Suggesting (More) Friends Using the Implicit Social Graph[*]

Maayan Roth   MROTH@GOOGLE.COM
Tzvika Barenholz   TZVIKAB@GOOGLE.COM
Assaf Ben-David   ABENDA@GOOGLE.COM
David Deutscher   DUDO@GOOGLE.COM
Guy Flysher   GUYFL@GOOGLE.COM
Avinatan Hassidim   AVINATAN@GOOGLE.COM
Ilan Horn   ILAN@GOOGLE.COM
Ari Leichtberg   ARIL@GOOGLE.COM
Naty Leiser   NATY@GOOGLE.COM
Yossi Matias   YOSSI@GOOGLE.COM
Ron Merom   RONME@GOOGLE.COM

Google Inc, Israel R&D Center

Abstract

Although users of online communication tools rarely categorize their contacts into groups such as "family", "co-workers", or "jogging buddies", they nonetheless implicitly cluster contacts, by virtue of their interactions with them, forming *implicit groups*. In this paper, we describe the *implicit social graph* which is formed by users' interactions with contacts and groups of contacts, and which is distinct from explicit social graphs in which users explicitly add other individuals as their "friends". We introduce an interaction-based metric for estimating a user's affinity to his contacts and groups. We then describe a novel friend suggestion algorithm that uses a user's implicit social graph to generate a friend group, given a small seed set of contacts which the user has already labeled as friends. We show experimental results that demonstrate the importance of both implicit group relationships and interaction-based affinity ranking in suggesting friends. Finally, we discuss two applications of the Friend Suggest algorithm that have been released as Gmail Labs features.

---

[*] This is an updated version of Roth et al. "Suggesting Friends Using the Implicit Social Graph" that appeared in SIGKDD, 2010. Please see that paper for all references.

1. Introduction

One benefit of many online communication channels over offline methods is that they enable communication among groups of people, rather than restricting communication to be peer-to-peer. Email is just one format that supports group conversations, but there are many others, such as photo- and link-sharing, and collaborative document editing. In fact, group communication is so prevalent that our analysis of the Google Mail email network shows that over 10% of emails are sent to more than one recipient, and over 4% of emails are sent to 5 or more recipients. Within enterprise domains, group communication is even more critical. An analysis of the email network of Google employees showed that over 40% of emails are sent to more than one recipient, and nearly 10% are sent to 5 or more recipients.

As opposed to broadcast-style media, such as blogs[1] and micro-blogging platforms like Twitter[2], the information communicated by an individual to a limited group is generally carefully targeted, and may be private. The recipient lists for small-group communications such as emails are selectively constructed by the message senders. We have observed that users tend to communicate repeatedly with the same groups of contacts. This observation has prompted many online communication platforms to provide their users with tools for creating and saving groups of contacts. Some examples are the Google Mail Contact Manager[3], or custom friends lists on Facebook[4].

---

[1] e.g. www.blogger.com, www.wordpress.com
[2] www.twitter.com
[3] http://mail.google.com/support/bin/answer.py?hl=en&answer=30970
[4] www.facebook.com/help/#/help.php?page=768

Despite the prevalence of group communication, users do not often take the time to create and maintain custom contact groups. One survey of mobile phone users in Europe showed that only 16% of users have created custom contact groups on their mobile phones (Kuhn & Wirz, 2009). In our user studies, users explain that group-creation is time consuming and tedious. Additionally, groups change dynamically, with new individuals being added to multi-party communication threads and others being removed. Static, custom-created groups can quickly become stale, and lose their utility.

In this paper, we present a friend-suggestion algorithm that assists users in the creation of custom contact groups, either implicit or explicit. This algorithm is based on analysis of the *implicit social graph*, which is the social network that is defined by interactions between users and their contacts and groups of contacts. We differentiate the implicit social graph from explicit social graphs that are formed by users explicitly adding other individuals as "Friends". The implicit social graph is a weighted graph, where edge weights are determined by the frequency, recency, and direction of interactions between users and their contacts and groups. Our measure of tie strength differs from previous work in that we consider group interactions, as well as peer-to-peer.

We use the implicit social graph to identify clusters of contacts who form groups that are meaningful and useful to each user. Unlike some previous research on contact clustering, we do not consider the content of interactions. Additionally, because the email network that we have studied is private, we do not consider any friend-of-friend ties, either when computing edge weights for the graph, or when computing contact clusters.

