

Rating Friends Without Making Enemies

Lada A. Adamic^{1,2}, Debra Lauterbach¹, Chun-Yuen Teng¹, Mark S. Ackerman^{1,2}

¹School of Information, ²EECS
University of Michigan
Ann Arbor, MI 48109, USA
{ladamic,dlauter,chunyuen,ackerm}@umich.edu

Abstract

As online social networks expand their role beyond maintaining existing relationships, they may look to more faceted ratings to support the formation of new connections between their users. Our study focuses on one community employing faceted ratings, CouchSurfing.org, and combines data analysis of ratings, a large-scale survey, and in-depth interviews. In order to understand the ratings, we revisit the notions of friendship and trust and uncover an asymmetry: close friendship includes trust, but high levels of trust can be achieved without close friendship. To users, providing faceted ratings presents challenges, including differentiating and quantifying inherently subjective feelings such as friendship and trust, concern over a friend's reaction to a rating, and knowledge of how ratings can affect others' reputations. One consequence of these issues is the near absence of negative feedback, even though a small portion of actual experiences and privately held ratings are negative. We show how users take this into account when formulating and interpreting ratings, and discuss designs that could encourage more balanced feedback.

Introduction

Many online systems rely on human ratings to enhance their product and content recommendations. Typically, users have little incentive to provide anything but a honest rating, and research has centered on addressing either rating sparsity or manipulation by those few who stand to gain by altering them. Little is understood, however, about the challenges that users face in rating other *users*. Most online social networking sites (OSNS) simply ask a user to specify whether or not someone is a friend. Acknowledging the awkwardness users feel in declining an unwanted friend request, OSNS typically do not alert the other user of a rejection or removal from a friend list. This, along with connection-specific privacy settings, has helped smooth interactions.

Binary designations of "friend" may be sufficient for typical social network interaction tasks for an individual user, but they are not especially informative to *other* users. Knowing that someone has 1000 Facebook friends does not reveal whether that person is particularly trustworthy. But trust is an important criterion in many emerging social computing

applications, ranging from local resource (e.g. tool) sharing sites to hospitality networks. So far, little is understood about how one would be able to capture the richness and multidimensionality of human relationships, because very few OSNS incorporate multifaceted ratings.

CouchSurfing.org (CS), the largest hospitality exchange network, is one site that does ask its users to quantify their relationships, using more finely grained and faceted ratings than most other social networking sites. Established in 2003 (before Facebook, MySpace, or LinkedIn), it was among the first social networking sites and has since steadily grown to over 2 million members, representing every country in the world*. CS fits the description of a social networking site (Boyd and Ellison 2008): it allows users to construct a profile, articulate those with whom they share a connection, and view their list of connections and those of others. Its primary purpose, however, is to enable new connections: to connect travelers with like-minded hosts who can provide a couch to sleep on. A host or traveler can use the multifaceted rating system in deciding with whom to connect.

In this paper we use the CouchSurfing setting to conduct complementary large-scale data analysis, survey and interviews, in order to analyze the problem of eliciting and quantifying human relationships. We start by trying to understand the relationship between two aspects of a relationship that one might wish to quantify: friendship and trust. We find that despite variation in how individuals differentiate and evaluate the two concepts, some large scale patterns emerge, including a correlation and asymmetry in how friendship and trust grow. We further examine how quantifiable these concepts are, and how the design of the rating system shapes the honesty and usefulness of such ratings. We discover operational difficulties in quantifying human relationships, including using one-dimensional labels and accounting for reputation and reciprocity. We conclude with a discussion of our findings and implications for design of rating systems on online social networking sites.

Related work

The present work is motivated by our previous study of several online reputation systems, including CS, that found publicly shown ratings of other users to be more reciprocal and

positive than ratings that were held private, or given anonymously (Teng, Lauterbach, and Adamic 2010). The findings, based only on quantitative analysis of ratings, left several questions unanswered. Were friendship ratings more reciprocal than trust ratings because they were shown publicly, or is there a fundamental difference in the nature of trust and friendship? Are negative ratings largely absent because of fear of reciprocal action, or are there other causes? These questions are addressed in the present paper through a survey and in depth interviews, which informed additional quantitative analysis of rating data.

