

Assessing Physical Location as a Potential Contextual Cue For Adaptive Mobile Contact Lists

A. Plessas, O. Georgiadou, V. Stefanis, J.
Garofalakis

Department of Computer Engineering & Informatics
University of Patras
Patras, Greece

{plessas, georgiadou, stefanis, garofala}@ceid.upatras.gr

A. Komninos

Computer and Information Sciences
University of Strathclyde
Glasgow, UK
andreas.komninos@strath.ac.uk

Abstract—The retrieval of the appropriate contact in order to start a new communication session from the contact repository of mobile devices can be a time consuming procedure since mobile contact lists usually contain hundreds of items. Several researchers have focused in the past on predicting the next contact a user is likely to call, a task that could prove useful in designing adaptive context-aware interfaces for the mobile contact list. Most of the researchers propose several contextual dimensions that could be used to predict the next callee, location being one of them. However, none of these research works have ever examined the impact of location on mobile communications and only few have actually incorporated this contextual dimension on their implementations. In this paper, we examine physical location as a contextual cue for adaptive mobile contact lists by analyzing call logs from the Nokia Mobile Data Challenge dataset. Our work indicates that, contrary to previous literature, the consideration of physical location as a context dimension does not necessarily lead to improvements in the accuracy of predicting the likelihood of communication with contacts for all types of users included in the dataset under review. Finally, we also discuss the possible reasons behind this limited impact.

Keywords—physical location; context awareness; call prediction; adaptive contact list

I. INTRODUCTION

As mobile phones are primarily communication devices [1], a common task that their users often accomplish in order to start a new communication session is that of contact retrieval. There exist several mobile interfaces from which a contact can be retrieved: contact list, recent call list, speed dial widgets, notifications etc. From all these alternatives, the contact list appears to be the most popular choice of users when placing a new call [2]. However, the mobile contact list repository may contain hundreds [3] or even thousands (especially when importing contacts from social networking and 3rd party communication services) [4] of items and at the same time users are reluctant to delete unused contacts [4][5]. As a result and taking also into account the screen size limitations of handheld devices, the cognitive load required for the task of contact retrieval increases when performed on a mobile phone.

Adaptive interfaces that take advantage of mobile context could offer considerable help to the users while searching for a contact. For the case of making an external call, such interfaces

could include adaptive contact and speed call lists. Such solutions require a context aware algorithm that predicts at any time the next contact to be called. In previous work [6][7], we proposed such an algorithm that uses the context dimensions of frequency and recency of communication and whose top suggestions feed a speed dial list, facilitating contact retrieval [8].

Several context dimensions have been proposed for adaptive algorithms, location being one of them. It is unanimously accepted in literature that location appears to be one of the most important aspects of context in mobile communication [9], while mobile users rank location as the second more important aspect of context [10]. However, for the case of adaptive contact interfaces this consensus seems to be based on little more than an assumption, since none of the research efforts have ever examined the actual impact of the caller location on selecting the contact to start a new call with.

In this paper, we attempt a thorough examination of physical location as a context dimension for predicting the next callee. We analyze call logs from the Nokia Mobile Data Challenge dataset [11][12] and we present interesting patterns regarding the effect of physical location, on a city level, on mobile communication. We show that using a location dependent frequency dimension for predictions instead of normal frequency, does not necessarily lead to improvements in the accuracy of predicting the likelihood of communication with contacts for all types of users included in the dataset under review and we also discuss the possible reasons behind this limited impact.

II. RELATED WORK

Several researchers have focused on proposing context aware call prediction algorithms that could be used to create adaptive contact or speed dial list interfaces. In [13], an approach to predict outgoing calls analyzing historical mobile call log data is described. Three context dimensions (number of calls with each contact, monthly and daily regularity) are combined to capture the frequency and regularity of communication behavior. Geographic location is not used; however the authors refer to it as a possibly important factor for calculating the weight of each contact in outgoing communication and present a scenario where it would add value to the prediction algorithm.

In [14], the researchers propose a prediction algorithm that also incorporates the dimensions of frequency and recency of communication with each contact. To improve the usability of the contact list, they present an interface for a dynamic speed dial widget with four contact positions. Again, location is mentioned as a possible contextual dimension (although not used in this paper) and the authors highlight the need to investigate it in their future work.

Moreover, in [15] the authors present a probabilistic model for call prediction that is incorporated in a dynamic Intelligent Address Book that can be used to save time from searching for a phone number in a typically lengthy contact book. Location is once again identified as a possibly appropriate contextual cue for the task of contact retrieval, although without being used as a context cue or without further investigation on its importance.

Another related work can be found in [16], where a method for modeling social and personal context is presented and semantic location is part of the context model. The model is used by a smart phonebook application that recommends contacts based on the user's current situation. The researchers performed an evaluation of their application with only one participant and only for a short period (12 days) without actually examining the impact of any of the individual contextual cues to the prediction performance. Although such a limited evaluation cannot provide reliable results, we should note that the inclusion of location did not seem to improve performance, as the reported prediction accuracy is close to that presented in [15] and [7] (around 70% for 5 suggestion entries).