Given a user's social network with weighted edges and an initial seed of a few contacts, our friend-suggest algorithm builds a custom contact group that accurately expands the seed. We evaluate the efficacy of our algorithm by comparing to baseline approaches via precision-recall measurements. We show two applications of this algorithm, implemented as Gmail Labs features, called "Don't forget Bob!" and "Got the wrong Bob?" Although our discussion centers around an email network, the network analysis that we have done is applicable to any social graph that is formed by interactions between users and their contacts.

This paper is an extended and revised version of a paper that appeared in KDD 2010. This paper provides additional experimental results about the performance of our seed expansion algorithm. We also provide preliminary empirical results about the performance of the "Don't Forget Bob!" and "Got the Wrong Bob?" Gmail features, along with additional discussion.

2. Characteristics of the Email

Implicit Social Graph

The Google Mail implicit social graph is composed of billions of distinct nodes, where each node is an email address. Edges are formed by the sending and receiving of email messages. For the purpose of our work, we consider a message sent from a user to a group of several contacts as forming a single edge, thereby constructing a directed hypergraph. We call the hypergraph composed of all of the edges leading into or out of a single user node that user's *egocentric network*. We call each hyperedge an *implicit group*, even though it may consist of a single contact. On average, a typical 7-day active user has 350 implicit groups in his egocentric network, with groups containing an average of 6 contacts. Note that this does not imply that the average user has thousands of distinct contacts. Rather, each implicit group is a unique combination of one or more contacts with whom the user has interacted.

Edges in the implicit social graph have both direction and weight. The direction of an edge is determined by whether it was formed by an outgoing email sent by the user, or an incoming email received by the user. There may be both outgoing and incoming edges joining a user and an implicit group, if the user has both sent and received email from the group. We consider a user to have received mail from a group by joining the sender of the mail and the other co-recipients into an implicit group. Thus, if a contact c_1 sent mail to the user u and contacts c_2 and c_3, this is represented in u's egocentric network as an incoming edge from the group $\{c_1, c_2, c_3\}$ to u.

The weight of an edge is determined by the recency and frequency of email interactions between the user and the group. In Section 3.1, we propose one metric for computing edge weight, which we call *Interactions Rank*. We claim that edge weight is an important indicator of the strength of the relationship between the user and a particular group. In the remainder of this paper, we use the terms *edge weight*, *group weight*, and *group importance* interchangeably.

In our work, we draw a sharp distinction between each user's egocentric network and the global or *sociocentric* network that is formed by combining the networks of all users. Although other researchers have found value in clustering contact groups by looking at friend-of-friend edges (e.g. (Kahanda & Neville, 2009)), we restrict our algorithm to look only at a single user's egocentric network during friend suggestion. By showing users suggestions based only on their local data, we are able to protect user privacy and avoid exposing connections between the user's contacts that might not otherwise have been known to him.

The social graph studied in this paper is constructed using the metadata (i.e. timestamp, sender, and recipients) of outgoing and incoming messages set or received via Google

Mail; message content is not included or examined. For the purposes of this research, we used a random sample of the metadata from thousands of interactions, and data was looked at exclusively in aggregate. The experimental results in Section 4 were gathered with the same privacy protections that are used in all Google software development[5] to ensure that developers do not intentionally or unintentionally access contact information about specific users without their explicit consent.

3. Friend Suggest

Our algorithm is inspired by the observation that, although users are reluctant to expend the effort to create explicit contact groups, they nonetheless implicitly cluster their contacts into groups via their interactions with them. For example, while a user may have multiple, possibly overlapping, subgroups of coworkers with whom he exchanges emails, he is unlikely to include his family members in those interactions. The Friend Suggest algorithm, described in this section, detects the presence of implicit clustering in a user's egocentric network by observing groups of contacts who are frequently present as co-recipients in the same email threads. The input to Friend Suggest is a *seed*, which is a small set of one or more contacts that belong to a particular group. This seed could be labeled by the user selecting a few contacts, e.g., as an initial list in the "To:" field of an email. Given this seed, Friend Suggest finds other contacts in the user's egocentric network who are related to the seed, meaning that they are present in the same implicit clusters. Friend Suggest also returns a score for each suggested contact, indicating the goodness of its fit to the existing seed.