One approach to understanding ratings is to attempt to predict them from other factors. Bialski et al. (2007) correlated trust on CS with the origin, duration and context of the friendship. Outside of CouchSurfing there have also been efforts to derive trust and friendship strength from information other than direct ratings. Gilbert and Karahalios (2009) tied perceived tie strength to user behavior on Facebook. Caverlee et al. (2008) utilized users' behavior and feedback to generate a trust score. Golbeck et al. (2009) showed that trust scores between users are correlated with similarity of users' profiles and preferences. Skopik et al. (2009) proposed determining trust relationships by mining communication patterns between users. Robert et al. (2009) formulated trust mediated by information technology as being based on initial information and accumulated experience.

Several studies have aimed to predict trust ratings between users by applying prediction algorithms to the trust network itself. For example, Jøsang et al. (2005) examined transitivity in trust networks, while Leskovec et al. (2010) utilized balance theory to predict link valance in a number of empirical networks with signed edges. Guha et al. (2004) predicted trust and distrust on the Epinions dataset with an algorithm that propagated trust along trusted edges, but halted at distrusted ones. However, Lauterbach et al. (2009) showed that the network-level variables were poor predictors of ratings on CS relative to variables reflecting direct experience between individuals. Therefore, in this study we focus on how these experiences affect trust and friendship, and how well online ratings and reputation systems are able to harness and quantify such concepts.

There is a rich literature on trust in a range of contexts, from intimate relationships (Rempel, Holmes, and Zanna 1985) to online transactions between strangers (Cheshire and Cook 2004). Our work distinguishes itself in two ways: it is the first to provide large-scale quantitative data on the relationship between trust, friendship, and other variables, and it addresses how design choices in online reputation systems can affect their expression.

Data and methods

Our study was conducted in three phases, intended to provide complementary insights into how users give and interpret ratings on CouchSurfing. In the first phase, we analyzed data which was provided to us by CS in October 2009. This anonymized data includes basic profile information for each CS user (city, country, member since, profile views, etc.) and data for each friendship connection (how they met, when they met, how many times they have hosted or

been hosted by each other, along with trust and friendship ratings. *Friendship degree* ranges on a 7-point scale from 'never met' to 'acquaintance' to 'CouchSurfing friend' to 'friend' to 'good friend' to 'close friend' to 'best friend'. The ordinal *trust rating* is on a 5 point scale, from not trusting the person to trusting the person with one's life, with a separate option stating that one does not know the person well enough to judge. A user's profile page shows all of their friends and how those friends rated their friendship degree with the user. The trust rating, however, is not shown to any other users. In all, the dataset contained 648,099 user profiles and 3,011,487 reciprocal ratings (out of 3,115,548 directed edges in the dataset).

Separately from the initial ratings, one can leave written *references* for people whom one has had experience with. These references are usually given by hosts and guests to each other after the guest has left. References also include ratings of the experience as either positive, neutral, or negative. Finally, once *vouched* for by 3 other users, a user can start vouching for others, indicating that they are highly trustworthy. Vouches, along with reference ratings and text, are also shown on a user's profile.

Table 1: CouchSurfing Rating types

characteristic	type	shown on profile
friendship	ordinal	yes
trust	ordinal	no
vouch	binary	yes
reference	pos/neg/neutral + text	yes

In a second phase, we conducted a survey of CS users, posing questions aimed at understanding the rating patterns in the large scale data. We pilot tested the survey with members of a local CS group. The survey was in English, and contained 35 questions, including several opportunities for open-ended responses. For the open-ended questions, two raters coded the questions. Cohen's kappa for inter-rater reliability of the coding ranged from .79 to .86, all values that are considered very acceptable. The survey link appeared on the CouchSurfing.org homepage for a random sample of visitors to the site over a three-day period in March 2010. The survey received 527 complete responses, and 12 partial responses which were omitted from the analysis.

In our third phase we wanted to understand more about the cultural context and rationales for users' decisions about ratings. We conducted follow-up interviews with a subset of the survey participants who had provided their email addresses. Eighteen interviews were conducted via Skype in English, 17 using voice and 1 using IM. The interviews were semi-structured and lasted between 20 and 45 minutes. With the participants' permission, the calls were recorded and later transcribed. We used standard qualitative techniques (Miles 1994) to analyze the interview data.

All data reported here have been made anonymous. We have done minor edits on the quotations for readability.