Finally, exploiting co-occurrence of behavior patterns of mobile users is another approach for predicting outgoing call events and providing UI call shortcuts [17]. Semantic place (e.g. home, work, outside) is considered part of the user's context in this case. As the prediction engine provides suggestions according to a confidence prediction level and not when this level is below a specific threshold (no guess situation), the researchers present the trade-off between recall and precision. To directly compare with other research efforts on call prediction, one should consider the case where recall is 100% and we can see that in this case the accuracy is close or even lower to other outgoing call predictors (around 30% for 1 suggestion, 45% for 3 suggestions and 55% for 5 suggestions).

While location is widely considered as a candidate context dimension for adaptive contact list interfaces, research on its effect on mobile communication, and especially for the case of outgoing phone calls, is limited. We have found just three examples in literature where this problem is considered. Firstly, in [18], the authors analyzed one year of mobile telecommunications data and investigated the relationship between people's calls and their physical location. They considered pairs of users that communicate with each other and they discovered that more than 90% of users who called each other have also shared the same space (cell tower), even if they live far apart. Moreover, they found that 69% of users who call each other frequently (at least once per month on average) have shared the same space at the same time. Additionally, in [19] the effect of semantic place (home, office, other meaningful,

elsewhere, abroad) on smartphone communication services is examined. Regarding outgoing voice calls, it is reported that they are placed least intensively in the home context where, nevertheless, their length is the longest. Finally, in [20] a study that involves a questionnaire completed by mobile users investigated location-based phone usage, without however specifying the feature of the phone that was used. According to the reported results, most users were using their phone in the office during weekdays (with home being the second most popular answer) and while driving or at home during weekends.

In summary, it can be seen that although user location seems to be regarded as an important contextual cue for predicting the next contact to be called, it has been utilized together with other context cues without being analyzed for its effect and without investigating its efficacy.

III. METHODOLOGY

A. Dataset

To investigate the role of location context in predicting calls, we used the NOKIA Mobile Data Challenge Dataset [12], which is probably the richest mobile dataset currently available, especially regarding collected location attributes [11]. Recorded data were collected from mobile phones in the framework of the Lausanne Data Collection Campaign that took place from 2009 to 2011 and tracked users' social interactions (calls, short messages, scanned bluetooth devices), location, media creation and behavior (applications usage, activity detection etc.). The number of participants reached 185 (38% female) and more than 240.000 calls were logged, while more than 26.000.000 location points were recorded [12]. Location sources include GPS traces providing raw spatial coordinates, wireless LAN access points whose location was estimated by available GPS data and GSM cell towers. However, to protect participants' privacy, position accuracy of GPS tracks was diluted. Also, since users could switch GPS or Wi-Fi off at will during the collection, many events in the dataset are not spatially tagged.

In order to overcome this obstacle and as we wanted to assign a location to each call, since this information was not directly available, we decided to consider as a call's representative position the closest location record within a timeframe of x minutes before or after the moment that the call took place. For this reason, we calculated the total percentage of calls for which a location was successfully assigned for different values of the timeframe x (Table I). When the timeframe equals to 15 minutes this percentage is higher than 70%, which seemed to be a fair threshold value. As it can be observed, expanding the timeframe window by 15 minutes results only in an additional 1% of geo-located calls, while, undoubtedly, uncertainty about the actual position of the user at the call time increases. Consequently, we decided to adopt the timeframe of 15 minutes, as we believe that it is quite a short period to consider that the location of the user remains practically unchanged, at least at a city level. Obviously, for some calls it was not possible to assign a location due to lack of location records within this timeframe. Setting a shorter timeframe would result in even more calls without a matched

position, while as already discussed a longer timeframe would introduce a higher degree of uncertainty as regards the accuracy of the assigned location, without significant profit regarding the number of geo-located calls.

TABLE I. TOTAL PERCENTAGE OF GEO-LOCATED CALLS FOR DIFFERENT VALUES OF TIMEFRAME

timeframe x (minutes)	total % of geo-located calls
5	69.4%
10	69.8%
15	71.3%
30	72.4%
45	73.4%
60	74.5%




Fig. 1. Percentage of geo-located calls for each user in descending order

After this procedure we noticed that there were significant variations to the volume of calls that were attributed to a location among the different users. There were users with more than 90% of successful matching between call and location records, but also there were users (probably switching off frequently the recording software) for which only 20% of call records could be tagged with position coordinates. We decided to reject from our analysis users with less than 70% of geo-located call records as we believe that keeping users below this threshold could prove harmful for the reliability of our prediction algorithm results. Moreover, we excluded users that could not be considered valid for call log analysis, since some of them had a call log covering an extremely short period of time (less than 30 days), while others had too few call records (less than 100) or a very sparse call usage of their phone (less than 2.25 calls per day), a procedure also followed while carrying out the experiments described in our previous work [7]. These extreme characteristics can be explained, since there were users who left the experiment early or others who often decided to turn off the recording software, according to the dataset's documentation. At the end, we were left with 54 users, which is not very far away from the number of 80 users distinguished as those with the highest quality location data for the mobile data challenge [12], taking into account that we additionally consider communication quality as an exclusion factor. The average percentage of calls that were tagged with location information for these users is 82,05% (sd=5.42%),

while in Fig. 1 we show the distribution of geo-located calls' percentages for each user.

