‘MASTerful’ Matchmaking in Service Transactions: Inferred Abilities, Needs and Interests versus Activity Histories

Hyunggu Jung[1, 2], Victoria Bellotti[2, 4], Afsaneh Doryab[3], Dean Leitersdorf[6], Jiawei Chen[5], Benjamin V. Hanrahan[5], Sooyeon Lee[5], Dan Turner[2], Anind K. Dey[3] and John M. Carroll[5]

[1]University of Washington
School of Medicine
Division of Biomedical and Health Informatics
Seattle, WA 98195 USA
hyunggu@uw.edu

[3]Carnegie Mellon University
HCI Institute
NSH 3519
5000 Forbes Ave.
Pittsburgh, PA 15213 USA
{adoryab, anind}@cs.cmu.edu

[5]The Pennsylvania State University
College of Information Sciences and Technology
University Park, PA 16802 USA
{jzc245, banrahan, sul131, jcarroll}@ist.psu.edu

[2]Palo Alto Research Center
3333 Coyote Hill Rd.
Palo Alto, CA 94304 USA
bellotti@parc.com
ddt@twoangstroms.com

[4]UC Santa Cruz
Jack Baskin School of Computer Engineering
Santa Cruz, CA 95064 USA

[6]Technion-Israel Institute of Technology
Haifa 3200003 Israel
dean.leitersdorf@gmail.com

ABSTRACT
Timebanking is a growing type of peer-to-peer service exchange, but is hampered by the effort of finding good transaction partners. We seek to reduce this effort by using a Matching Algorithm for Service Transactions (MAST). MAST matches transaction partners in terms of *similarity* of interests and *complementarity* of abilities and needs. We present an experiment involving data and participants from a real timebanking network, that evaluates the acceptability of MAST, and shows that such an algorithm can retrieve matches that are subjectively better than matches based on matching the category of people’s historical offers or requests to the category of a current transaction request.

Author Keywords
Timebanking; reciprocal recommenders; matching algorithms; experimental evaluation.

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g. HCI): Miscellaneous.

INTRODUCTION
Local peer-to-peer marketplaces, particularly timebanks, that use alternative currencies, are part of a growing phenomenon [43] that promises to help communities become more deeply interconnected and robust [3,8,31,33,41,42,43]. This is especially true of timebanks, which focus on person-to-person (P2P) *service* transactions.

CHI'16, May 07 - 12, 2016, San Jose, CA, USA

ACM 978-1-4503-3362-7/16/05…$15.00
DOI: http://dx.doi.org/10.1145/2858036.2858263

Figure 1. An offer from the hOurworld timebank and a request from the Time Republik timebank. One is categorized as “Classes…” and the other as “Programmer.”

Timebank members earn and spend an alternative currency (time dollars) in service transactions. Any member can earn a time dollar for an hour’s work, such as mowing another’s lawn, and spend it on any service from any other member. In this way, people provide support to each other through the provision of services, which low-income members in particular might otherwise not have access to, with opportunities to develop skills and self-respect [8].

Timebanks must ensure that members find good transaction opportunities and respond promptly to posted requests and offers; failure to find matches is a significant demotivator for members [45]. We seek to reduce this problem by more actively promoting possible transactions via an automatic transaction-partner matching algorithm, which we present along with its experimental evaluation in this paper.

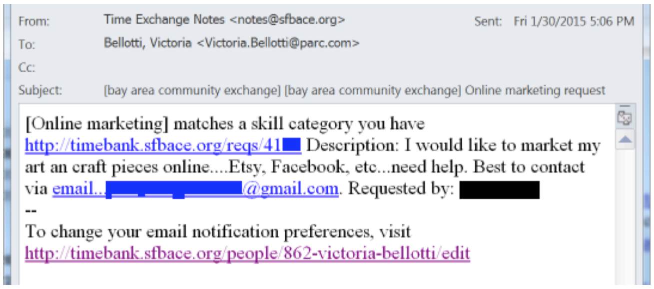

From: Time Exchange Notes <notes@sfbace.org> Sent: Fri 1/30/2015 5:06 PM
To: Bellotti, Victoria <Victoria.Bellotti@parc.com>
Cc:
Subject: [bay area community exchange] [bay area community exchange] Online marketing request

[Online marketing] matches a skill category you have http://timebank.sfbace.org/reqs/41[redacted] Description: I would like to market my art an craft pieces online....Etsy, Facebook, etc...need help. Best to contact via email..[redacted]@gmail.com. Requested by: [redacted]
--
To change your email notification preferences, visit http://timebank.sfbace.org/people/862-victoria-bellotti/edit

Figure 2. An email from BACE that alerts a member to a service match

Currently, timebanking technology is limited in facilitating transactions. Members must spend time *browsing* a large hierarchy of categories, (e.g. *Classes/Lessons/Tutor: Outdoor Activities/Sports)* or tags (e.g. *Online Marketing* (see Figure 1; hOurworld, a large timebanking network, has 78 categories, each with sub-categories). Members can also *search*, but it takes time to think up viable offers or requests. Bay Area Community Exchange (BACE) timebank supports keyword matching to a member's offered services (though not to a user's needs), so when a user classifies a request, it is *pushed* via email to members with profiles whose service categories match (see Figure 2). But members never get matching suggestions of people to transact with at the time they are posting offers or requests, which is when they are most motivated to reach out to contact others to propose a transaction.

We seek to increase transaction rates with *context-aware transaction facilitation* in timebanking and other P2P services [3], in contrast to context-aware *social* matching [27]. One of our methods is to find providers and receivers who have compatible (i) shared interests (defined as transaction categories mentioned in any part of a timebanker's profile), and (ii) complementary needs (inferred from categories of requests previously posted) and abilities (from categories of offers previously posted).

In the rest of this paper, we review related work on P2P systems, community-oriented social media and matching (recommender) systems. We then describe our MAST system and its evaluation. Finally, we discuss implications of our finding for timebanking and P2P services in general and conclude with limitations and future work.

RELATED WORK AND BACKGROUND FOR OUR STUDY

Timebanking

Ideas similar to timebanking have been around since Josiah Warren established the 'Cincinnati Time Store' in 1827. In the 1980's Edgar Cahn [7,8] coined the term 'time banking' and created the TimeBanks USA timebank network. There are now many networks in existence (e.g. hOurworld, TimeRepublik, LinkAges, Community Exchange).

Timebanks provide a valuable public service that alleviates burdens on the state and the need for professionals in services such as transportation, deliveries, help-in-the-home, tutoring, and so on. The key ethos of timebanking is that every person's time is worth the same amount [8,12]. Another key feature of timebanks is that most are engaged in local community building. When Seyfang et al. [40] surveyed the motivations of members of a timebank, their top two responses were "to help others" and "to get more involved in the community."

Being part of a community has many benefits for members [5]. Given timebanks' focus on community, these benefits fall to members as well. In addition to providing needed services, timebanks crucially provide opportunities to make meaningful service contributions, which promotes physical health and a sense of belonging [23]. These benefits particularly impact older and lower-income persons. In fact, rapid growth in timebanks was seen in both Greece [15] and Spain [30] during their economic crises.