Table 2: Demographics of CS users and our survey sample

	Survey sample	Entire population
Gender	Male: 53% Female: 45%	Male: 50.2% Female: 42.0%
Age	Average: 27.3 min:18, max:85	Average: 28 min: 18, max: 89
World region		
Europe	58.3%	50.7%
North America	16.3%	27.9%
South America	4.9%	6.1%

Participants

The demographics of our survey respondents (n=527) roughly match the demographics of the CS population as a whole, which are listed publicly on the CS site[†] (see Table 2). There were nearly as many females as males (53% male, 45% female, with 2% not identifying), and most respondents were in their early-to-mid twenties (M = 27.3, SD = 7.55, min = 18, max = 85). Respondents hailed from 68 countries, with the most frequently reported countries being those in Europe (14% Germany, 10% France, 5% Italy, 4% Poland), and North America (12% United States, 4% Canada). Half of the respondents reported being a member of CS for 1-3 years, with 36% being a member for less than a year, and 13% for more than 3 years. The majority of respondents had at least one experience of having hosted or been hosted by another user, while 15% had not yet done either activity.

We interviewed 10 Americans and 9 Europeans, who ranged in age from 21 to 52. Ten of the interviewees were female. All but one had either surfed or hosted, and most were seasoned CouchSurfers.

Couchsurfing culture and norms

The CouchSurfing organization’s core belief, as stated in their mission, is that travel should be about connecting people and building meaningful connections across cultures. The level and type of involvement in CouchSurfing differs from user to user. Some prefer to host as it allows them to meet interesting people from all over the world without leaving home. Others use the site mainly for travel. Some join the site to find hosts for an upcoming trip but do not stay involved after, while the most committed core of users live nomadic lives, surfing from couch to couch. One participant we interviewed had been couchsurfing nonstop for a year.

The CS culture fosters trust, which incorporates faith, dependability, predictability, and a belief that others’ motives are altruistic (Rempel, Holmes, and Zanna 1985). As P9 discovered in her first experience being hosted, CouchSurfers have faith that others will behave in the CouchSurfing spirit:

We arrived in Brussels at 9 o’clock in the morning and the guy that we were staying with, our host, had to come to his work and he just handed his house keys right over to us and said, “This is where I live. Go here, go here and my wife and I won’t be home until after 5.”

[†] <http://www.couchsurfing.org/statistics.html>

Overwhelmingly, the survey respondents and interviewees cited positive experiences, in line with CS’s culture. None mentioned personally experiencing theft. However, several women did mention a general concern about unwanted sexual advances and safety.

Negative experiences recounted to us, such as a host leaving a traveler stranded at their doorstep or a traveler not showing up as agreed, tend to violate the expectation of predictability and dependability inherent in trust. Some smaller, but significant faux-pas were committed by individuals who did not share in the cultural consensus, especially if they were new to CS. Several interviewees stated that they would not host people who were clearly looking for “just a place to sleep” [P13] or showed little interest in the host.

He lacked the couchsurfing spirit (not the sharing type of person) and was using it merely as a cheap way to travel... [S35]

On the other side, there were instances of hosts providing unclean or crowded accommodations. Sometimes the problem was difficult to pinpoint: the host or surfer was just “weird”.

Results

Overall, the rating system on CouchSurfing seems to work well, as can be seen by the large number of couchsurfing experiences that have taken place, and the general satisfaction of participants with the rating system. However, while participants may view the rating system positively overall, our findings uncovered several key issues affecting how easily and accurately ratings are given, and their perceived utility.

First, we discuss participants’ differing definitions and views of friendship and trust at a semantic level. We then discuss why friendship, trust, and experience can be difficult for users on CouchSurfing to rate. We claim this is hard for four reasons: *difficulty quantifying relationships on a pre-defined scale*, *concern over a friend’s reaction to a rating*, *knowledge of how ratings can affect one’s reputation* and *timing of ratings relative to the evolution of the relationship*. These problems are related to both the inherent complexity of human relationships, and the design of the rating system.

We begin with two sections discussing why users found rating friendship and trust conceptually difficult, and follow these with two sections discussing how users were concerned with the effects of their ratings on other CS users. Implications for designs to address these problems are discussed at the end.