B. Physical vs Semantic Location

Physical location refers to an absolute geographic position on the surface of the earth, usually expressed as a pair of geographic coordinates. An address like Downing Street 10, London, UK can be derived from location coordinates and from this, one is able to understand in which country, city and specific area the geographic point belongs. On the other hand, semantic location implies the attachment of a semantic tag to a location, such as "home", "work", "university" etc.

Due to the fact that the available location data from GPS receivers were modified to protect the participants' privacy and since the accuracy from sources such as wireless LAN access points and cell towers is low, we decided to examine the impact of physical location to outgoing call communication as a first step of our research and leave semantic location for future research. This context dimension could prove useful for scenarios like the following:

"George grew up in a small town of western Greece. After graduating from high-school he moved to Athens as he found a good job opportunity. From time to time, he visits his hometown where his parents still live. When being there, he usually calls his childhood friends and arranges to meet them at their favorite pub. Although they do not communicate often when he is in Athens, they still feel close to each other."

"Maria works for an ICT company based in London. However, she often travels abroad, to the company's office in Madrid. When in Madrid, she often calls her local colleagues from her mobile phone, while when she is at home they mostly communicate via Skype or email."

A similar situation is also described in [13], where location is proposed as a useful contextual dimension for outgoing call prediction for users that spend periods of their lifetime moving between different countries. These scenarios emphasize how communication behavior may change when users move to a different physical location at the level of another city in the same or another country. In this work, we focus on the effect of location changes at this level.

C. Location Clustering

To perform our analysis, we clustered for each user all geo-located calls at a city level. There are several algorithms for spatial clustering of geographic tracks, with one of the most widely used in bibliography being DBSCAN [21].

In DBSCAN, the definition of a cluster is based on the notion of density-reachability. A point q is directly density reachable from a point p , if it is within a given distance ϵ from p . A point q is density reachable from a point p if there is a sequence of points p_1, \dots, p_n with $p_1=p$ and $p_n=q$ where each p_{i+1} is directly density-reachable from p_i . Two points p and q are density connected if there is a point o such that both p and q are density reachable from o . A cluster satisfies the following properties: i) all points within the cluster are mutually density-connected and ii) if a point is density-reachable from any point

of the cluster, it is part of the cluster as well. Points that are not part of any cluster are considered as noise.

The algorithm requires two parameters: the minimum distance ε to consider two points directly density reachable and the minimum number of points $minPts$ required to form a cluster. The appropriate values for these parameters depend on the dataset and on the purpose of clustering and are usually empirically defined. After experimenting with different values of these parameters and having in mind that we wanted a city level clustering, while there was no point for clusters with a very small number of calls, we decided to use an ε value of 5 km and set the value of $minPts$ to 2.5% of the total calls of the user. Higher values of the ε parameter resulted in clusters covering much larger areas (e.g. the clusters depicted in Fig. 2 where $\varepsilon = 10$ km), while lower values resulted in distinct clusters within the same city (as shown in Fig. 3 where $\varepsilon = 3$ km) or higher percentage of points characterized as noise. On the other hand, higher values of the $minPts$ parameter resulted in higher percentage of noise calls (which however often seemed optically to form clusters at specific areas) and lower values led to the formation of clusters consisting of a very small number of calls (often < 10). An example of the formed clusters for one of the participants is shown in Fig. 4. In Fig. 5, another example of a clustering with a high percentage of noise data is depicted. We can see that noise calls seem to be either isolated calls at places where the user does not have frequent presence or calls between two clusters, probably while moving from one location to another.




Fig. 2. Calls clustering with $\varepsilon = 10$ km




Fig. 3. Calls clustering with $\varepsilon = 3$ km



Fig. 4. Map of the clustered calls (4 clusters) for a user in the dataset. Markers with the “N” symbol represent calls classified as noise.



Fig. 5. Map of the clustering for the calls of a user with a high percentage of noise data.

IV. DATA ANALYSIS

A. Clustering Results

After applying DBSCAN for the spatial clustering of each user's calls, we observe that twenty users appeared to have voice communication only while being in one location (cluster). These users do not present any significant mobility during the period that they participated in the experiment. The rest of them have calls distributed from two to five clusters as shown in Table II.

