'm the mayor of my house: examining why people use foursquare - a social-driven location sharing application," in *Proceedings of the 2011 annual conference on Human factors in computing systems*, ser. CHI '11. New York, NY, USA: ACM, 2011, pp. 2409–2418. [Online]. Available: http://doi.acm.org/10.1145/1978942.1979295
[5] E. M. Daly and W. Geyer, "Effective event discovery: using location and social information for scoping event recommendations," in *Proceedings of the fifth ACM conference on Recommender systems*, ser. RecSys '11. New York, NY, USA: ACM, 2011, pp. 277–280. [Online]. Available: http://doi.acm.org/10.1145/2043932.2043982
[6] P. e. a. Hardy, *Italy Travel Guide*. Lonely Planet, 2012.