Understanding trust vs. friendship

Much work on online communities, in particular work on algorithmic approaches to automated reputation ranking, assumes that trust is inherent in friendship and that one can serve as a proxy for the other. By gathering the two ratings separately, CouchSurfing provides a unique opportunity to study the two concepts in relation to one another and to other factors. Generally, the two quantities show a moderate correlation ($\rho = 0.57$) across 3 million paired trust and friendship ratings. However, some distinct asymmetries emerge, as shown in Figure 1. When one user rates another as a close

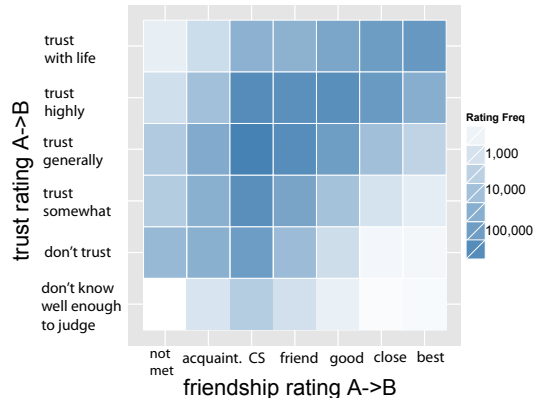


Figure 1: Correspondence between the friendship and trust levels one individual assigns to another.

or best friend, the corresponding trust level tends to be high, with little variance. However, a high trust level can occur with a range of friendship levels.

The interviews and surveys elaborated this. One interviewee’s comments mirrored the findings exactly:

I think close friends you trust, but I don’t think everyone you trust is a close friend. [P12]

Several others made similar observations, one realizing that friendship was more ambiguous for her.

Friendship includes trust. You can trust someone, but still without that person being a friend. I guess friendship is a more elusive concept and therefore more difficult to judge. [S312]

On the other hand, another felt high levels of trust were reserved for just close friends:

In general I only *really* trust my closest, real life friends. That takes years to earn. [S256]

There is also a time element for some users. Some concurred that trust takes time, while mentioning that friendship can happen instantly:

Trust takes, you know, in a lot of cases years to build, whereas friendship and that sort of thing can sort of happen instantly [P8]

But others followed their gut feeling about trust:

I have a gut feeling about who I could trust, but not so much about who actually counts as my friend. [S10]

Given the widely varying responses our survey and interviews elicited about how friendship and trust evolve, we turned to the data on ratings. Table 3 shows that, on average, both friendship and trust grow with time, but even more so with familiarity. Naturally, the connection may go both ways: staying friends and getting to know someone if one likes and trusts them, or liking and trusting someone more the more one gets to know them. Table 3 also suggests that co-habitation builds or requires trust, but friends *choose* to travel together, since the number of days traveled together

	friendship	trust
how well	0.719	0.670
log(how long)	0.592	0.365
days traveled	0.404	0.250
same country	0.212	0.075
abs. age difference	-0.106	-0.068
days surfed	0.268	0.203
days hosted	0.279	0.193
same gender	0.078	0.045

Table 3: Pearson correlations between trust, friendship, and other attributes of the relationship (all $p < 0.001$)

is more revealing of friendship than trust. Trust is the more unpredictable of the two, with duration of acquaintance and familiarity accounting for 41.7% of the variance in trust ratings, relative to 57.7% of the variance in friendship ratings. This unpredictability of trust ratings is perhaps captured by the following comment:

In part, it’s a question of hours that we stay together. If we still stay together the level of trust increases... Not always, but increases. And sometimes you understand some limitations so maybe your level of trust can arrive to a certain level and that’s it, and doesn’t increase over that level. [P10]

Finally, the variance in ratings by recipient is slightly higher for friendship than trust, even controlling for number of friends, suggesting that between friendship and trust, trust is slightly more associated with the individual, while friendship is more a property of the relationship ($\Delta z_{var} = 0.015, t = 12.95, p < 10^{-10}$).

The next section discusses how these types of ambiguities, uncertainties, and differences in individual interpretations, along with the manner in which ratings are elicited, affect whether trust and friendship can be easily quantified by users. We discuss each of the four issues identified in turn, beginning with the problem of quantifying characteristics like friendship and trust on a pre-defined scale.

Quantifying relationships on a pre-defined scale

Although, on average, survey respondents evaluated trust and friendship as being equally difficult to rate, several considered trust to be more amenable to being put into “levels” and friendship to be more subjective and multi-dimensional. One respondent thoughtfully pointed to the complexity of the underlying concept and the difficulty in reducing this to a unidimensional scalar:

Social networks (and CS has some aspects of it) are... the objectification of something as subjective as a human relation. One “level” is never enough to point to the correct tone of a human relation. [S291]

Others felt that designating others as “friends” was not meaningful.