TABLE II. DISTRIBUTION OF CALLS IN THE CLUSTERS FOR EACH USER

User ID	C1%	C2%	C3%	C4%	C5%	Noise %
5927	76.03	2.32	17.33			4.32
5928	31.94	61.84				6.22
5940	27.46	62.47				10.07
5943	86.96	4.35				8.69
5947	72.63	2.05	5.75	3.32	2.56	13.69
5948	68.45	24.47				7.08
5953	70.92	10.46				18.62
5964	54.09	37.74				8.17
5972	79.28	2.62	12.07			6.03
5975	32.26	57.66	4.03	4.44		1.61
5976	61.77	26.77	4.11			7.35
5987	14.37	72.75				12.88
5988	33.17	32.19	9.09	9.09		16.46
5989	35.77	37.91	8.44			17.88
6012	88.3	9.55				2.15
6014	60.43	16.53	6.5	2.98	3.79	9.77
6020	66.15	9.73	7.39	4.67		12.06
6026	37.5	28.26	21.74			12.5
6028	49.47	37.46	5.12			7.95
6032	86.15	5.66				8.19
6035	67.15	30.69				2.16

6038	19.76	35.87	3.44	27.42	2.66	10.85
6045	49.2	35.35	3.37			12.08
6063	88.03	4.21				7.76
6082	5.8	91.67				2.53
6106	44.86	39.51				15.63
6109	62.39	16.28				21.33
6168	13.76	82.07				4.17
6175	82.35	6.62	3.68			7.35
6176	39.75	40.75	11.75			7.75
6178	65.22	24.71	3.98			6.09
6192	65.6	32.48				1.92
6194	60.77	22.88	5.96	7.69		2.7
6197	79.1	19.73				1.17

In Table III we present a summary of our findings regarding the number of distinguished clusters. We observe that users usually communicate from a small number of geographic locations, which is in line with the results presented in [22]. Another important observation from Table II is that for the majority of users with more than two clusters, there is a cluster with a high percentage of calls, ranging from around 60% to 90%. We could assume that this cluster represents the place where they live and spend most of their time. For some users, we can see that there is a second cluster with a significant percentage of calls around 30%, which could be a location that they spend a considerable fraction of their time. A plausible explanation for these users might be that they may be commuting daily to another city in order to work or study. For example, we found that user 6028 has 99% of all communication within cluster 2 (37.46% of her total communication) performed during weekdays. Other examples include users 5976 and 5989 with 98% and 100% of their communication within the second larger cluster taking place on weekdays. Finally, clusters with a very low percentage of calls indicate places that are not often visited and could be indicators of traveling for business or leisure.

TABLE III. SUMMARY OF CLUSTERING RESULTS

Number of Clusters	Number of Users
1	20
2	17
3	10
4	4
5	3
Total	54

B. User Types

Our next step after clustering each user's communication based on location was to analyze the percentage of outgoing calls with the contacts that each user calls per cluster. After this analysis we observed that there are different types of users regarding their outgoing communication patterns per location. In this section we present these user types and some figures representing the respective patterns.

As we have already mentioned, there are users that all their communication was clustered in just one geographic area. These users seem not to have traveled (or not have had mobile phone voice communication when travelling or even not having enough communication in another place to form a cluster) during the period that they participated in the campaign, so we could say that they are characterized by low

mobility for that time interval. Using physical location context to predict the most probable contacts to call will not have any effect on prediction accuracy for this type of users, so they were omitted from further analysis.

Another observed type includes users that while they make phone calls from different clusters, their communication pattern remains practically unaffected when changing physical location, i.e. they call the same contacts regardless of their current location. As a typical example, we can see in Fig. 6 a graph showing the percentage of outgoing calls to each contact broken down by location cluster for the user with ID 6168. This user makes phone calls from two clusters, while from Table II we can see that the majority of this communication for this user takes place within cluster 2. For all contacts that are called in both clusters we observe percentages of outgoing calls that are very similar, irrespectively of the location cluster. Of course, in the cluster with the majority of communication there are calls to contacts that are not called from the other cluster. However, for this user we observe that the percentage of calls to each of these contacts is very low, below 5%. For this type of users one would again expect that physical location context would not add much value to the prediction procedure.

On the other hand, it seems that there is a third user type: users for which we observe important deviations regarding outgoing communication patterns among the different clusters. Such a case is presented in Fig. 7, where we can see the percentage of outgoing calls to each contact per cluster for user 5947. It is obvious that this user has a completely different communication behavior within each cluster. However, this user type could also include users with observed deviations only in some of the clusters, which means that their behavior changes only in some of the places that they make phone calls from. Consequently, one could expect that prediction accuracy might improve for this type of users if the physical location context is considered in a prediction algorithm.