The algorithm described in this section is applicable to the problem of group clustering in any interaction-based social graph. For clarity and convenience, we describe it in terms of email interactions.

3.1. Interactions Rank

The first requirement of the Friend Suggest algorithm is an implicit social graph with edges whose weights represent the relationship strength between a user and his implicit groups. We wish to compute edge weights that satisfy the following three criteria:

1. Frequency: Groups with which a user interacts frequently are more important to the user than groups with which he interacts infrequently.

2. Recency: Group importance is dynamic over time.

3. Direction: Interactions that the user initiates are more significant than those he did not initiate.

Regarding recency, we observe that a group with which the user is actively interacting now is more important than one with which the user last interacted a year ago. Overall, recent interactions should contribute more to group importance than interactions in the past. We also note that receiving an email from a contact, a passive interaction, is a weaker signal of closeness than the active interaction of sending an email to that contact. In the most extreme case, we want to be able to rank spammer contacts, from whom the user receives many emails but to whom he sends none, very low in importance.

To satisfy these criteria, we propose Interactions Rank, a metric computed by summing the number of emails exchanged between a user and a particular implicit group, weighting each email interaction as a function of its recency. Interaction weights decay exponentially over time, with the half-life, λ, serving as a tunable parameter. An additional parameter that can be tuned in Interactions Rank is ω_{out}, the relative importance of outgoing versus incoming emails.

Interactions Rank (sometimes abbreviated $\mathcal{IR}$) is computed over a set of email interactions $I = \{I_{out}, I_{in}\}$, according to the following equation:

$$\mathcal{IR} \leftarrow \omega_{out} \sum_{i \in I_{out}} \left(\frac{1}{2}\right)^{\frac{t_{now} - t(i)}{\lambda}} + \sum_{i \in I_{in}} \left(\frac{1}{2}\right)^{\frac{t_{now} - t(i)}{\lambda}}$$

where I_{out} is the set of outgoing interactions between a user and a group, and I_{in} is the set of incoming interactions, t_{now} is the current time, and $t(i)$ is the timestamp of an interaction $i \in I$. Note that according to this equation, an interaction from the current time has a contribution of 1 to a group's Interactions Rank, whereas an interaction from one half-life λ ago contributes $\frac{1}{2}$ and so on.

Interactions Rank is related to the Recency metric proposed by Carvalho and Cohen (Carvalho & Cohen, 2008). However, Interactions Rank calculates the weight of each interaction according to its timestamp, while Recency sorts interactions in chronological order, and weights them on an exponentially decaying scale computed over their ordinal rank. Additionally, Recency does not take into account the direction of each interaction. Ting *et al.* propose an edge-weight metric that considers the role of the interaction participant, but does not take into account the time of the interaction (Ting et al., 2009).

It should be noted that Interactions Ranks do not easily allow for comparisons across several users. A very active user, who sends and receives many emails per day, will have overall higher Interactions Ranks for his implicit groups than a relatively inactive user. However, within a single user's egocentric network, Interactions Rank allows for a clean ordering of the user's implicit groups by

estimated relationship strength. We are actively working on incorporating other signals of importance, such as the percentage of emails received from a contact that the user chooses to read.

3.2. Core Routine

The core routine of the Friend Suggest algorithm, EXPANDSEED is shown in Table 1.

---

function EXPANDSEED(u, $\mathcal{S}$):
  input: u, the user
          $\mathcal{S}$, the seed
  returns: $\mathcal{F}$, the friend suggestions

  1. $\mathcal{G} \leftarrow$ GETGROUPS(u)
  2. $\mathcal{F} \leftarrow \emptyset$
  3. **for each** group $g \in \mathcal{G}$:
  4.   **for each** contact $c \in g$, $c \notin \mathcal{S}$:
  5.   **if** $c \notin \mathcal{F}$:
  6. $\mathcal{F}[c] \leftarrow 0$
  7. $\mathcal{F}[c] \xleftarrow{+}$ UPDATESCORE(c, $\mathcal{S}$, g)

Table 1. Core algorithm for suggesting contacts that expand a particular seed, given a user's contact groups.

The EXPANDSEED function takes as inputs a user, u, who is the mailbox owner of a single egocentric network in the implicit social graph, and a seed, $\mathcal{S}$, consisting of a set of contacts that make up the group to be expanded. EXPANDSEED returns a set of friend suggestions, $\mathcal{F}$, which maps each suggested contact to a score. Each contact's score indicates the algorithm's prediction for how well that given contact expands the seed, relative to the other contacts in u's network. Note that not all contacts from u's network are guaranteed to be returned in $\mathcal{F}$.