I think it would be much better if it was called with another name than friends’ list. ... I don’t feel the need

to "officialize" the fact that we're friend with someone. The label "friend" is completely useless [P16]

Even for users comfortable in assigning friendship to a linear level, a further complication arises from the "Couch-Surfing friend" designation, applied to a full 45% of all connections. The label "CS friend" is intended to capture the origin and context of a friendship. However, to many surveyed, the implied strength of this label was ambiguous:

It's hard to state what a 'CS friend' really means. I have no problems with determining higher level friendships though. [S68]

Interviewees had diverging rules for designating someone as a CS friend. To some it signaled a lack of closeness:

A CouchSurfing friend sounds like you wouldn't be friends with them if they hadn't gone and stayed on your couch [P12]

I use it sometimes, when I really didn't have many things to do with the other one. [P18]

Others had no difficulty applying the "CS friend" label to anyone they met through CS.

Friendship level - easy, those who I hosted I mark as CS friends. I've never had negative experience but I cannot say how trustful is one who I know 1 or 2 days. [S288]

The effect of language and culture on ratings A further problem with designating fixed labels is that an expression may not translate well into all languages. A CS user from Italy commented:

Best friend, I understand what that means but between good friend and close friend, I don't know exactly what is the higher level [P10]

A Spanish interviewee said:

Some of these things are ... directly translated from English to our ... language. To say I trust this person with my life, I think that's such a hyperbole. [P5]

Indeed, Spanish users used this highest trust designation 8.6% of the time, while Americans used it 13.9% of the time, and all Europeans 10% of the time. This brings us to the question whether there may be cultural differences in norms in expressing friendship, in addition to language-specific idioms. For example, one European interviewee thought Americans expressed themselves differently:

Americans have an interesting way of putting things. they can write "awesome, great, super" but don't mean it. [P19]

while another commented:

American people, [laughter] they tend to be very open ... Yeah, you're my friend even though you meet them for the first time. [P11]

However, it should be noted that while world-wide there are country-by-country differences in average ratings (Teng, Lauterbach, and Adamic 2010), Americans had a negligibly higher average friendship rating (3.91) relative to western Europeans (3.86), and the same held true for trust.

In short, users' difficulty in rating friendship and trust suggests caution to those interpreting these ratings. Not only do CS users consider friendship and trust differently, they also have considerable difficulty making the fine distinctions that the rating scale demands. In the next two sections we discuss how concern over the effect of ratings on other CS users can further impact how ratings are given.

Reciprocity

Even if a user is able to interpret given categories and sort their feelings into them, when these selections are public, they have to take into consideration the fact that the other person might notice how they were rated.

Reciprocity in friendship The expectation is that friendship is mutual ("trusting someone else is more of a one-way feeling than deciding level of friendship (which is two-way [S121])", but one must choose a friendship level before they are shown how the friend rated them. Nine survey respondents attributed the difficulty in selecting friendship level to this matching problem, e.g.:

It can be difficult to select a friendship level if I am unsure of how the other person may react or if I think they may see our friendship as being at a different level. [S114]

Cause sometimes you don't want to be unpolite by just adding the person as "Couchsurfing friend" ... you want to have the person the same friendship level. [S175]

The trust level is anonymous, and I tend to trust people more easily. that's why friendship level is more difficult: everyone can see it [S276]

Previously, we found that mutual friendship ratings on CS tend to be more aligned than mutual trust ratings (Teng, Lauterbach, and Adamic 2010). Here we see that there are two possible explanations for this, the first being that the ratings reflect a true difference in mutuality between friendship and trust, the second being that users take greater care to select an appropriate friendship level because it is public.

In fact, only 41.7% of our survey respondents ever noticed selecting a friendship level different from what their friend had designated. When asked how they felt after noticing a difference in friendship level, the responses varied. Seventy-nine of the 201 freeform responses expressed neutral sentiment. They felt "OK", "neutral", "didn't care", "it's normal", "it doesn't matter". Twenty-six had felt "weird", "strange", "confused" or "awkward", with some feeling curious as to why there is a difference, but not curious enough to ask. Some individuals who received higher ratings than they gave felt "good", "happy", "grateful", "honoured" and "flattered" (12 cases), but sometimes also guilty or ashamed (6 cases). On the other hand, 25 individuals who had rated someone more highly than that person rated them in turn felt "disappointed", "not good", or "bad", e.g.

I once said one girl was a "good friend" - however, she added me as an acquaintance. It actually made me feel quite bad to hear that she didn't even consider me as her friend [S491]

Not a big deal but yes it feels not great. Because you see that the feelings about the friendships is not really mutual. [S31]

Those who were unbothered by the difference or grew to accept it after initially feeling awkward or disappointed understood that friends experience a relationship differently (39 cases), or realized that the same level, e.g. “close friend” may be defined differently by different individuals (17 cases). A further 7 mentioned that they were not bothered because the difference was small or they thought their friend had not updated the friendship level to reflect how the friendship had evolved.