Fig. 6. Percentage of outgoing calls to contacts per cluster for user with ID 6168




Fig. 7. Percentage of outgoing calls to contacts per cluster for user with ID 5947

Wanting to validate our visual observations, we decided to define a metric that indicates variations of the user's communication within the geographic clusters. Since low communication contacts are insignificant for predictions, we excluded those with a percentage of communication below 10% in all clusters. For the remaining contacts, we computed for all cluster pairs the Euclidean distance between the contacts' communication frequency vectors and we set the metric equal to the average of these pair values. A lower value of the metric indicates a lower variation of communication among the different clusters. Next, we applied the k-means algorithm on this metric for our users (excluding those that had not traveled at all) for $k=1,2,3$ and 4. An attempt to run the algorithm for three clusters provided better results by reducing the within-clusters sum of squares. This means that there is another user type between the two that we described in the previous two paragraphs. This type represents users in an intermediate state regarding their communication deviation within different clusters.

A summary of the results of the k-means clustering algorithm is presented in Table IV. Ultimately, our dataset of 54 users is split in 4 user types as follows: 'no mobility' user type has 20 users, 'low communication deviation' user type has 13 users, 'medium communication deviation' user type has 15 users and 'high communication deviation' user type has 6 users.

TABLE IV. NUMBER OF USERS FOR EACH USER TYPE

User Type	Number of Users
No mobility	20
Low communication deviation	13
Medium communication deviation	15
High communication deviation	6
Total	54

C. Prediction Results

In this work, we wanted to examine the effect of physical location context, exclusively, on outgoing call prediction. For this reason, we modified our original prediction model [6][7] to use only frequency as a contextual cue, in order to exclude the effect of recency dimension. For the purpose of establishing a performance baseline, we define the prediction metric to be the

total frequency of outgoing calls for each contact. Such a choice simulates the "frequently called" screen of the contact list application on some Android devices. In order to compare with the case when physical location context is enabled, we perform a second experiment where we replace this dimension with the frequency of outgoing calls with each contact within the cluster that the user is at the time of prediction. In other words, we only take into account calls made from the place where the user is. As explained in the previous section, we would expect a noticeable impact of this contextual cue mainly for users with a deviated communication frequency distribution among different location clusters.

In our original work, we used a temporal training window to assist prediction. This approach would have been problematic here, since it would result in empty training sets for low communication activity clusters that the user had not visited during the training period. To overcome this obstacle, we made a slight modification of the algorithm and switched from a temporal training window to a numerical one, by computing for each user the calls made per day and multiplying with the number of days we wanted to include in our training. Using a training window of 15 days (as in our previous work) we traverse the users' call logs and for each call after this period we simulate the prediction procedure that provides weighted suggestions for presentation in an adaptive interface with 1, 3 and 5 suggestion positions respectively. The prediction accuracy results for each user type group are presented in Table V where for each number of suggested contacts, the left column shows the percentage of successful predictions for the base experiment (B1, B3 and B5 respectively) and the right column shows the impact of physical location as the difference of the two percentages ($\Delta L1$, $\Delta L3$ and $\Delta L5$ respectively). We should note that users who present no mobility were omitted from this table, since there is no difference in prediction accuracy when running the two experimental simulations.

Overall, the inclusion of location context had an average impact differential of 0.09% on our users ($sd=2.90\%$) for 1 suggestion, -0.47% ($sd=2.10\%$) for 3 suggestions and -0.41% ($sd=3.35\%$) for 5 suggestions. However, a statistical analysis shows that these differences are not statistically significant.

When considering the three categories for users that show mobility, we have the following results: For 'low communication deviation' users there is an average impact differential of -1.00% ($sd=2.16\%$) for 1 suggestion, -1.15% ($sd=2.23\%$) for 3 suggestions and -1.30% ($sd=1.75\%$) for 5 suggestions. Results are statistically significant (paired samples t-test) only for the case of 5 suggestions ($t(12)=2.694$ $p=0.02$). For 'medium communication deviation' users there is an average impact differential of 0% ($sd=2.24\%$) for 1 suggestion, -0.20% ($sd=1.86\%$) for 3 suggestions and -1.20% ($sd=2.11\%$) for 5 suggestions. Results are statistically significant (paired samples t-test) only for the case of 5 suggestions ($t(14)=2.201$ $p=0.045$). Finally, for users belonging to the 'high communication deviation' group there is an average impact differential of 2.67% ($sd=4.41\%$) for 1 suggestion, 0.33% ($sd=2.34\%$) for 3 suggestions and 3.50% ($sd=5.68\%$) for 5 suggestions. Results are not statistically significant for any suggestion list size.

As we can see, the results indicate that incorporating the physical location context in the prediction algorithm does not necessarily lead to improved prediction accuracy. When the differences are not statistically significant, the prediction results are not impacted by location. In most cases where the results are statistically significant, the differences are only marginal, while there are users for which the algorithm provides slightly worse or slightly better results. For users that belong to the ‘low communication deviation’ group, with similar call patterns within all clusters this seems normal.