And timebanking is on the increase [43]. However, when compared to similar more commercial systems in the sharing economy [16], the growth is modest. One reason is limited technical support [20], which could be improved upon [2]. But there are also barriers around understanding trade mechanisms and how these contribute value to the community [33,42] and some users do not understand that asking for help gives other users the benefits of earning hours and feeling valued [2,33,42]. Also the promotion and coordination of services takes considerable effort from organizers who often get burnt out [2,11,45]. So technology that can increase the rate of successful transactions and reduce burdens on overworked organizers is greatly needed.

Community Building and Reciprocity

Communities where people help one another promise a number of wellness benefits. Unprompted altruism seems to contribute to happiness [14]. Providing social support to one's community can even reduce mortality [5]. Seligman [39] likewise observed that gratification is obtained by striving for noble goals by applying our personal strengths.

So one of our main aims is to connect people who have similar interests and are likely to form nurturing social ties. There is little benefit to simply having virtuous individuals around if they are isolated [36]. But as well as finding good matches for a service transaction we also want to connect people with *complementary* abilities and needs to provide *multiple* opportunities for transactions and *return* of favors. Even though providing favors has many benefits, people are loath to only accept services for fear of seeming needy [2,46]. The *reciprocity rule* [48] dictates that individuals feel compelled to return favors *to the person who dealt them*, even if they do not particularly care for that person [37]. Although both Seyfang and Ozanne argue that timebanks help to reduce the stigma of asking for help [33,42], it appears that they may not go far enough [2].

We seek to reduce the stigma of asking for help by framing that experience as a successful match between people who *want to transact **on both sides** of the service equation*. What is novel in our work is that, if we match members in a *complementary* manner, the parties to the match can be and ideally will be *both providers and receivers* over time, as timebanks encourage. Thus, our matching approach is intended to help to forge valuable social and helping relationships with *high potential for future reciprocity*.

Matching and Recommenders

Person-to-person matching technology per se is not novel. There has been considerable work on how to recommend people, tasks, or jobs, albeit in different circumstances. Terveen and McDonald [47] outline special considerations for matching people, not least of which is that, by definition, some personal information has to be disclosed. For brevity's sake we identify just a few *applications* and *approaches* to person-to-person and person-to-role matching to compare with our own—***reciprocal service-oriented P2P matching***.

Dating

The most obvious application for *person-to-person matching* is online dating where profiles are matched for relationship compatibility. Pizzato et al. [34] derived *implicit preferences* from the interactions that people had with each other. Those preferences were then matched with the profiles of other members. In a similar approach, Diaz et al. [13] had members identify their *preferences* in an ideal partner, which were used to match and rank others. If people who were matched exchanged messages, the *match was used as positive reinforcement* for the algorithm.

Person-to-Person Matching (Social Matching or Reciprocal Recommenders)

IBM's enterprise social networking service, Beehive [10] allows users to connect to friends and co-workers, post new information or comment on shared information, using content-based and collaborative filtering algorithms. This means using similarity of the content posted by users and other contextual information to match people.

The recommender for a social networking website proposed by Han et al. [21] uses two-way interaction for exchanging messages between users, where interests of both sender and receiver are considered. The method extracts interests from both user profiles and interactions and then uses a weighted harmonic mean to make recommendations.

Reciprocal recommender systems also match people [24,34]. Different from the traditional content-based recommendation, these systems provide recommendations by considering the preferences of both parties involved.

Practical Matching (Talent- and Expertise-Matching)

Reciprocal recommenders [24,34] can also be used for *talent-matching*. For instance, in an online recruiting system, a job seeker would search jobs that match his/her preferences, e.g. special skills and salary; and a recruiter might seek suitable candidates to fulfill job requirements. Other illustrative examples of reciprocal recommenders include online mentoring systems, customer-to-customer marketplaces, as well as online dating services.

Another approach for matching people to jobs [26] considers *preferences* of both the job *seeker* and *recruiter*. Two separate recommender systems used expectation maximization to build prediction models: one to match job seekers to jobs and the other to match jobs to candidates. To create a model for matching job seekers to jobs, the recruiter manually labels a profile or resume as either fitting or not fitting different job descriptions. The features considered included demographics, education, experience, skills and language. To get job recommendations, candidates were asked to rank a set of job descriptions indicating how well each job fit their preferences.

Expertise-seeking research is concerned with how people search for expertise and choose whom to contact for a specific task [38,50]. Researchers have focused on developing content-based algorithms similar to document search. These algorithms identify experts primarily based on the text of documents associated with those experts. Fazel-Zarandi [19] studied social drivers including the level of expertise, homophily and social exchange to predict scientific collaboration. They found the combination of these different drivers better-predicted collaborators than network structure. Alves [1] aimed to support collaboration among researchers with common interests by analyzing their profiles and hierarchical relationship history with other researchers in their network.

The i-Help system [6] was designed for *mentor-mentee matching* to help students with courses. It matched mentors with mentees by considering their attributes and preferences, using information from different sources, including self-evaluation and peer feedback in previous help sessions, to infer attributes of mentors such as interests, helpfulness, availability, and knowledge of various topics. The ranked list of potential mentors was refined by considering preferences of the mentees, e.g. the importance of a mentor's helpfulness and availability.

To sum up, previous approaches are tailored to the contexts they were designed for and focus mainly on analyzing profiles for matching text or entities and, in the case of service-provision, assuming pre-defined roles of provider and receiver. In our case, each person can be a *provider* ***or*** a *receiver* in any transaction, depending on the type of post they create or respond; indeed, *this* ***role switching*** *is the very point of timebanking.* This motivates our customized MAST approach; *reciprocal service-oriented P2P matching*, which presumes (in line with the reciprocity rule [48]) that a timebanker, A, will prefer to receive a service from another timebanker, B, who happens to be in need of a reciprocal service that A is able to provide (i.e. has complementarity with A), rather than timebanker C who needs nothing that A can provide.

OUR TECHNICAL APPROACH

Before we describe our technical approach, we note that our algorithm is not the only way to do transaction partner matching and better algorithms may soon be developed.

Our aim is simply to determine initially whether *any* reasonable matching algorithm can improve on a baseline of matching based on timebank members' indicated offer and request categories. This baseline is a two-way version of what the timebank BACE does, as mentioned previously, which is the state-of-the-art in timebanking at present.

The Matching Algorithm for Service Transactions (MAST) that we designed for our experiment is summarized in Figure 3 and relies on matching via the following:

1. Self-description (similar to prior work, Figure 3, top) and classified interest categories (Figure 3, middle).
2. Extracted abilities and needs (Figure 3, near bottom).
3. Explicit offers and requests (Figure 3, bottom).

Items 2 and 3 above are possible due to the special nature of timebanking where both abilities and needs are important and can be inferred from the category people choose for each of their offers and requests, as we shall explain.

We generate matches between timebankers, based on a ***user's profile***. Critical to our discussion, this consists of:

- ***A bio*** containing a *free-text field for self-description*
- ***Past service offers and requests***, each including a *free-text field* and, critically, a *category label*.

