Getting Places:
Collaborative Predictions from Status

Mohamad Monibi and Donald J. Patterson

University Of California, Irvine, USA
{mmonibi,djp3}@ics.uci.edu

Abstract. In this paper we describe the use of collaborative filtering to make predictions about *place* using data from custom instant messaging status. Previous research has shown accurate predictions can be made from an individual's personal data. The work in this paper demonstrates that community data can be used to make predictions in locations that are completely new to a user.

1 Introduction and Related Work

The practice of online status setting has evolved into a form of informal interaction and information exchange. It is an indicator of interruptibility in instant messaging (IM) clients [1] and expresses current mood in social networking applications [2]. The popularity of status messages has given rise to services dedicated to status setting and sharing (e.g., twitter.com, ping.fm). We are primarily interested in its use to indicate the current place of a user (i.e., "micro-presence").

To provide an accurate and useful indication of place, users must frequently update their status, for example every time they enter a new location [3] or engage in a new activity, which is cumbersome. Automating this task can greatly assist the user in maintaining up-to-date status. Accuracies of 80-90% can be achieved in automatically predicting a user's place and activity from their *own* history of status setting behavior [4]. However, these techniques do not address how to make predictions when a user arrives at a place for the first time.

In this paper, we describe the use of collaborative filtering techniques for recommending socially appropriate and relevant place descriptions for mobile IM status using the history of an entire community of users. These place labels often incorporate colloquial understandings of place, putting them outside the ability of commercial point-of-interest databases for choosing appropriate labels.

Recent work has combined sensing with semantic labeling to provide lightweight interpretation of sensor data. WatchMe is a tool based on interpreting sensors as graphical icons to communicate remote context between members of close relationships [5]. Reno allowed users to associate labels to cell tower connections and then to activate rules based on entering those zones [6]. Finally, Connecto allows users to associate labels with combinations of cell-phone towers [7] and IMBuddy with combinations of WiFi access points [8].

Zonetag [9] is a context-aware media annotation system that uses sensor data to provide tag suggestions to annotate photos taken with a cell phone.

Zonetag suggestions come from the tag history of the user, the user's social circle, and the public. Scoring of tags generated by others depends on the social distance between users. Our research builds on ideas of using community labels from ZoneTag but uses wifi access points instead of cell towers and collaborative filtering for label prediction.

2 Methodology

Dataset To develop effective techniques for predicting status we first obtained data collected for the Nomatic*IM project [10]. The data was collected over approximately 3 months by 72 users. It contains 19,664 status entries each of which consists of a set of sensor readings paired with 390 unique place labels. Although several sensors were available, we utilized only wifi access point mac addresses. Our data contained sightings of 2352 unique wifi access points. Users had *connected* to 501 of those access points.

Collaborative Filtering We applied collaborative filtering as a method for making place label predictions from this data. Collaborative filtering is a method for recommending items of interest to a user based on the interests of other similar users [11]. It had much early success in movie recommender systems [12] and remains an active area of research. Our approach treats place labels as the item which is "recommended."

Collaborative filtering assumes that given the data, a *ranking* function exists for ordering recommendations for one user based on one other user. Then, a user *similarity* function describes how to order recommendations from multiple users for a global recommendation.

As a baseline we created a recommendation based on Equation 1, which simply calculates the most **probable** label for a wifi access point across all users. $\psi(l, w, u)$ is a function which counts the number of times a label, l, is used at a given wifi access point, w, by a user, u.

$$P_B(l|w) = \frac{1}{|U|} \sum_{u \in U} \left[\frac{\psi(l, w, u)}{\sum\limits_{l' \in L} \psi(l', w, u)} \right] \tag{1}$$

Ranking Functions When making a recommendation for a user, u^*, the first ranking function, Equation 2, ranks the most probable labels used at his current **single wifi** access point, w_{u^*} based on one other user, u.

$$P_{R_1}(l|w_{u^*}, u) = \frac{\psi(l, w_{u^*}, u)}{\sum\limits_{l' \in L} \psi(l', w_{u^*}, u)} \tag{2}$$

The second ranking function, Equation 3, ranks the most probable labels using the same wifi access point, w_{u^*} and then incorporates the chance that the same label was generated by **all visible wifi** access points, $\widehat{w_{u^*}}$, which u^* can also currently see. The trade-off between w_{u^*} and $\widehat{w_{u^*}}$ is managed by $0 < \alpha < 1$.