TABLE V. EFFECT OF PHYSICAL LOCATION ON PREDICTION ACCURACY (ΔL_X) ON THE RESULTS OF BASELINE PERFORMANCE (BX) FOR 1, 3 AND 5 SUGGESTIONS

User ID	B1%	$\Delta L1\%$	B3%	$\Delta L3\%$	B5%	$\Delta L5\%$
<i>Low communication deviation users</i>						
5927	91	0	96	-1	97	0
5940	37	2	66	2	74	-1
5948	57	0	77	-3	81	-1
5953	19	-1	34	0	44	2
5964	12	0	28	0	34	-2
5988	61	-5	78	-1	83	-4
6012	71	-1	91	0	95	-1
6082	57	0	81	0	86	-1
6106	19	-4	42	-3	52	-2
6168	32	-1	70	0	89	-1
6176	53	-4	71	-7	77	-5
6192	14	-1	33	0	47	0
6197	37	2	68	-2	79	-1
<i>Medium communication deviation users</i>						
5928	77	-2	93	-1	95	0
5943	11	0	29	0	41	0
5972	19	2	42	0	51	0
5976	32	6	66	2	77	-2
5987	16	2	40	0	56	-1
5989	52	1	71	0	78	0
6014	18	-1	35	2	49	1
6026	28	0	60	-5	75	-7
6035	22	-1	42	2	58	-1
6038	43	-2	82	0	96	0
6045	51	-1	63	1	71	1
6063	35	1	53	0	62	-1
6109	35	0	51	0	60	-1
6178	21	-2	45	-3	58	-4
6194	27	-3	56	-1	70	-3
<i>High communication deviation users</i>						
5947	25	1	51	0	60	1
5975	51	11	67	-3	75	13
6020	22	1	42	3	52	8
6028	49	4	74	3	82	0
6032	58	0	76	0	84	0
6175	15	-1	38	-1	47	-1

Next, we analyze some typical examples, representative of the majority of users existing in the dataset under review, in order to explain the results for users who follow different communication patterns in different clusters (medium and high deviation user types). Firstly, we observe that some of them place the majority of their outgoing calls mainly from one cluster, thus deviations of communication patterns within the other clusters play a minor role in prediction accuracy. Such a typical case is user 5947 (high deviation) which mainly communicates from cluster 1, as shown in Table II. In Fig. 8, we observe that for this user the contact frequency distribution for cluster 1 almost coincides with contact frequency distribution for the total outgoing communication.

A different explanation stands for participants such as 5976 (medium deviation). This participant places phone calls from 3 clusters, however only 2 clusters have significant impact in his communication behavior, since communication from cluster 3 accounts only for around 4% of total communication (see Table II). If we disregard contacts with a call frequency below 10% (grey area in Fig. 9), which are usually difficult to predict from call log analysis, we can see that there are only 3-4 remaining contacts. Calls to such a small group of contacts are easy to predict even without taking into account context like location when providing 3 or 5 suggestions. In this example, this is even clearer as the top contacts for cluster 1, cluster 2 and total communication are the same in a different ranking.




Fig. 8. Frequency distribution of outgoing communication for user with ID 5947




Fig. 9. Frequency distribution of outgoing communication for user with ID 5976

An observation of the call logs of our users and an examination of failures of the prediction procedure when location is used compared to the normal frequency dimension, shows that these failures usually occur when there is a transition from one cluster to another and especially when the destination cluster is not the one with the most communication activity. When location is used, we observed that the training window for the “secondary” clusters, which users do not visit often or spend short time periods within, may contain older calls from previous visits of the user. As a result, this training

window fails to appropriately represent recent communication pattern of users, which provides better results for future prediction [15]. When users stay for a long time (>15 days) within a cluster, the two methods gradually provide the same results (as they both analyze the same training sets to generate suggestions). Consequently, these transition failures are critical for the performance of the algorithm when location is taken into account.

V. CONCLUSIONS

In this paper we examine the impact of physical location as a contextual cue for predicting the next callee, a function useful for the design of adaptive contact list interfaces. As already discussed in the introduction and related work sections of this paper, location is considered in literature to be an important contextual cue for mobile users and is often highlighted as a possible candidate dimension for outgoing call prediction. However, as far as we know, this is the first attempt to assess the effect of location on call prediction and we believe that this is the main contribution of our work.

We showed that using physical location for outgoing call prediction does not necessarily lead to prediction accuracy improvements, while in some cases (5 suggestions, low and medium communication deviation groups) the statistical analysis showed a statistically significant slight decline. We presented our explanation for this fact based on the distinct user types that we distinguished in our dataset. While the results were expected for users that follow similar communication patterns in all places they visit, it was a surprise to note that the prediction accuracy was not improved for users that seem to change their behavior while moving between different areas. We analyzed call patterns for some typical users which belong to this category and we provided our arguments to explain our observations. We found that there were users for which communications within the clusters where they present deviated behavior represent only a small percentage of their total communication and as a result possible prediction accuracy improvements for calls from these places do not affect significantly total prediction accuracy. Other users were found to mainly interact with a small number of contacts, which are easy to predict independently of the addition of location. In this case, interactions with other contacts are infrequent, so these calls usually fail to be predicted. This is in line with the results presented in [4] where it was found that 84% of phone calls are exchanged with the top 5 contacts in terms of frequency, while we have also concluded in similar observations in [23]. Moreover, for the case of suggestion lists with 1 and 3 entries [24] suggests that the most frequently contacted people are those with the higher tie-strength and as such it is normal to be used independently of location.