Friend suggestions are computed as follows: The user u's egocentric network is extracted from the implicit social graph. The network, $\mathcal{G}$, is represented as a set of contact groups, where each group $g \in \mathcal{G}$ is a set of contacts with whom u has exchanged emails. Each group g has an Interactions Rank, computed as described in Section 3.1, indicating the strength of u's connection to the group g. The goal of EXPANDSEED is to find, among all the contacts in $\mathcal{G}$, those whose interactions with u are most similar to u's interactions with the contacts in the seed $\mathcal{S}$.

EXPANDSEED iterates over each group g in $\mathcal{G}$, computing a score for each contact c that is a member of g. The algorithm does not suggest contacts that are already members of the seed $\mathcal{S}$. Scores for each contact are computed iteratively via a helper function, UPDATESCORE, which takes the contact being considered, the contact's score so far, $\mathcal{F}[c]$, the seed $\mathcal{S}$, and the group g. In the following section, we dis-

cuss several possible scoring heuristics that were considered for UPDATESCORE.

3.3. Scoring Functions

UPDATESCORE is a function template that takes a single contact, c, from a user u's egocentric network and an implicit group g to which c belongs, and returns an incremental score based on the group g's similarity to the seed group, $\mathcal{S}$. The sum of UPDATESCORE for a contact c over all of the implicit groups to which it belongs is an estimate of c's fitness to expand the seed. Because both the implicit groups making up an egocentric network and the seed group that is the input to Friend Suggest are unordered sets of contacts, they can be compared via standard measures of set similarity (Tulloss, 1997). In this work, we look only at set member intersection, leaving more complex metrics for future exploration. We define below several implementations of UPDATESCORE. In the next section, we evaluate their relative merits.

The most basic instantiation of UPDATESCORE is INTERSECTINGGROUPSCORE ($\mathcal{IGS}$), which simply returns a group g's Interactions Rank if the group has a non-empty intersection with the seed set:

$$\mathcal{IGS} = \begin{cases} \mathcal{IR} & \text{if } g \cap \mathcal{S} \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

Intuitively, INTERSECTINGGROUPSCORE finds all the contexts in which the proposed contact c exchanged emails or was a co-recipient with at least one seed group member. However, a larger intersection between the members of the seed group and the members of a given implicit group seems to indicate a higher degree of similarity. INTERSECTIONWEIGHTEDSCORE ($\mathcal{IWS}$) takes this into account:

$$\mathcal{IWS} = \mathcal{IR}(g) \times k|g \cap \mathcal{S}|$$

We investigate the contribution of group importance to friend suggestion by comparing against a metric, INTERSECTINGGROUPCOUNT ($\mathcal{IGC}$), that simply counts the number of groups a contact c belongs to that have some intersection with the seed $\mathcal{S}$. This metric ignores Interactions Rank entirely, and treats all implicit groups as having equal value to the user:

$$\mathcal{IGC} = \begin{cases} 1 & \text{if } g \cap \mathcal{S} \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

Finally, to highlight the importance of using a seed of contacts that characterize a distinct friend group, we compare against an UPDATESCORE instantiation, TOPCONTACTSCORE ($\mathcal{TCS}$), that ignores the seed and always suggests the top-ranked contacts. Contact ranks are computed

by summing the Interactions Ranks of the implicit groups containing each contact:

$$\mathcal{TCS} = \mathcal{IR}(g)$$

In each metric, the final friend suggestion scores are normalized with respect to the highest-ranked contact, so that a single threshold can be used across all users, to cut off the list of suggested contacts.

4. Evaluation

In this section, we evaluate the quality of the Friend Suggest algorithm on real user data. We compare the different scoring functions discussed in the previous section, and explore the impact of seed size of friend prediction.

4.1. Methodology

Evaluation is one of the major challenges of developing algorithms that make predictions based on online social network data. Often, researchers build their data sets by surveying a small set of users who are willing to provide the ground truth about their online social relationships (Gilbert & Karahalios, 2009; Kahanda & Neville, 2009; Yoo et al., 2009). By asking users to categorize their contacts into groups, or rate contacts as "close to me" or "not close to me", researchers can build a labeled data set that serves for both training and testing. However, the nature of this type of survey necessarily limits the number and variety of users who can be included in an experiment. Small sample size and user selection bias can harm the accuracy of the evaluation.