Several of our interview subjects did not recall how they had rated friends when prompted.

It's not that important. I couldn't even tell you how I've graded a lot of them to be honest. So it's not something I really think about too much. [P12]

A handful of the survey respondents mentioned that what matters is the real experience:

I do not care, the real relationship between [two] people is something that goes beyond a label in CS website. [S324]

The relaxed attitude about differences in level is also reflected in the small percentage of the survey respondents, 15.8%, who adjusted their friendship level to match after noticing a difference. In fact, only 37.3% of those surveyed recalled having ever modified a friendship rating. A contributing factor may be that one is not able to adjust one's rating on the sly:

About updating the friendship level - The other person gets a message about the update, which feels a bit weird. You can't help but ask yourself what you've done recently to cause this. [S252]

Awkward. Afterwards the friend changed the friendship level to match my choice. Which also felt awkward. [S387]

Reciprocity in references While friendship ratings and reciprocity therein may not be deemed critical by many CS users, most were attuned to reciprocity in references. Reciprocity starts with the CS norm, mentioned though not held by all, that both the host and surfer will leave a reference. 55.3% of the survey respondents who had surfed or hosted always leave a reference, and another 35.6% leave a reference most of the time.

I usually don't write references to those I hosted and didn't leave a reference on my profile [S399]

We hosted them & because they didn't leave a reference, we thought it was rude & decided not to leave one also [S59]

Reciprocity extends to the content of the reference:

Generally, I prefer the host or the surfer to leave a reference first, so I can ... gauge on the way that their reference was. If their reference was very detailed and

inclusive of our expense then I will meter that, and if it was a little more scoped. [laughter] then I will narrow mine down too. [P9]

Reciprocity not only plays a role in whether users leave references for each other, but appears to be a contributing factor in the near absence of negative references on the site: The ratio of positive to negative references on CouchSurfing is 2,500:1 (Teng, Lauterbach, and Adamic 2010). Of our survey respondents, 84.9% had never left a negative or neutral reference, while according to the ratings dataset users leave a positive reference for 87.7% of those they host, and for 90.1% of those who host them. This does not mean that all experiences were positive, and our survey revealed that at least some of those “missing” references reflected negative or neutral experiences that were not reported. When prompted why they sometimes did not leave a reference, 51.3% listed being too busy as one of the reasons, but 31.7% had had a neutral experience and another 12.1% had had a negative experience.

Of the 50 freeform answers explaining why respondents had not felt comfortable leaving a neutral or negative reference, 9 of them attributed this to recognizing that the person for whom they are leaving a reference can reciprocate with a negative reference in turn.

But the big problem is that if you leave a bad reference, what happens then. What will that person say about you. You leave a bad reference and he can do the same. And its not true. [S37]

He'd written vitriolic references to anyone else who said anything negative about him and I was unaware that you could have those references removed, so I was scared of him leaving some crazy, untrue rant about me on my profile. [S421]

I've had so many (i.e. 30+) great experiences with CouchSurfing. I'm also afraid he'll leave me a negative review I don't deserve and thus scar my profile. [S256]

Reciprocity in vouches Perhaps because CouchSurfing warns that vouches are only to be used sparingly, or because only a minority of users are able to issue them (one must have three vouches in order to vouch for another), reciprocal vouches are not always expected. Exactly 1/3 of our survey respondents were able to vouch for another, and of those, 44.9% had received a vouch that they did not reciprocate, with just 8% reporting that they had felt pressured to reciprocate a vouch. Of the 34.48% of users who had been asked directly to vouch for someone, 41.3% did not, and only 14% felt uncomfortable not vouching in that situation.

Of the 74 freeform responses given by those who had at some point not reciprocated a vouch, roughly half explained that they did not know the person well enough or could not trust them, usually because they received the vouch after only a single meeting. One individual had even “once ‘rebuked’ a very nice city amb[assador] for vouching [her] at first sight...” Many reported that they have to feel ready to vouch for someone, and 8 mentioned that the system is not an exchange. Another 8 mentioned not being familiar

enough with the system or not initially being able to vouch back.

And no one ever asked me and I've never asked anyone to vouch for me. It's kind of like a... taboo thing. You hope that they do [P9]

The fact that individuals feel comfortable not reciprocating a vouch corresponds to unreciprocated vouches containing a useful signal about a lower private trust level (Teng, Lauterbach, and Adamic 2010).