A profile is matched with profiles of others who need or can provide a *matching service*. Match quality depends on *similarity* of interests (i.e. homophily, which inspires liking [32]) and *complementarity* of needs, and abilities (which make people more practically useful to each other).

We want to infer common interests from timebank users' profiles for similarity, and service matches between profiles for complementarity. Since each profile includes offers and requests, we consider offer categories as skills and request categories as needs. However, as interests of users are not explicit entities in the database, we must classify text in bios, offers, and requests to generate them. We do this by compiling all the category labels of the offers and requests and removing repeated labels to produce a set of categories that we deem to be 'interests.' Next, we extract interests from unstructured text in bios by performing the following five steps: 1, pre-process unstructured text from user profiles to anonymize the data, 2, extract textual features from all offers and requests to characterize categories textually, 3, train and build a classifier of categories from the offers and requests dataset, 4, evaluate classifier performance, and 5, apply the classifier to the unstructured text in user bios to derive a richer set of interests (i.e. categories). Finally, the MAST algorithm aggregates the interests from bios, offers, and requests for each user and generates matched profiles.

Data Pre-processing

Our dataset is 3,943 real hOurworld user profiles (bio, offers, requests and some metadata) and their historical transactions (i.e. 10,615 offers and requests). The data was obtained as a data dump from hOurworld where each profile was only indexed by its member ID number

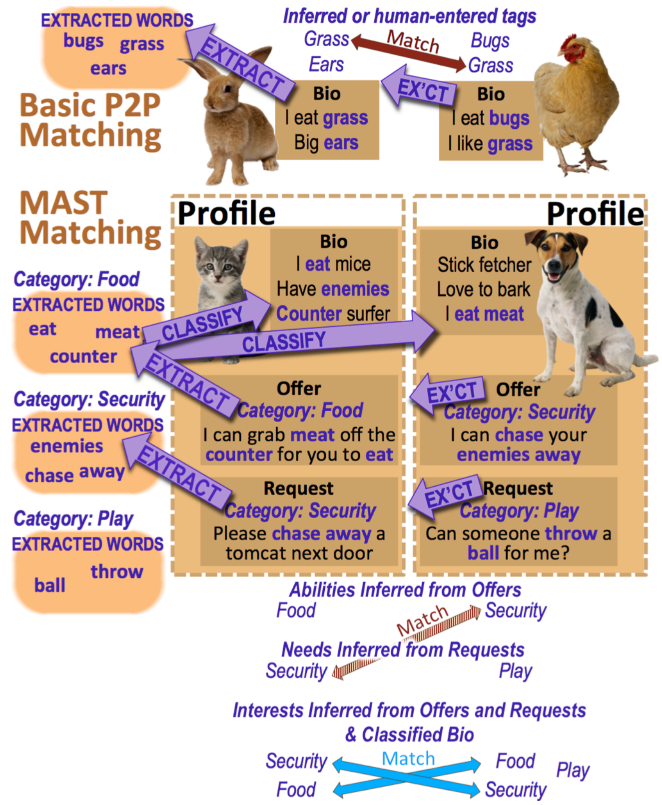

Figure 3. Simplified basic reciprocal recommendation above, between Rabbit and Chicken. Our MAST transaction matching below, between Cat and Dog. Purple arrows are 'extract' (from text) and 'classify' (with category). Dark red arrow is a 'match' on bio text. Striped/pale red arrow is a 'match' on inferred complementary abilities and needs. Blue arrows are a 'match' on explicit posted offer or request.

(member names were removed). We cleaned the dataset to remove any identifying information in member bios, offers and requests (e.g. names, phone numbers, email addresses and locations). The data was then stored in our database.

We needed ground truth category labels from offers and requests for feature extraction and to train our service category classifier that extracts category labels from text. Recall that each offer or request in hOurworld includes an unstructured textual description and a *category label* (e.g. *Help in the Home*) chosen by the user at the time of posting.

Now, we describe in more depth how we extracted features for our classifier, and then how we trained and validated it.

Extracting Textual Features from Offers and Requests

Using the LightSIDE text-mining tool [28,29], we extracted unigram (single) words from text in offers and requests and filtered out stop words (e.g. the, is, that) and other irrelevant elements like numbers and punctuation. We excluded unigram words that occur fewer than 5 times across all categories in the entire training set to remove uncommon/rare words and reduce the size of the feature set. This threshold was empirically derived to maximize accuracy of the trained classifier. The feature extraction step generated a feature table that was used to train this classifier (described next) to discover latent patterns in those features.

Building a Classifier to Categorize Member Bios

As mentioned, users choose a category for each timebank offer and request. In our supervised learning approach, we used these labels as the ground truth to train our classification algorithm and to build a model that is later used to classify member bios into interest categories.

We used logistic regression (also known as maximum entropy or log-linear) integrated with regularization (i.e. introducing additional information to avoid overfitting). This machine learning approach optimized the feature set to use the fewest possible features for prediction by dropping the weighted input of as many features as possible to zero.

The feature table generated in our feature extraction step is used to train the customized logistic regression algorithm on our (labeled) offer and request dataset. The number of extracted features/words is large (9,058) for the number of training instances – 10,989 sentences of offer and request descriptions. Despite the latter being a larger number, it is small relative to the ideal for training. This resulted in low classification performance, and therefore, we increased the size of the training set by using the technique of three-times sampling with replacement [18]. This bootstrapping helped the learning algorithm find discriminative features for each category to improve the classification accuracy. Later, when text in a member's bio is provided to the classifier, it outputs the probability of the text being placed into each of the 78 categories. We choose the categories with the highest probabilities as interest categories for that particular member.

To evaluate the performance of our classifier and before applying it to classify member bios into interest categories, we used 10-fold cross-validation on the 10,615 offers and requests dataset, which included all extracted features. We built 10 models, each on 90% of the data and tested on the remaining 10%. The performance of the classifier is the average accuracy and the average improvement over chance (measured by Kappa) of the 10 models. Our classifier obtained an overall accuracy of 90.3% and an overall Kappa of 89.9%. This compares with a baseline chance of guessing correctly of 7.9% by choosing the most popular category (Classes/Lessons/Tutor) every time. Note that, the training set used to build and evaluate the classifier contained only text from offers and requests. Member bios, the targets of our classifier, were not included in the training and model-building phase.

Classification of Member Bios into Interest Categories

We then applied the built classifier to the dataset of member bios to categorize each bio into service/interest categories. Each member's bio was divided into individual sentences for classification and the classifier output the probability of that sentence belonging to each of the 78 categories. The probabilistic categories for each sentence were filtered to only include those that had a probability greater than 65% (we arrived at this threshold by iteratively reviewing how well the extracted interests matched what human readers would agree with). If there were duplicate categories in an entire bio, we only counted them once. In the rest of this paper, we refer to the aggregated set of service categories of a member's offers and requests, and those inferred from the member's bio as the member's ***interests***.

The following section describes how the inferred interests along with users' past history of offers and requests are used to match people according to their needs and skills.