We therefore propose a novel evaluation methodology. From a stream of real email traffic, we randomly sampled 10000 email interactions with between 3 and 25 recipients. Each recipient list is, in essence, a group of contacts that was implicitly clustered by the user. We can test the accuracy of the Friend Suggest algorithm, and compare the relative success of different scoring functions, by sampling a few recipients from each group, and measuring how well Friend Suggest is able to recreate the remaining recipient list. Our approach is similar to the evaluation methodology used by Pal and McCallum (Pal & McCallum, 2006), but whereas they removed one recipient from each interaction and verified whether their algorithm could restore him, we begin with small seeds and attempt to generate multiple additional recipients per email.

To generate the 10000 random test interactions, we first sampled 100000 interactions, and then defined rules that aggressively filtered this set to produce a set of email interactions likely to have been generated by active human users. We define an active user as a user with a minimum of 5 implict groups in his social network, who has sent at least one other email in the 7 days prior to the sampled interaction. We attempt to limit our data set to human users by excluding, via regular expression matching, bots and auto-reply addresses such as "info@domain", and "noreply@domain".

Our experiment tests the ability of our algorithm to use a user's existing social graph to predict his future group interactions. Therefore, when testing our algorithm's ability to predict the remaining recipients in a given email interaction, we use a snapshot of the user's egocentric network based only on interactions that occurred earlier than the sampled interaction.

4.2. Results

The graph below shows a precision-recall curve for the Friend Suggest algorithm using the different scoring functions defined in Section 3.3, with seed groups of size 3. (Precision-recall curves for seed groups of other sizes can be found in (Roth et al., 2010).) For the purposes of our evaluation, we measure precision as the percent of correct suggestions out of the total number of contacts suggested for each seed group, and recall as the percent of correct suggestions out of the total number of email recipients who were not already members of the seed group. A correct suggestion is any contact who was a recipient of the email being evaluated.

Note that the scoring functions that take into account both group membership and relative group importance, INTER-SECTINGGROUPSCORE and INTERSECTIONWEIGHTED-SCORE, significantly out-perform TOPCONTACT-SCORE, which ignores the similarity of the seed contacts to the implicit groups and always suggests the top-ranked contacts, and INTERSECTINGGROUPCOUNT, which ignores the Interactions Ranks of the groups and simply counts the number of groups in which a contact was a co-recipient with at least one seed contact.

Overall, the scoring function with the best performance is INTERSECTIONWEIGHTEDSCORE. For small seed sizes, its performance is similar to INTERSECTINGGROUP-SCORE. However, as the size of the seed contact group increases, INTERSECTIONWEIGHTEDSCORE's performance remains fairly constant, while INTERSECTINGGROUP-SCORE's ability to correctly predict email recipients decreases. Because it includes in each contact's score the score of *every* implicit group that contains at least one member of the seed group, INTERSECTINGGROUPSCORE is noisy and prone to false positives. By taking into account the size of the intersection between each implicit group and the seed group, INTERSECTIONWEIGHTEDSCORE is able to discount the impact of spurious implicit groups that have low similarity to the seed group.

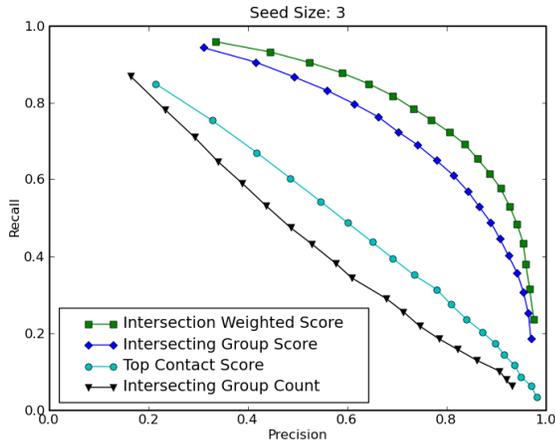These experimental results demonstrate that the Friend

Figure 1. Precision/recall curves for the Friend Suggest algorithm with a seed of 3 contacts, run over the four different scoring functions defined in Section 3.3.

Suggest algorithm, with a correctly chosen scoring function, is able to predict the remaining recipients of an email with high accuracy, given the first few contacts who were added by the user.