$$P_{R_2}(l|w_{u^*}, \widehat{w_{u^*}}, u) = \alpha \; P_{R_1}(l|w_{u^*}, u) + \frac{(1-\alpha)}{|\widehat{w_{u^*}}|} \sum_{w \in \widehat{w_{u^*}}} P_{R_1}(l|w, u) \tag{3}$$

Similarity Functions To combine the results of the ranking functions, we formulated two distinct user similarity metrics. Since we were looking to predict place labels, the first metric, Equation 4, asserts that similar users are often **co-located**. The function, $\phi(u, w)$, returns the number of times user u, entered *any* label at a wifi access point, w, from the set of all wifi access points, $\widehat{w}$.

$$\mathrm{S}_1(u_1, u_2) = \frac{1}{|\widehat{w}|} \sum_{w \in \widehat{w}} \left[1 - \frac{|\phi(u_1, w) - \phi(u_2, w)|}{max(\phi(u_1, w), \phi(u_2, w))} \right] \tag{4}$$

In this method the similarity between users is the average of the percentage of times that two users entered any status at the same location.

The second metric, Equation 5 combines how often two users were in the same physical location and *also* labeled the location the same way. This metric captures **place agreement**. We calculate this as the product of the probability of u_1 and u_2 using the same label, l, at the same location and sum that over all labels at all wifi access points. $\widehat{l_w}$ is the set of labels used at an access point, w.

$$\mathrm{S}_2(u_1, u_2) = \frac{1}{|\widehat{w}|} \sum_{w \in \widehat{w}} \frac{1}{|\widehat{l_w}|} \sum_{l \in \widehat{l_w}} \left[\frac{\psi(l, w, u_1)}{\sum_{l' \in \widehat{l_w}} \psi(l', w, u_1)} \frac{\psi(l, w, u_2)}{\sum_{l' \in \widehat{l_w}} \psi(l', w, u_2)} \right] \tag{5}$$

While it is possible (and may indeed be useful) to incorporate external user profile or demographic data into similarity metrics, for the purposes of this research and evaluation, we assumed the only data available to us for determining similarity is the data actively generated by the users, namely the status messages and the sensor data that accompanies them.

3 Results and Discussion

In our dataset, there were 746 user-location combinations (location here is based on connected access points). Only 344 of these instances consisted of a location that was visited by more than one user; these are the only cases where community based suggestions were possible. A recommendation for a user, u^*, at a wifi access point, w, was calculated by combining similarity and ranking scores as follows:

$$R(u^*, w) = \arg\max_{l \in L} \left(\sum_{u \in U} P(l|w, u) S(u^*, u) \right) \tag{6}$$

Figure 1 shows the percentage of times the label that the test user chose was in the top n recommended labels by combinations of ranking functions and similarity functions.

Our best case prediction accuracy is 23%, much better than the 0.2% which would be expected by random. Our technique shows a clear improvement over our baseline and it appears that our similarity metrics are a more important consideration than our ranking techniques.

To more deeply understand our results we categorized and counted the prediction errors as follows:

Syntactical These errors result from different phrasings of the same label. For example, "peets" and "peet's coffee". Errors resulting from use of abbreviations or acronyms also fall into this category, e.g. "technology garden" and "tech garden."

Technical These errors result from technical limitations of the system such as the limited accuracy of our positioning mechanism, which resulted in multiple locations being recognized as one. For example, "tech garden" and "galen lab" are across from each other in the hallway of our department, beneath the resolution of our location recognition.

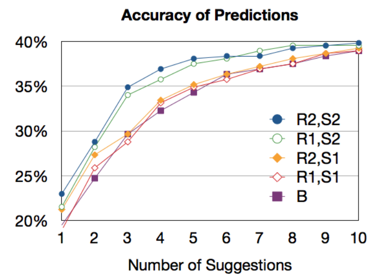

Fig. 1. Prediction accuracy graphed against the number of suggestions included.

Conceptual Four error categories are due to variations in conceptualization of place among different users and situations:

Activity Bleed The concept of place is often intermingled with activity, resulting in labels that actually describe activity rather than place. E.g. "dbh 1300" and "class", "austria" and "ubicomp."

Specificity Many locations have labels describing both a larger place and a specific subsection of that place, for example a room within a building. E.g. "engineering tower" and "et 201", "austria" and "Innsbruck."

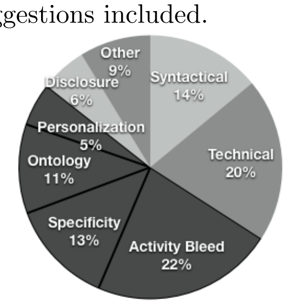

Fig. 2. Types of errors.