Matching Algorithm

Our MAST algorithm takes *inferred interests* from timebank user bios, and user-indicated *explicit interests* from the categories of their offers and requests in the hOurworld database as input and produces matched profiles as output. To determine matched profiles, for each user, MAST recommends users with the highest calculated match scores. Intuitively, the profiles that have the highest matching scores have the most interests in common. A match score is calculated by a linear combination of levels of similarity and complementarity as follows:

$$score = level\ of\ similarity + level\ of\ complementarity$$

Level of similarity is indicated by the number of inferred interests that two users have in common. Mathematically, we represent this as the ratio of the number of common interests between two users to the total number of interests of the user being matched. Likewise, level of complementarity is indicated by the ratio of the number of service matches (when a category of a request matches an offer or vice-versa) to the total number of offers and requests for the user being matched. So our algorithm combines inferred and explicit interests into measures of similarity and complementarity. After human subjective review of the goodness of matches, similarity was weighted by multiplying by 3 (complementarity was unweighted) in the final match score for best results (higher score is better).

We now present our MAST algorithm more formally:

$Input: \{Profile_1, \ldots, Profile_n\}$
$Output: \{MatchedProfile_1, \ldots, MatchedProfile_{n-1}\}$

Let the profile of a user be $Profile_u$ and *n* be the number of profiles in our database.

$O(u) \leftarrow Interests\ from\ category\ in\ offers\ of\ Profile_u$
$R(u) \leftarrow Interests\ from\ category\ in\ requests\ of\ Profile_u$
$B(u) \leftarrow Interests\ from\ text\ in\ bio\ of\ Profile_u$
$I(u) \leftarrow O(u) \cup R(u) \cup B(u)$

for [each $Profile_k$ in $\{Profile_1, \ldots, Profile_n\} \setminus \{Profile_u\}$,
$O(k) \leftarrow Interests\ from\ category\ in\ offers\ of\ Profile_k$
$R(k) \leftarrow Interests\ from\ category\ in\ requests\ of\ Profile_k$
$B(k) \leftarrow Interests\ from\ text\ in\ bio\ of\ Profile_k$
$I(k) \leftarrow O(k) \cup R(k) \cup B(k)$

$$score_{common\_interests}(k) \leftarrow \frac{|I(u) \cap I(k)|}{|I(u)|}$$

$$score_{service\_matches}(k) \leftarrow \frac{|O(u) \cap R(k)| + |R(u) \cap O(k)|}{|O(u)| + |R(u)|}$$

$$score_{final}(k) \leftarrow \alpha \times score_{common\_interests}(k) + \beta \times score_{service\_matches}(k)$$

$$ScoredProfile_k \leftarrow (Profile_k, score_{final}(k))$$

] end for

Sort the $ScoredProfile\ 1$ through n in decreasing order.

$$MatchedProfile_1, \ldots, MatchedProfile_{n-1} \leftarrow Sorted\ ScoredProfiles$$

We have already provided a high-level schema for how this algorithm works in Figure 3. Essentially, this algorithm is developed for *reciprocal service-oriented P2P matching* and combines inferred interests with inferred abilities and needs from offers and requests to enable *complementarity* as well as *similarity* matching.

EVALUATION STUDY

Study Overview

While a real-world deployment of MAST in a timebanking service is our ultimate goal, here we simply evaluate whether MAST is more effective than just matching based on historical offers and requests, which is a commonsense approach that aligns with people's needs and abilities. So we conducted a web-based evaluation of the quality of MAST's matches with timebankers as participants.

Participants and Recruitment

The participants in our study were anonymous hOurworld timebankers. They were recruited via email from a subset of the hOurworld timebank network (i.e. the ten largest exchanges). Emails contained a link to a pre-evaluation MAST acceptability survey. The only prerequisite for participation was that they have an active hOurworld account and an email address so that we could send the survey link. They were compensated with half a time-dollar, equivalent to half an hour of effort in the timebank. Timebank membership tends to be skewed towards older, better educated, but less well-off, females [22,23,44,49] so we also assume that our participant pool is similarly skewed, although we did not collect such data.

MAST Acceptability Survey

In our MAST acceptability survey, we asked volunteers to rate agreement with three statements about their willingness to use profile analysis technology that would match them to other timebankers and collected the volunteer's email address, which we used to contact them later with a link to the study web page. We also anticipated that some individuals' hOurworld profiles would not be filled in and this survey gave volunteers an opportunity to enter profile bio information that they would like us to use for finding good matches for them.

MAST Evaluation Method

Our MAST algorithm evaluation method was implemented via the following steps for each evaluation volunteer:

Pre-evaluation Set-up

1. Apply our technical approach, as described previously.
2. Obtain volunteer email address from MAST acceptability survey and assign her to the control condition or to the MAST matching condition (to be explained).
3. Use email address to find her account, either in the data dump or if a volunteer was not a member of one of the ten largest timebanks, by query from hOurworld through an API that was previously built.
4. Save the ten most recent offers and ten most recent requests (or as many up to ten of each as they have).
5. If no bio or offers or requests were obtained from volunteer accounts and nothing was provided in the MAST acceptability survey, we asked volunteers for this information via email (with limited success).
6. Create a set of ten offers made up of a set of examples taken from each of the ten most popular categories but replaced by up to ten of the user's own (depending on how many we were able to get). Repeat for requests.
7. Run our matching algorithm on all ten offers and ten requests and the volunteer's profile and generate two transaction partner profile matches for each one.
8. Add the volunteer's own bio and the interests we extracted from it with the sets of ten offers and ten requests (step 6), each with their two matches inferred for that specific volunteer (step 7) to a row in a csv file.
9. Repeat for each volunteer.
10. Create a Qualtrics survey that allows the user to rate their own extracted interests, pick two offers and two requests (i.e. items stored in rows from the csv file), and to rate the matches for those offers and requests.

Evaluation Procedure

1. Volunteer logs into the survey with email address.
2. Qualtrics identifies the corresponding entry in the csv file and volunteer is presented with his own profile and asked to respond using a Likert scale (1=Strongly Disagree to 7=Strongly Agree) to the following request; *"Please rate how much you agree with this statement: These interests match my real interests."*
3. Volunteer is then asked to pick two offers and two requests that are most similar to something he would post (he can choose from the sets of ten offers and ten requests, each set including up to three of his own if he posted any).
4. Volunteer is then asked to rate each offer and each request in terms of how typical it is. This was done by indicating (using the same Likert scale) agreement with the following statement: *"Please Rate Your Offer Selections in terms of how typical they are of offers you have posted or might post."*
5. Volunteer is then presented with each offer and each request he selected in sequence along with two matched profiles (previously compiled in the csv file).
6. Volunteer rates each profile (same Likert scale) to indicate agreement with the following statements:

i. *This profile is a good match for me to transact with in general.*
ii. *This profile is a good match for this offer/request.*
iii. *This person's interests are similar to my own.*
iv. *This person's abilities and needs are complementary to my own.*

7. Volunteer is thanked and closes the survey window.

Conditions

There were two conditions in our experiment, although each one looked exactly the same to participants.

CONTROL—Basic Category Matching: Profiles are randomly matched to each offer or request, based only on whether or not they included a past request in the offer's category or an offer in the request's category. For example, if a profile in our database contains an offer classified as Garden/Yard Work, that profile is a candidate for being randomly matched to a request tagged with the same category. This baseline or control for our experiment resembles the BACE timebank approach mentioned earlier.