4.3. Seed Size Effect

Our experiments showed that for every seed size between 1 and 5, INTERSECTIONWEIGHTEDSCORE outperforms the other three scores. In this subsection we focus on this score, and analyze the effect of different seed sizes. Figure 2 shows the precision/recall curve for different seed sizes. The size of a seed, with the exception of a seed consisting of a single recipient, has negligible effect on the performance of the algorithm. The high success rates, leads us to believe that given the names of the first two email recipients, one can guess a small set of people, and the email will be sent to a subset of them. For example, if Snow White sends an email to Dopey and Grumpy, it is likely that the rest of the recipient list will only contain dwarfs. However, since she is likely to send messages to different subsets of the dwarves, it is hard to know whether the email be exclusive (for example she may add just one other dwarf), or inclusive, containing most of the dwarves or all of them. Further study is required to address this, perhaps by using better thresholding.

5. Applications

We use the Friend Suggest algorithm in two Gmail Labs features, "Don't forget Bob!", and "Got the wrong Bob?"
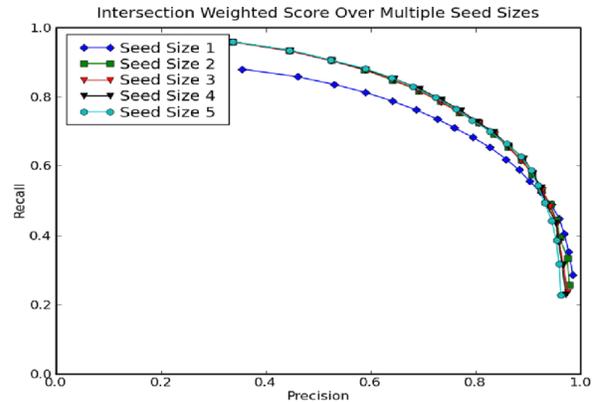

Figure 2. Precision/recall curves for Intersection Weighted Score with seeds of different size.

5.1. Don't Forget Bob!

"Don't forget Bob" is a straightforward user interface on top of the Friend Suggest algorithm. As seen in Figure 3, "Don't forget Bob" operates when a user is composing an email message. The lab treats the first contacts added by the user as the seed set, and uses them to generate a set of possible suggested recipients that the user may wish to add to the email. Once the user has added at least two contacts, the application queries the implicit social graph to fetch the user's egocentric network, and uses Friend Suggest to generate up to four contacts who best expand the seed set of existing contacts. These contacts are displayed as clickable links below the "To:" input field. If the user clicks on a suggestion, or types in another email address, it is added to the list of recipients, and a new set of suggestions is generated.

"Don't forget Bob" has been enabled and used by millions of users, and overall, the user response has been positive. One user posted to the lab's feedback group[6], "This is incredibly helpful for work/school/family groups without having to create contact groups."

Coming up with a good quantitative measure for the performance of "Don't forget Bob" is non-trivial. One possible metric is to count the number of times users clicked a suggestion, and normalize it by the number of suggestions (or the number of times a suggestion was shown). However, since most users do not click on the suggestions but rather use the keyboard to generate the recipient list, this does not capture the effect of the product. Another alternative is to take the final recipient list, and consider the ratio between the number of recipients who appeared in a suggestion, and the number of suggestions (or the number of times a suggestion was shown). The problem with this type of metric is that over 80% of the email messages are sent to a single
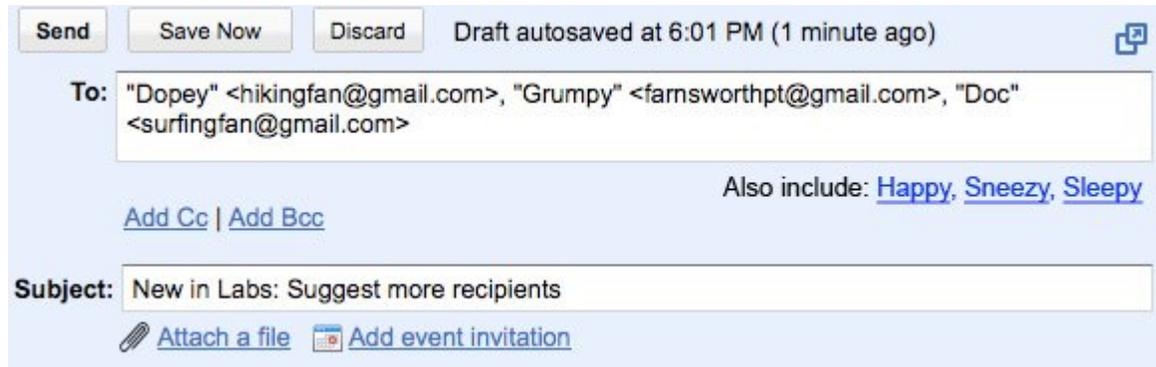
---