An observation of users' call logs and a step by step application of the algorithm on them revealed that there are often failures when users are moving from the main communication cluster to other locations, since the training sets used in these 'secondary' clusters fail to capture recent call trends. A possible improvement could arise from a combination of the location dimension with the recency

dimension and this is a direction for future work we intend to follow.

Of course, we cannot unreservedly generalize the conclusions from this analysis. Although at this time there does not exist a more comprehensive publically available dataset on which to test our algorithms, it is probable that our dataset did not contain users that fell into the situations described in the scenarios cited in section III.B. It is possible that physical location plays an important role in communication behavior for people that their lifestyle involves frequent travelling and communication with more contact than the average user. We should emphasize at this point that results are encouraging for users belonging to the high communication deviation group, which counts only 6 members, a number too low to lead to statistically safe conclusions. However, we believe it is important that this analysis indicates that physical location may not be as useful a context dimension for call prediction for the majority of mobile users as one might expect and as literature on this subject suggests. Although location in our work is based on a spatial granularity level that can be assumed to carry semantic interpretation (city), it is possible that finer-grained semantic representations of location (e.g. "home", "work", etc.) might yield better results.

In a previous publication [7] we performed a preliminary experiment to evaluate the effect of temporal context on prediction accuracy. Temporal context can be used to detect patterns of movements and semantic information about location. Again, the results were not encouraging as prediction accuracy slightly decreased. If we combine this conclusion with the results presented in this paper, it becomes apparent that adding contextual dimensions in a prediction algorithm is not an undertaking that can be taken lightly. Our work shows that vector context models require a thorough investigation of the effect of vector dimensions and that the results might be counter-intuitive. This was the reason for proposing a prediction procedure based on frequency and recency of communication, which are two dimensions that we have thoroughly examined [23]. A similar conclusion that confirms our argument can be inferred from [2], where although not directly underlined by the authors, it is cited that not all contextual dimensions used for prediction had the same contribution in the prediction output.

In the future, apart from mixing location with recency, we intend to analyze other available datasets, including one dataset collected in an experiment that we organized in order to evaluate our adaptive user interface. We hope to be able to get further insight regarding the effect of physical location and other contextual cues (semantic location, activity, scheduled calendar events etc.) on call prediction.

ACKNOWLEDGMENT

The research in this paper used the MDC Database made available by Idiap Research Institute, Switzerland and owned by Nokia.