MAST—Smart Matching: Profile matches are made by prioritizing the profiles that best match the participant's inferred *interests*, *abilities* and *needs as well as* matching the offer or request category as in the control condition.

Running the Study and Data Collection

The MAST acceptability survey was released a month before experimental data collection and 121 individuals responded to the three receptivity questions. Of those individuals, only 104 gave email addresses indicating that 17 individuals chose not to participate in the study, but still provided input to the receptivity questions (see Figure 4).

Of the remainder, a number had to be disqualified from the study because they had no information in their hOurworld profile or gave us an incorrect email address that could not be resolved by the start of the data collection period. In total, 89 participants gave sufficient information to participate and were randomly assigned to each of the two conditions and then sent the link for the web-survey-based study. Participants were not informed of which condition they were in, nor were they told that there was a control condition with no matching algorithm.

Of the 89 who received the study link, 20 did not respond, and 63 completed the study (32 in the control- and 31 in the MAST matching algorithm condition; 6 additional individuals were dropped for not rating any matches). Participants were included, if they rated 6 or more matches out of the total of 8. One participant in the control condition skipped rating two profiles, two participants in the MAST condition skipped rating two profiles and one participant in the MAST condition skipped rating one profile (there are more skips in the MAST condition due to a bug, which was intercepted and fixed soon after survey release). Thus, 32 participants in the control condition rated 122 matches and 31 participants in the MAST condition rated 115 matches (237 in total) between them.

We collected all participant scale ratings and the associated profiles (bios, offers and requests of the participant and the matches) as well as relevant selected offer and request text.

ANALYSIS

Our data analysis has been mainly quantitative, including some descriptive statistics and t-tests for simple differences in scales i-iv between conditions. To dig deeper we ran a more sophisticated mixed-effects linear regression with nested variables (with eight observations per participant and two observations per each offer and each request for all profile match rating scales; i-iv). The regression analysis was carried out to determine whether there were influences on the ratings in scales i-iv due to other variables we were able to collect (see Results and Commentary Section).

In addition to our quantitative data, we also collected some qualitative data on attitudes to transaction partner matching in our MAST acceptability survey. These were simply collated for review and some example quotes are presented below.

RESULTS AND COMMENTARY

MAST Acceptability Survey

Our MAST acceptability survey included three receptivity questions to determine acceptability of our matching technology to timebankers. On a 7-point Likert scale, most of the 121 timebankers indicated support for profile matching technology (see Figure 4). However, a minority, about 10%, did not see the advantages of matching and about 20% were concerned about the matching technology. We invited open-ended comments in the survey and the bulk of these were favorable, explicitly stating that privacy was not a concern or that they trusted hOurworld to take care of it for them or that they kept private information entirely out of their profile and offers and requests. Some examples of the favorable comments are:

"As long as my privacy is protected, I have no objection. I take HIPAA training at least once a year and know how important it is to maintain privacy."

"I'd like to have matching made MUCH easier AND I put nothing on my profile that is a security problem, and would advise others to do the same."

"I'm not trying to hide anything. I post information to facilitate trades. Some people find it, but I suspect that there are people who don't see it but might be interested. If the software is helping people find information I have deliberately posted, I have no objection."

However, a substantial minority of the comments voiced concerns about privacy. One example is this:

"I generally have some caution because I am a professional psychotherapist in the community, so participating at all in these online communities has a boundary risk for me."

In this case there is a serious professional concern, as this person's services are, by nature, confidential.

Another respondent mentioned the loss of opportunities to meet diverse types of people with filtering for similarity:

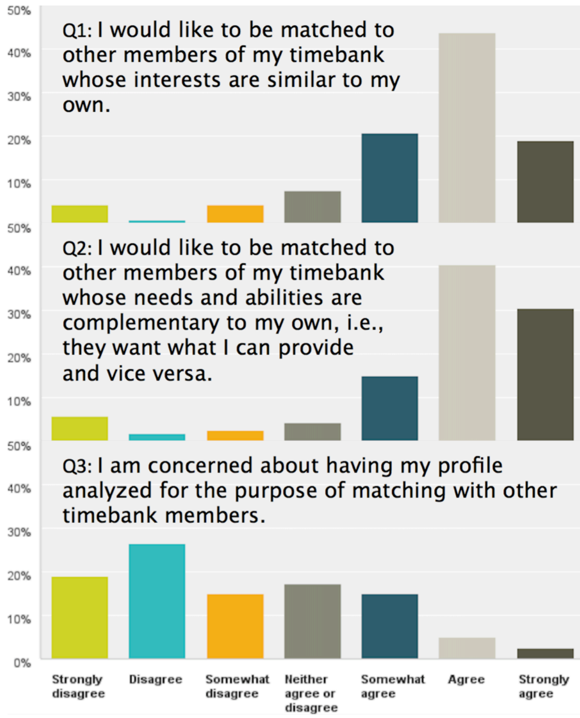

Figure 4. Responses to survey questions about acceptability of matching personal profiles for service transactions. From 'Strongly disagree' (left) to 'Strongly agree' (right). N=121.

"I am a private person. Even though I'd like to meet compatible people, there's a way I'd prefer my matches just to be based around the tasks required. I feel like I'm living in a world where corporations are constantly trying to extract my "consumer profile" to figure out what to offer me to buy. Even though this is not exactly the same, it feels similar. I don't necessarily want things to be so fine-tuned by an outside source. I'd prefer the connections through timebanking have a little bit of a random aspect. This broadens my world."

Yet another offered a subtle, but interesting point about the criteria that might be used to make matches being limited. This is true and worth reflecting upon:

"My guess is that people will look for political correctness, this is [PlaceName]; or focus on overt characteristics; and miss the more interesting aspects of what we have to offer."

These two latter quotes appeal for a more human approach to sizing up compatibility that could be lost with matching technology and we return to this in our discussion.

These kinds of articulately expressed concerns are both valid and useful. At a minimum it is clear that timebank users should not have a matching algorithm turned on automatically and that an opt-in procedure is advisable.

Matching Algorithm Evaluation

We now focus on the evaluation of the MAST algorithm itself. The analysis presented in this section makes use of the following *measures of interest:*

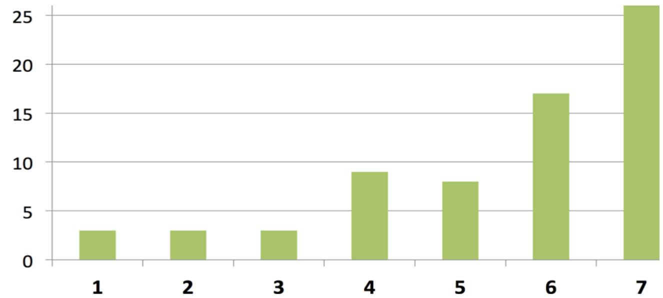

Figure 5. MAST condition only, participant ratings of inferred interests on a scale of 1-7 'strongly disagree' to 'strongly agree' with the statement "These interests match my real interests."