Ontology A place may have different names depending on the ontology used by users. For instance, one user, e.g. a college student, may elect to label a place "seminar room" while another user, e.g. from building's facilities staff, may label the same place is "5011."

Personalization Some places are attributed to a unique person or group of persons, resulting in a personalized approach to labeling. E.g. "office" and "André's office."

Disclosure In some instances, users may elect to leave the place label blank when updating status which we do not allow as a suggestion.

Other Errors that we could not interpret/categorize due to lack of information.

Figure 2 shows the percentage of each type of error when evaluating the top ranking suggestions ($n = 1$) generated using the R2, S2 prediction method.

Based on our results, we hypothesize that actually suggesting collaboratively filtered suggestions for place labels will lead to a convergence of place labels among groups of similar users driving accuracy numbers even higher. We expect this to eliminate virtually all syntactical and many conceptual errors.

4 Conclusion

Existing point-of-interest databases don't provide much flexibility in addressing individual users' place-naming style, and existing techniques based on personal histories don't provide labels in new places. The collaborative filtering methods described in this paper allow for both of these capabilities. By combining personal and community data sources, it should be possible to make good predictions when a user first visits a place.

References

1. Fogarty, J., Lai, J., Christensen, J.: Presence versus availability: the design and evaluation of a context-aware communication client. Int. J. Hum.-Comput. Stud. **61**(3) (2004) 299–317
2. Smale, S., Greenberg, S.: Broadcasting information via display names in instant messaging. In: GROUP '05: Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work, New York, NY, USA, ACM (2005) 89–98
3. Lehikoinen, J.T., Kaikkonen, A.: PePe field study: constructing meanings for locations in the context of mobile presence. In: MobileHCI '06: Proceedings of the 8th conference on Human-computer interaction with mobile devices and services, New York, NY, USA, ACM (2006) 53–60
4. Patterson, D.J., Ding, X., Kaufman, S.J., Liu, K., Zaldivar, A.: An ecosystem for learning and using sensor-driven IM messages. IEEE Pervasive Computing: Mobile and Ubiquitous Systems (to appear)
5. Marmasse, N., Schmandt, C., Spectre, D.: Watchme: Communication and awareness between members of a closely-knit group. In Davies, N., Mynatt, E.D., Siio, I., eds.: Ubicomp. Volume 3205 of Lecture Notes in Computer Science., Springer (October 2004) 214–231
6. Smith, I.E., Consolvo, S., LaMarca, A., Hightower, J., Scott, J., Sohn, T., Hughes, J., Iachello, G., Abowd, G.D.: Social disclosure of place: From location technology to communication practices. In Gellersen, H.W., Want, R., Schmidt, A., eds.: Pervasive. Volume 3468 of Lecture Notes in Computer Science., Springer (May 2005) 134–151
7. Barkhuus, L., Brown, B., Bell, M., Sherwood, S., Hall, M., Chalmers, M.: From awareness to repartee: sharing location within social groups. In: CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, New York, NY, USA, ACM (April 2008) 497–506
8. Hsieh, G., Tang, K.P., Low, W.Y., Hong, J.I.: Field deployment of IMBuddy : A study of privacy control and feedback mechanisms for contextual im. In McCarthy, J., Scott, J., Woo, W., eds.: UbiComp08, New York, NY, USA, ACM (September 2008) 91–108
9. Naaman, M., Nair, R.: Zonetag's collaborative tag suggestions: What is this person doing in my phone? IEEE MultiMedia **15**(3) (2008) 34–40
10. Patterson, D.J., Ding, X., Noack, N.: Nomatic: Location by, for, and of crowds. In Hazas, M., Krumm, J., Strang, T., eds.: Location- and Context-Awareness, Second International Workshop, LoCA 2006, Dublin, Ireland, May 10-11, 2006, Proceedings. Volume 3987 of Lecture Notes in Computer Science., Springer (2006) 186–203
11. Goldberg, D., Nichols, D., Oki, B.M., Terry, D.: Using collaborative filtering to weave an information tapestry. Commun. ACM **35**(12) (1992) 61–70
12. Miller, B.N., Albert, I., Lam, S.K., Konstan, J.A., Riedl, J.: Movielens unplugged: experiences with an occasionally connected recommender system. In: IUI '03: Proceedings of the 8th international conference on Intelligent user interfaces, New York, NY, USA, ACM (2003) 263–266