Considering how ratings affect others' reputations

References and reputation As mentioned, lack of negative references can in part be attributed to a fear that the other person will reciprocate in kind. But many users felt that they simply did not want to harm another's reputation, recognizing that their experience may not be representative of how others might interact with that same person.

The few times when I had a neutral experience, I believe it was because of personal character differences, and not because I had complaints against the person in question. Somebody else might have a positive experience, why write them a negative/ neutral one and prevent people from considering the person? [S83]

21 of the 50 explanations of why a neutral or negative reference was not left, expressed that the person's behavior was not bad enough to warrant a negative reference.

Negative reference... would be only in case of the person would be not respectful or dangerous. [S668]

It may have been because we have a different vision of CouchSurfing. For me bad references are only for very bad experiences (when someone doesn't host you as promised without telling you, etc.) [S553]

Newcomers, not accustomed to CS norms, were especially likely to be cut some slack:

I didn't leave a neutral or negative reference, because it was his first CS experience and I considered the possibility that CS is just not for him [S252]

Vouches and reputation Vouches were also perceived to benefit users' reputation, especially for newcomers to CS:

... better to vouch for people [...] who are new and still need their first vouches [S366]

As is typical in point systems, a few users turn vouching into a popularity contest, but for most vouches were reserved for individuals who had earned their trust.

I didn't vouch back to a guy that vouched me just because he wanted to get some popularity and vouch(es) back - I think [S74]

Discomfort in leaving negative references

Several users reported not feeling comfortable in leaving negative references. Sometimes it was because they were contrary to references others had left:

I didn't feel comfortable to leave a negative reference, cause he had only positive references and I thought that I shouldn't be so bad with him. [S59]

One interviewee expressed remorse about not leaving a negative reference:

I chose not to leave a reference because I just felt uncomfortable... then I actually ended up speaking to the country ambassador... and she told me that several other surfers have had the same experience with this particular host, and I just chose on a personal level not to leave a reference for him. Today, I regret that. I wish I had left a negative reference so no one else would have been put in that situation that we were. [P9]

Another interviewee who reported being too scared to leave a reference, but downgraded her private trust rating to "I do not trust this person" instead:

I was almost raped and felt too scared to write something about it, but i did write to CS about it. [P5]

Finally, even if a user overcomes the aforementioned issues that might deter them in leaving a negative reference, the reference might not stick:

CS isn't a date site. The guy tried to kiss me and I gave a negative. But he had ambassador friends and they took my negative out. [S585]

The fact that some users have experiences that are negative enough to report in private, but fail to disclose them in public, points to a need to construct incentives for leaving negative references, e.g. by providing an alternate, anonymous feedback mechanism.

A workaround: hidden signals Given the reluctance of CS users to leave negative or neutral reference, one might wonder whether references can even serve the purpose of steering users away from undesirable contact. In fact, CS users seem adept at both writing and interpreting seemingly positive references that imply a neutral experience:

I either neglect to reference, or write a "positive" response but in a neutral tone. [S67]

I've gone pretty keen on what certain references mean, and you can tell a... you-were-a-nice-person-reference: "[She] was great. She was very hospitable. She's a great host." That can mean in a sense you might be kind of boring. [P9]

Conclusions and future work

Ratings of friends by friends can be useful to other users evaluating whether they should trust someone who is a stranger to them. Users trying to provide such ratings do face a series of challenges. They are often uncertain how to best translate inherently subjective and multidimensional feelings to a pre-labeled, linear scale. Individual interpretations of trust and friendship vary, and in aggregate the two concepts correlate to other characteristics of a social tie and to each other in a non-symmetric way.

Beyond wrestling with labels, the rater also tends to consider both how the ratee's reputation might be affected, and

how they might react. Together these factors help to explain patterns such as higher reciprocity in public ratings, and the near absence of negative ratings. They also help to explain why, according to our survey, a majority of users think that the number of vouches, and number and types of friendships a user has, are unimportant in deciding whether to host them. Instead, they attach much more value to textual references (47.6% rated them as “very important” and another 40.8% as “important”), which directly address the likely experience they will have (Cheshire and Cook 2004) and wherein a broader range of signals can be embedded.