A. Participant rating of interests inferred from their own profile (see Figure 5).
B. Offer and Request ratings by participant in terms of their representativeness for that participant.
C. Ratings by participants of the matches they received for the offers and requests they selected in terms of the scales i-iv itemized in the Evaluation Procedure subsection.

Rating One's Own Inferred Interests (Measure A)

Figure 5 above shows that on the whole, the MAST algorithm was reasonably successful at inferring participant interests. However, in some cases it failed, so we clearly have room for improvement in our interest inferencing.

Rating the Representativeness of Offers and Requests Used in the Experiment (Measure B)

Our participants rated the typicality of offers and requests—that they selected in the study to obtain matches to—as being on average 5.7 on a scale of 1=extremely untypical to 7=extremely typical (sd=1.6).

MAST Matches Compared with Control Matches, Based on History of Offers and Requests (Measure C)

To evaluate the performance of MAST (matching on inferred interests from text in the bio, offers and requests) compared to the control condition (matching on historical offer and request category alone), we performed (stricter two-tailed) t-tests for each of the scales i-iv. The results are summarized in Figure 6 and Table 1 with Bonferroni corrections for having performed 4 tests (which increased the probability of a Type 1 error or false positive).

A prior power analysis indicated that we needed about 30 participants per condition and we barely made this threshold. This means that statistical significance depended on a reasonably sized effect. As shown in Figure 6, MAST outperformed the control by one point on 3 out of four median (black lines) and all mean results (red '+' signs) were better by at least about half a point with MAST. On similarity in particular, 75% of volunteers in the MAST condition rated matches as above or equal to the median in the control group. The smallest advantage for MAST is in the complementarity measure and this is consistent with the fact that MAST relied heavily on past offers and requests for this characteristic as did the control condition, although

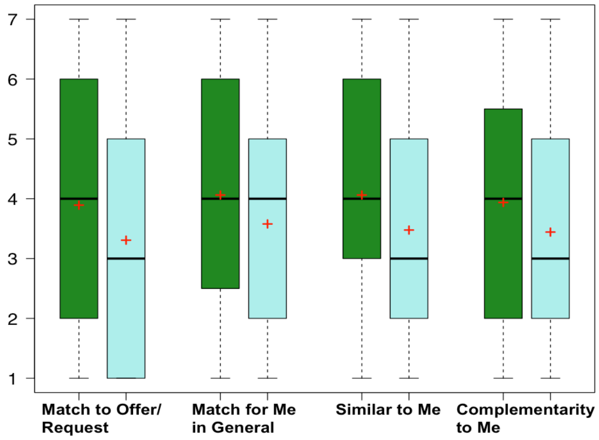

Figure 6. Box and whisker plot for participant rating of matches in MAST condition (dark green) and control condition (light blue) on scale of 1=strongly disagree with match and 7=strongly agree with match N=63. Median indicated with black lines and mean with a red "+".

Scale of Agreement	Match to Offer/Request	Match for Me in General	Similar to Me	Complementarity to Me
Average MAST	3.89	4.06	4.06	3.94
Average Control	3.31	3.58	3.48	3.44
Standard Deviation MAST	1.23	1.14	1.11	1.16
Standard Deviation Control	0.87	0.94	0.96	0.95
t-value	2.28	1.88	2.26	1.89
p-value	0.01	0.03	0.01	0.03
Bonferroni corrected p-value	0.05	0.13	0.05	0.13

Table 1. Results of two-way (stricter) t-tests for differences of averages on four scales with Bonferroni-corrected significance levels (df=61 in all cases).

matching on interests too meant that MAST-inferred complementarity was likely to be better tuned.

Our mixed-effects linear regression with nested variables did not yield any significant results. That is to say that no variable such as the number of words in a bio or in offers and requests was correlated with the four rating scales, i-iv.

GENERAL DISCUSSION

Most participants in our experiment felt that MAST did a reasonable job of inferring their interests and that the offers and requests that they were required to use in the experiment, were, for the most part, fairly typical of real offers and requests they might post. From their ratings of the matches they received, MAST performed better than matching on historical offers and requests and statistically significantly better on the scales of goodness of match to the offer or request, and similarity of matched profiles to self. Our small experimental groups, combined with two-way (stricter) t-tests and Bonferroni corrections, make the significance more compelling as it relies on effect size, rather than number of participants (more participants could lead to statistical significance for even a minute effect).

So from a technical standpoint, our effort has been quite successful. Our evaluation of MAST demonstrates that a matching algorithm can help people find promising transaction partners easily and can do so more effectively than just matching on a person's history of offers and requests. But from a social standpoint, our work raises some significant issues to address on the way to a fully working matching system. We address these briefly here.

Interpersonal Relationships: It is worth emphasizing the technical distinction between our recommendations and the state-of-the-art in timebanks today. This is that, rather than recommending transaction requests, as does BACE, we are recommending *people* (using and presenting their entire profile as a result). Someone who creates a new offer or request sees many details about others (bio, interests, past offers and requests) with whom they have things in common and with whom they can be both a provider and receiver of services. According to the reciprocity rule [48] these will be the most satisfactory people to transact with because kindnesses can always be repaid with kindnesses rather than just time dollars. However, this benefit comes with an inherent downside for timebanks, which is that, if people find ideal similar and complementary matched transaction partners, they are likely to depend on the timebank less for the service exchanges those partners can engage in with them (they will simply reach out to their new connections via other means of communication). This could reduce visible activity and evidence of a timebank's value [2] as people would develop independent mutually helpful relationships with those individuals and no longer go through the timebank to exchange those services. But we have been studying the issue of dropping out of P2P marketplaces recently [4] and our work suggests that the convenience and value of easy and well-targeted matches will continue to attract users when they are looking for new kinds of transactions. Thus, it is important to improve our matching capabilities so that users are always delighted with the results when they post new offers and requests.

Human Factors: As a quote that we cited implied, MAST matches are made on limited criteria relative to rich human inferences that can be made from reading profiles. To mitigate this problem, there should always be multiple results (probably many more than two) presented to allow users to make the final choice. This approach would compensate for any imperfections in an algorithm's matching capabilities. And there is no practical limit to the number of results we return. We might simply rank them in terms of our algorithm's match score. Users could then apply other more diverse considerations when picking a transaction partner, such as bio text that cannot be classified in terms of the algorithm's service-related categories, for example, that a person mentions that she loves reading biographies or going on long walks.

Serendipity: As one or two of our survey-takers implied, matching should be optional. Or it could be blended with a 'roulette'-style feature to maintain the ability to meet many 'random' different types of people when timebanking. Our algorithm might also 'remember' previous recommended profiles and downgrade them in subsequent match results for a while. This would also avoid repeatedly recommending the same people over and over.

Critical Mass: In a small timebank, MAST might not be as helpful as in a large one, since many active members will already know each other and others' interests, needs and abilities. However, 68% of hOurworld's members are in timebanks of 150 or more members; a number that begins to exceed people's ability to keep track of everyone [16] (especially as most will have many acquaintances outside of their timebank). There also exists a very large non-location-centered service, TimeRepublik, with 30,000 members who can all transact with each other. People in these timebanks would be most likely to benefit from a matching capability. However, our intention is also to *increase* the activity in timebanks by reducing problems such as non-responses to posted offers and requests that tend to demotivate participation [45]. So, in our approach, a person is recommended to people they are compatible with and with whom can communicate when they post an offer or request. We hypothesize that personalized communications with individuals with whom one has much in common are likely to be more successful at prompting responses than timebank members receiving impersonal transaction requests with no information about requestors and no guarantee of having anything else in common. If our work is successful in provoking more transactions, this would make timebanking more rewarding and tend to attract more users through word-of-mouth, so increasing the size of smaller timebanks.