[6]http://groups.google.com/group/gmail-labs-help-suggest-more-recipients

Figure 3. Example screenshot of the "Don't forget Bob!" lab in action. Given the user's initial seed contacts, "Dopey", "Grumpy", and "Doc", the Friend Suggest algorithm suggests additional recipients "Happy", "Sneezy", and "Sleepy".

recipient, and at every stage it is more likely that the user will hit the send button, than add another name. Thus, to improve the algorithm in this metric, one might be tempted not to show suggestions after the first recipient.

The main metric we use to evaluate the lab counts the number of suggestions accepted, over the number of times at least one suggestion was shown, except in the last step, where the user hits the send button. Suggestions shown at this point do not increase the denominator. Note that when the algorithm shows suggestions it can show up to three suggestions, and still it is only charged once in this benchmark. The results for this metric are very good - the ratio between the number of accepted suggestions and the number of times a suggestion was shown (ignoring the last recipient) is above 0.8. Moreover, this precision comes at a good coverage, and suggestions are shown for more than half the email messages.

5.2. Got the Wrong Bob?

A more complex use of the Friend Suggest algorithm can be found in the "Got the wrong Bob?" lab, shown in Figure 4. "Got the wrong Bob" addresses the known problem of email autocompletion errors (Carvalho & Cohen, 2007). While previous approaches have relied heavily on message content, "Got the wrong Bob" uses the Friend Suggest algorithm to detect the inclusion of contacts in a message who are unlikely to be related to the other recipients.

The WRONGBOB algorithm, shown in Table 2, works as follows: From the current recipients of an email that have been entered by the user, the algorithm attempts to find a single contact whose removal and replacement with another contact from the user u's egocentric network would lead to a more coherent recipient list. For each contact c_i in the current recipient list L, WRONGBOB builds a seed set that includes all of the members of L except c_i (lines 4-5). This seed is expanded via EXPANDSEED to generate a set of contacts that are similar to the current members of

the seed. If the excluded contact c_i is a member of the suggestion set, it is considered to be related to the other recipients and unlikely to be a mistake (lines 7-8). WRONGBOB therefore stops searching for a replacement for c_i.

If, however, c_i is not returned as a suggestion from EXPANDSEED, it is a potential mistake. WRONGBOB searches for another contact that could replace c_i. Each contact c_j in the result set returned by EXPANDSEED is compared to the error candidate c_i via a helper function IS-SIMILAR. In our implementation, we measured similarity by checking to see if c_j was listed as an autocomplete suggestion at the time that the user entered the contact c_i. If c_i and c_j are similar, and c_j's score as a member of the seed expansion is higher than the current maximum, c_i and c_j are saved as the current candidate pair (lines 10-13). After examining all contacts in L, the candidate pair with the highest score is returned and displayed to the used as "Did you mean Contact A instead of Contact B"?

For example, consider the recipient list $L = \{a, b, c\}$. Assume that when removing a to create the seed list $\{b, c\}$, EXPANDSEED generates the suggestion set $\{a, d\}$. In this case, because the excluded contact a is a member of the suggestion set, WRONGBOB determines that it is not a mistake. Then, when removing b, if the algorithm observes $\{b', d\}$, where b' is similar to b but d is not, the algorithm will consider $\{b, b'\}$ as candidates for replacement. If, after removing c, the algorithm generates another candidate pair $\{c, c'\}$, then it will return the pair with the highest score.