One limitation of our study is that it pertains to a single community with a specific culture. For example, Tan (2010) found that for CS users trust went beyond physical safety to also include a feeling of membership in the community. Either personal safety or sense of membership may not be central to all communities. However, we believe that many online communities will have similar issues about ratings. They will not be the same cultural effects, but they will all have cultural effects that lend themselves in similar directions. An example is the discomfort users feel in publicly expressing distrust. Despite an overwhelmingly high proportion of positive references on CS, 6% of private trust ratings are “I do not trust this person” and another 12% are “I trust this person somewhat”. A rather different site, Epinions.com, focused around product ratings, has users publicly rating other users reviews as very helpful or helpful 97% of the time, consistent with CouchSurfing references. But privately Epinions users assign a distrust score to other users 15% of the time, again consistent with our findings for CS.

It is therefore useful to draw on some current design decisions in CouchSurfing that were found to offset the pressure to reciprocate and give positive ratings. Stern warnings about the meaning of vouching, combined with a system where many vouches are not given due to initial constraints, allowed a norm to form where vouches are not requested and when requested are not always given. If desired, one could provide similar instructions for other types of ratings, to help users disambiguate the terminology and expectations:

the rating system . . . can be good or bad, it depends how responsibly the members use it. . . of course you can't force people to use it more responsibly, but maybe in larger letters, with a different wording, it would be possible to achieve it [P19]

Or one may want to foster ambiguity to smooth the awkwardness between friends who rated one another differently.

To help users give more accurate feedback, one might design an alternate, anonymous feedback system. One might also explore using faceted ratings relevant to aspects of the specific community, for example, whether someone, while being a nice person, might have a tendency to overstay their welcome. In future work we would like to experiment with different rating system designs, to probe how wording and dimensionality affect the ability to quantify relationships, to measure the effect of anonymity and privacy on truthfulness of ratings, and to devise ways in which explicit online ties can more closely track their offline evolution.

Acknowledgements We thank CouchSurfing for providing the data used in this study and posting a link to the survey. We are grateful to the survey respondents and interviewees for sharing their experiences with us. This work was supported by MURI award FA9550-08-1-0265 from the Air Force Office of Scientific Research.

References

- Bialski, P., and Batorski, D. 2007. Trust networks: Analyzing the structure and function of trust. In *Sunbelt conference poster*.
- Boyd, D., and Ellison, N. 2008. Social network sites: Definition, history, and scholarship. *JCMC* 13(1):210–230.
- Caverlee, J.; Liu, L.; and Webb, S. 2008. Socialtrust: tamper-resilient trust establishment in online communities. In *JCDL '08*, 104–114.
- Cheshire, C., and Cook, K. 2004. The Emergence of trust networks under uncertainty—Implications for Internet interactions. *Analyse & Kritik* 26:220–240.
- Gilbert, E., and Karahalios, K. 2009. Predicting tie strength with social media. In *CHI'09*, 211–220. ACM New York, NY, USA.
- Golbeck, J. 2009. Trust and nuanced profile similarity in online social networks. *ACM TWeb* 3(4):1–33.
- Guha, R.; Kumar, R.; Raghavan, P.; and Tomkins, A. 2004. Propagation of trust and distrust. In *WWW2004*, 403–412.
- Jøsang, A., and Pope, S. 2005. Semantic constraints for trust transitivity. In *APCCM'05*, volume 43, 68. Australian Computer Society, Inc.
- Lauterbach, D.; Truong, H.; Shah, T.; and Adamic, L. 2009. Surfing a web of trust: Reputation and Reciprocity on CouchSurfing.com. In *IEEE SocialCom*, volume 4, 346–353.
- Leskovec, J.; Huttenlocher, D.; and Kleinberg, J. 2010. Predicting Positive and Negative Links in Online Social Networks. In *WWW2010*. ACM.
- Miles, Matthew B. and Huberman, A. M. 1994. *Qualitative Data Analysis*. Sage.
- Rempel, J.; Holmes, J.; and Zanna, M. 1985. Trust in close relationships. *Journal of personality and social psychology* 49(1):95–112.
- Robert, Jr., L.; Denis, A.; and Hung, Y.-T. 2009. Individual swift trust and knowledge-based trust in face-to-face and virtual team members. *JMIS* 26(2):241–279.
- Skopik, F.; Truong, H.; and Dustdar, S. 2009. Trust and reputation mining in professional virtual communities. In *ICWE'09*, 90.
- Tan, J. 2010. The Leap of Faith from Online to Offline: An Exploratory Study of Couchsurfing.org. *Trust and Trustworthy Computing* 367–380.
- Teng, C.-Y.; Lauterbach, D.; and Adamic, L. A. 2010. I rate you. You rate me. Should we do so publicly? In *WOSN'10*.