To conclude, matched transaction partner recommendations represent an appealing idea with transformative potential for the efficiency of finding opportunities to transact in a timebank or any peer-to-peer marketplace. However, they raise some serious questions about social networking and community building that must be addressed in any practical application.

Limitations and Future Work

This study took place in an experimental setting and it remains to be seen how well matching works in a real timebanking context. We need to evaluate the effects it has at the individual level in terms of satisfaction with real transaction experiences and subsequent relationship formation, as well as how much of an impact MAST can have on timebanking activity overall and on community building. These measures of success are beyond the scope of the present evaluation, but will be a focus in future work.

There are two main areas where we want to improve on MAST. First, we only used top-level service categories from hOurworld but we need to pay attention to subcategories to avoid matches such as between *Offer: 'Art and Crafts: Furniture'* and Request *'Art and Crafts: Artwork.'* We saw many such top-level matches with sub-level mismatches. Secondly, we would like to take advantage of *all* the text that people put in their offers and requests, not just that which we can classify into service categories, as this could provide information critical to matching appropriately, particularly as some people misclassify their offers or requests. It could also help to improve our ability to infer common interests outside of timebank services.

We used a best practice text-mining approach to extract the interest categories from the text in member bios. This approach resulted in high classification accuracy and close matching of profiles. The improvement of text mining algorithm was out of the scope of this evaluation, however, in future work, we plan to refine our classification approach. For example, in the current method, we included the offers and requests of survey participants in the training of classifier. The built model was applied to a different dataset; the member bios. While this approach is technically sound, one might argue that the set of offers and requests of survey participants should be removed from the training set because of the possibility of overlap between text used in a bio and text in the associated service requests and offers. But our current training set included 10,615 offers and requests with only 564 coming from participants. Hence, inclusion of those data points could not significantly affect our results. However, refinement of our algorithm will as well involve refinements in the classification approach and evaluation of the significance in matching quality.

We may also explore entirely different approaches to matching timebankers' profiles that may be more generalizable to situations where there may not be a finite set of categories or clear differentiation between needs and abilities. For example, Facebook profiles could (with users' permission) be matched to service transaction opportunities if a different type of algorithm were implemented.

In addition, we plan to integrate matching with our ongoing development of a system that makes *dynamic* matches, based on location and travel patterns to find people who are nearby, available and heading in the right direction, to further increase the efficiency of transaction partner matching. Thus, MAST is one piece of a larger context-aware matching capability that is still currently under development to bring about the vision laid out in [3]. And ultimately, we hope that all this work will be applicable to other kinds of peer-to-peer matching including dating, friend-finding, skill-finding, and more.

ACKNOWLEDGEMENTS

The authors wish to thank Stephen Beckett, John Saare and the hOurworld leadership and members for help setting up and running our evaluation. We also thank Joe Konstan for advice on MAST's design. This work was funded by National Science Foundation under Grant No. IIS-1407630.

REFERENCES

1. Tatiana P. V. Alves, Marcos R. S. Borges and Adriana S. Vivacqua. 2013. An environment to support the discovery of potential partners in a research group. In *Proceedings of the 17th International Conference Computer Supported Cooperative Work in Design* (CSCWD '13), IEEE, 344-349.

2. Victoria Bellotti, Sara Cambridge, Karen Hoy, Patrick C. Shih, Lisa Renery Handalian, Kyungsik Han and John M. Carroll. 2014. Towards community-centered support for peer-to-peer service exchange: rethinking the timebanking metaphor. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), 2975-2984.
3. Victoria Bellotti, John M. Carroll and Kyungsik Han. 2013. Random acts of kindness: The intelligent and context-aware future of reciprocal altruism and community collaboration. In *Proceedings of Collaboration Technologies and Systems* (CTS '13), IEEE, 1-12.
4. Victoria Bellotti, Kamila Demkova, Dan Turner, Alexander Ambard and Amanda Waterman. In preparation. “Black Market” Activity in Peer-to-Peer Marketplaces. *In preparation, available from the first author on request*.
5. Stephanie L. Brown, Randolph M. Nesse, Amiram D. Vinokur and Dylan M. Smith. 2003. Providing social support may be more beneficial than receiving it: Results from a prospective study of mortality. *Psychological Science* 14, 4: 320-327.
6. Susan Bull, Jim Greer, Gordon McCalla, Lori Kettel and Jeff Bowes. 2001. User modelling in i-help: what, why, when and how. *User Modeling, Vol. 2109 of Lecture Notes in Computer Science*, Springer, Heidelberg 117-126.
7. Edgar S. Cahn. 1986. *Service Credits: A New Currency for the Welfare State.* Welfare State Programme Discussion Paper No WSP/8, Suntory-Toyota International Centre for Economics and Related Disciplines, London School of Economics.
8. Edgar S. Cahn. 2000. *No More Throw-Away People: The Co-Production Imperative*. Essential Books Ltd.
9. Edgar S. Cahn and Jonathan Rowe. 1992. *Time Dollars: The New Currency that Enables Americans to Turn their Hidden Resource – Time – into Personal Security and Community Renewal*. Rodale Press.
10. Jilin Chen, Werner Geyer, Casey Dugan, Michael Muller and Ido Guy. 2009. Make new friends, but keep the old: recommending people on social networking sites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '09), 201-210. http://doi.acm.org/10.1145/1518701.1518735
11. Ed Collom. 2005. Community Currency in the United States: The Environments in which it Emerges and Survives. *Environmental Planning* 37, 9: 1565-1587.
12. Ed Collom, Judith N. Lasker and Corinne Kyriacou. 2012. *Equal Time, Equal Value: Community Currencies and Time Banking in the US*. Ashgate Publishing.
13. Fernando Diaz, Donald Metzler and Sihem Amer-Yahia. 2010. Relevance and ranking in online dating systems. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval* (SIGIR '10), 66-73. http://doi.acm.org/10.1145/1835449.1835463
14. Ed Diener and Robert Biswas-Diener. 2008. *Happiness: Unlocking the Mysteries of Psychological Wealth*. Wiley Blackwell.
15. Rachel Donadio. 2011. Battered by economic crisis, Greeks turn to barter networks. *New York Times*, October 1, page A6. http://www.nytimes.com/2011/10/02/world/europe/in-greece-barter-networks-surge.html
16. Dunbar, R. I. M. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution* 22, 6, 469–493.
17. The Economist. 2013. The Rise of the Sharing Economy. (9 March): Retrieved October 6, 2014 from http://www.economist.com/news/leaders/21573104-internet-everything-hire-rise-sharing-economy.
18. Bradley Efron and Robert J. Tibshirani. An introduction to the bootstrap. *CRC Monographs on Statistics & Applied Probability*, 57. Chapman & Hall.
19. Maryam Fazel-Zarandi, Hugh J. Devlin, Jun Huang and Noshir Contractor. 2011. Expert recommendation based on social drivers, social network analysis, and semantic data representation. In *Proceedings of the 2nd International Workshop on Information Heterogeneity and Fusion in Recommender Systems* (HetRec '11), 41-48. ACM, New York. ISBN 978-1-4503-1027-7
20. Kyungsik Han, Patrick Shih, Victoria Bellotti and John M. Carroll. 2015. It’s Time There Was an App for That Too: A Usability Study of Mobile Timebanking. *International Journal of Mobile Human Computer Interaction* (IJMHCI) 7, 2: 1-22.
21. Yang Sok Kim, Ashesh Mahidadia, Paul Compton, Xiongcai Cai, Mike Bain, Alfred Krzywicki and Wayne Wobcke. 2010. People recommendation based on aggregated bidirectional intentions in social network site. Knowledge Management and Acquisition for Smart Systems and Services, vol. 6232/2010 of *Lecture Notes in Computer Science*, Springer, Berlin. 247-260.
22. Judith Lasker, Lauren Baldasari, Tara Bealer, Ethan Kramer, Zane Kratzer, Rachel Mandeville, Erin Niclaus, Julia Schulman, Danielle Suchow, Jessica Young, Abby Letcher and Anne Rogers. 2006. Building community ties and individual well-being: A case study of the community exchange organization. Paper presented at the 56th Annual Meeting of the Society for the Study of Social Problems, Montreal, Canada, August 10-12, 2006.
23. Judith Lasker, Ed Collom, Tara Bealer, Erin Niclaus, Jessica Young Keefe, et al. 2011. Timebanking and health: The role of a community currency organization in enhancing well-being. *Health Promotion Practice* 12, 1: 102-115.
24. Lei Li and Tao Li. 2012. MEET: a generalized framework for reciprocal recommender systems. In *Proceedings of the 21st ACM international conference on Information*