Like "Don't forget Bob", the "Got the wrong Bob?" lab has been enabled and used by millions of users. However, as in the previous lab, measuring the true performance is challenging. While some users do click on many suggestions, others prefer to edit email recipient lists via the input textbox after being notified by the lab that they may have made a mistake. As we do in the "Don't forget Bob" lab, we consider both actions as a success for the algorithm. In addition, some users only delete the extra recipient, or only
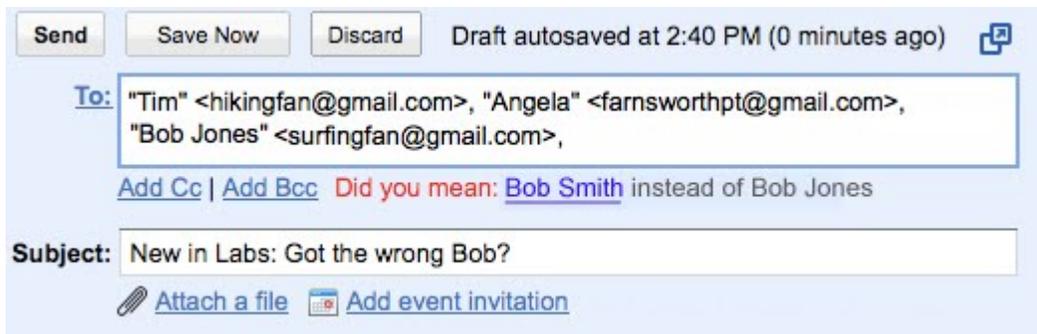
Figure 4. Example screenshot of "Got the wrong Bob?" In the context of an email to recipients "Tim" and "Angela", the Wrong Bob algorithm detects that the user may have intended to include "Bob Smith" instead of "Bob Jones".

```
function WRONGBOB(u, L):
  input: u, the user
         L, a list of the recipients of an email
  returns: a pair {c,s} where
         c is a contact ∈ L
         s is a suggested contact to replace c

  1. score_max ← 0
  2. wrongRecipient ← null
  3. suggestedContact ← null
  4. for each contact c_i ∈ L:
  5.    seed ← L \ c_i
  6.    results ← EXPANDSEED(u, seed)
  7.    if c_i ∈ results:
  8.       continue
  9.    for each contact c_j ∈ results:
  10.       if ISSIMILAR(c_i, c_j)
             and score(c_j) > score_max:
  11.          score_max ← score(c_j)
  12.          wrongRecipient ← c_i
  13.          suggestedContact ← c_j
  14. return {wrongRecipient, suggestedContact}
```

Table 2. The WRONGBOB algorithm which, based on the user's egocentric network, checks if one of the existing recipients would be a good candidate for replacement with another contact.

add the proposed recipient. In almost 70% of the time, the users accept both suggestions, deleting the wrong Bob and adding the correct one. In almost 90% of the cases, the user accepts at least one suggestion (usually adding the extra Bob). These metrics reflect high precision, but low occurrence, as only about 1% of the messages trigger the lab. This infrequent triggering is unsurprising, as over 80% of messages are sent to only one recipient, and thus should not trigger the lab. Out of messages which are sent to two or more recipients, almost all the messages have the correct list of recipients and again should not trigger the lab. We therefore believe the overall recall of the feature to be high.

6. Related Work

Please see related work in (Roth et al., 2010).

7. Conclusions

In this paper, we studied the *implicit social graph*, a social network that is constructed by the interactions between users and their groups. We proposed an interaction-based metric for computing the relative importance of the contacts and groups in a user's egocentric network, that takes into account the recency, frequency, and direction of interactions. We then defined the Friend Suggest algorithm which, given a single user's egocentric network with computed edge weights and a seed set of a few labeled contacts, finds other contacts who are related to the seed contacts, and therefore form a semantically meaningful group. We demonstrated the effectiveness of the Friend Suggest algorithm via a novel experimental methodology. Finally, we showed two applications of the Friend Suggest algorithm, the Gmail Labs "Don't forget Bob!" and "Got the wrong Bob?", and presented some data on their performance.

Although the experimental results described in this paper were performed by examining email interactions from the Google Mail system, the algorithms and approaches described in this paper apply to any interaction-based social network. Some other interaction types that could form similar implicit networks are photo and document sharing, instant messenger chatting, online calendar meeting invitations, or comments on blog posts. Even offline interactions, such as mobile text messages or telephone calls, form an implicit social graph between individuals and groups. Our future research is intended to study the relative importance of different interaction types in determining the social relationships between individuals. We are also interested in exploring other applications of the Friend Suggest algorithm, such as identifying trusted recommenders for online recommendation systems, or improving content sharing between users in various online contents.