REFERENCES

- [1] M. Böhmer, B. Hecht, J. Schöning, A. Krüger, and G. Bauer. "Falling asleep with Angry Birds, Facebook and Kindle: a large scale study on mobile application usage". In Proc. of the 13th international conference on Human computer interaction with mobile devices and services (August 2011). MobileHCI '11. ACM Press, New York, NY. pp. 47-56.
- [2] S. Lee, J. Seo, and G. Lee. "An adaptive speed-call list algorithm and its evaluation with ESM". In Proc. of the 28th Int. Conf. on Human Factors in Computing Systems (April 10-15, 2010). CHI '10. ACM Press, New York, NY, pp. 2019-2022.
- [3] A. Komninos, and D. Liarokapis. "The use of mobile contact list applications and a context-oriented framework to support their design". In Proc. of the 11th Int. Conf. on Human-Computer Interaction with Mobile Devices and Services (September 15-18, 2009). MobileHCI '09. ACM Press, New York, NY, Article No 79.
- [4] F. Bentley, and Y. Y. Chen. "The Composition and Use of Modern Mobile Phonebooks". In Proc. of the 33rd Annual ACM Conference on Human Factors in Computing Systems (April 2015). CHI '15. ACM Press, New York, NY. pp. 2749-2758.
- [5] O. Bergman, A. Komninos, D. Liarokapis, and J. Clarke. "You never call: Demoting unused contacts on mobile phones using DMTR". J. Personal and Ubiquitous Computing (August 2012), vol. 16 (6), pp. 757-766.
- [6] A. Plessas, V. Stefanis, A. Komninos, and J. Garofalakis. "Using communication frequency and recency context to facilitate mobile contact list retrieval". Int. J. of Handheld Computing Research (October-December 2013), vol. 4 (4), pp. 52-71.
- [7] V. Stefanis, A. Plessas, A. Komninos, and J. Garofalakis. "Frequency and recency context for the management and retrieval of personal information on mobile devices". Pervasive and Mobile Computing (December 2014), vol. 15, pp. 100-112.
- [8] V. Stefanis, A. Komninos, A. Plessas, and J. Garofalakis. "An interface for context-aware retrieval of mobile contacts". In Proc. of the 15th International Conf. on Human Computer Interaction with Mobile Devices and Services (August 27-30, 2013). MobileHCI '13. ACM Press, New York, NY, pp. 492-497.
- [9] I. Arminen. "Social functions of location in mobile telephony". J. Personal and Ubiquitous Computing (August 2006), vol. 10 (5), pp. 319-323.
- [10] L. Barkhuus, L. "Context information in mobile telephony". In Proc. of the 5th Int. Symposium on Human Computer Interaction with Mobile Devices and Services (September 8-11, 2003). MobileHCI '03. ACM Press, New York, NY.
- [11] N. Kiukkonen, J. Blom, O. Dousse, D. Gatica-Perez, and J. Laurila. "Towards rich mobile phone datasets: Lausanne data collection campaign". In Proc. of the 7th Int. Conf. on Pervasive Services (July 13-15, 2010). ICPS '10.
- [12] J.K. Laurila, D. Gatica-Perez, I. Aad, J. Blom, O. Bornet, T. Do, O. Dousse, J. Eberle, and M. Miettinen. "The mobile data challenge: big data for mobile computing research". In Proc. of Mobile Data Challenge Workshop in conjunction with Pervasive (June 18-19, 2012).
- [13] O. Barzaiq, and S. Loke. "Adapting the mobile phone for task efficiency: the case of predicting outgoing calls using frequency and regularity of historical calls". J. Personal and Ubiquitous Computing (December 2011), vol 15 (8), pp. 857-870.
- [14] C. Sun, Y. Wang, J. Zheng, and D.F. Hsu. "Feature fusion for mobile usage prediction using rank-score characteristics". In Proc. of 12th IEEE Int. Conf. on Cognitive Informatics and Cognitive Computing (July 16-18, 2013). ICCI*CC '13. pp. 212-217.
- [15] S. Phithakkittukoon, R. Dantu, R. Claxton, and N. Eagle. "Behavior-based adaptive call predictor". J. ACM Transactions on Autonomous and Adaptive Systems (September 2011), vol. 6 (3), Article No 21. ACM Press, New York, NY.
- [16] J.-K. Min, H.-T. Kim, and S.-B. Cho. "Social and personal context modeling for contact list recommendation on mobile device". In Proc. of Int. Conf. on Web Intelligence and Intelligent Agent Technology (December 9-12, 2008). pp. 381-384.
- [17] V. Srinivasan, S. Moghaddam, A. Mukherji, K. Rachuri, C. Xu, and E.M. Tapia. "MobileMiner: Mining your frequent patterns on your phone". In Proc. of the 2014 ACM Int. Joint Conf. on Pervasive and Ubiquitous Computing (September 13-17, 2014). UbiComp '14. pp. 389-400. ACM Press, New York, NY.
- [18] F. Calabrese, Z. Smoreda, V. Blondel, and C. Ratti. "Interplay between telecommunications and face-to-face interactions: A study using mobile phone data". PLoS ONE (July 2011), vol. 6 (7).
- [19] J. Karikoski, and T. Soikkeli. "Contextual usage patterns in smartphone communication services". J. Personal and Ubiquitous Computing (March 2013), vol. 17 (3), pp. 491-502.
- [20] E.M. LaRue, A.M. Mitchell, L. Terhorst, and H.A. Karimi. "Assessing mobile phone communication utility preferences in a social support network". Telematics and Informatics (November 2010), vol. 27 (4), pp. 363-369.
- [21] M. Ester, H.-P. Kriegel, S. Sander, and X. Xu. "A density-based algorithm for discovering clusters in large spatial databases with noise". In Proc. of the 2nd Int. Conf. on Knowledge Discovery and Data Mining. Eds. AAAI Press, pp. 226-231.
- [22] S.R. Subramanya. "Analysis of mobile phone call data to determine user characteristics and to enhance user experience". In IJCA Proc. of Nat. Conf. on Recent Trends in Computing (April 7-8, 2012). pp. 9-13.
- [23] V. Stefanis, A. Plessas, A. Komninos, and J. Garofalakis. "Patterns of usage and context in interaction with communication support applications in mobile devices". In Proc. of the 14th Int. Conf. on Human-Computer Interaction with Mobile Devices and Services (September 21-24, 2012). MobileHCI '12. ACM Press, New York, NY.
- [24] J. Wiese, J.K. Min, J.I. Hong, and J. Zimmerman. "You Never Call, You Never Write: Call and SMS Logs Do Not Always Indicate Tie Strength". In Proc. of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (March 14-18, 2015). CSCW 2015. ACM Press, New York, NY, pp. 765-774.