and knowledge management (CIKM '12). ACM, New York, NY, USA, 35-44.

25. Rensis Likert. 1932. A Technique for the Measurement of Attitudes. *Archives of Psychology* 22, 140: 1-55.
26. Jochen Malinowski, Tobias Keim, Oliver Wendt and Tim Weitzel. 2006. Matching people and jobs: a bilateral recommendation approach. In *Proceedings of the 39th Annual Hawaii International Conference on System Sciences* 6, p. 137c. IEEE Computer Society, Washington.
27. Julia M. Mayer, Starr Roxanne Hiltz and Quentin Jones. 2015. Making Social Matching Context-Aware: Design Concepts and Open Challenges. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). 545-554.
28. Elijah Mayfield and Carolyn Penstein Rosé. 2012. LightSIDE: Text Mining and Machine Learning User's Manual. © 2012 Carnegie Mellon University. Retrieved September 18, 2015 from http://www.cs.cmu.edu/~emayfiel/LightSIDE.pdf
29. Elijah Mayfield and Carolyn Penstein Rosé. 2013. LightSIDE: Open Source Machine Learning for Text In Handbook of Automated Essay Grading, Routledge Academic Press, XCITY, XCountry, 124-135.
30. Matt Moffett and Ilan Brat. For Spain's Jobless, Time Equals Money. *The Wall Street Journal*, August 27, 2012, pp. A1.
31. Stefan Molnar. 2011. Time is of the essence: The challenges and achievements of a Swedish time banking initiative. *International Journal of Community Currency Research* 15, Section A: 13-22.
32. R. Matthew Montoya, Robert S. Horton and Jeffrey Kirchner. 2008. Is actual similarity necessary for attraction? A meta-analysis of actual and perceived similarity. Journal of Social and Personal Relationships, 25, 6: 889-922.
33. Lucie K. Ozanne. 2010. Learning to exchange time: Benefits and obstacles to time banking. *International Journal of Community Currency Research* 14: 1-16.
34. Luiz Pizzato, Tomek Rej, Thomas Chung, Irena Koprinska and Judy Kay. 2010a. RECON: a reciprocal recommender for online dating. In *Proceedings of Recommender Systems*, 207-214.
35. Stephen G. Post. 2005. Altruism, happiness, and health: It's good to be good. *International Journal of Behavioral Medicine* 12, 2: 66-77
36. Robert D. Putnam. 2000. *Bowling Alone: The Collapse and Revival of American Community*, Simon and Schuster.
37. Dennis T. Regan. 1971. Effects of a favor and liking on compliance. *Journal of Experimental Social Psychology* 7, 6: 627-639.
38. Debbie Richards, Meredith Taylor and Peter Busch. 2008. Expertise recommendation: A two-way knowledge communication channel. In *Proceedings of ICAAS*, 35-40.
39. Martin E. P. Seligman. 2002. *Authentic Happiness: Using the New Positive Psychology to Realize Your Potential for Lasting Fulfillment*. Simon and Schuster.
40. Gillian Seyfang. 2002. Tackling social exclusion with community currencies: learning from LETS to Time Banks. *International Journal of Community Currency Research* 6, 3. Retrieved September 22, 2015 from http://ijccr.net/2012/05/23/tackling-social-exclusion-with-community-currencies-learning-from-lets-to-time-banks/.
41. Gillian Seyfang. 2003. With a little help from my friends: Evaluating timebanks as a tool for community self-help. *Local Economy* 18, 3: 257-264.
42. Gillian Seyfang. 2004. Time banks: Rewarding community self-help in the inner city? *Community Development Journal* 39, 1: 62-71.
43. Gillian Seyfang and Noel Longhurst. 2013. Growing green money? Mapping community currencies for sustainable development. *Ecological Economics* 86: 65-77.
44. Gillian Seyfang and Karen Smith. 2002. *The Time of Our Lives: Using Time Banking for Neighbourhood Renewal and Community Capacity Building*. New Economics Foundation.
45. Patrick C. Shih, Victoria Bellotti, Kyungsik Han and John M. Carroll. 2015. Unequal Time for Unequal Value: Implications of Differing Motivations for Participation in Timebanking. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 1075-1084.
46. Steven Tait, Chiara Leonardi, Massimo Zancanaro, Michele Caraviello, Bruno Lepri and Paolo Massa. 2014. Towards collaborative communities: a preliminary study on exchange of goods and services in local contexts. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational* (NordiCHI '14), 999-1002.
47. Loren Terveen and David W. McDonald. 2005. Social matching: A framework and research agenda. *ACM transactions on computer-human interaction (TOCHI)* 12, 3: 401-434.
48. Robert L. Trivers. 1971. The evolution of reciprocal altruism. *The Quarterly Review of Biology* 46, 1: 35-57.
49. Colin C. Williams, Theresa Aldridge, Roger Lee, Andrew Leyshon, Nigel Thrift and Jane Tooke. 2001. *Bridges Into Work? An Evaluation of Local Exchange Trading Schemes (LETS)*. The Policy Press.
50. Svetlana Yarosh, Tara Matthews and Michelle Zhou. 2012. Asking the right person: supporting expertise selection in the enterprise. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12), 2247-2256.