Bluetooth Familiarity: Methods of Calculation, Applications and Limitations.

Barry Lavelle, Daragh Byrne, Cathal Gurrin, Alan F. Smeaton, Gareth J.F. Jones
Centre for Digital Video Processing
Dublin City University
Dublin 9, Ireland
cathal.gurrin@computing.dcu.ie

ABSTRACT

We present an approach for utilising a mobile device's Bluetooth sensor to automatically identify social interactions and relationships between individuals in the real world. We show that a high degree of accuracy is achievable in the automatic identification of mobile devices of familiar individuals. This has implications for mobile device security, social networking and in context aware information access on a mobile device.

Categories and Subject Descriptors

H5.3 [**Information Interfaces and Presentation**]: Group and Organizational Interfaces

General Terms

Algorithms, Bluetooth, Familarity.

1. INTRODUCTION AND BACKGROUND

Bluetooth is a short-range wireless protocol by which enabled devices can exchange content. It is increasingly routinely included in a wide variety of devices from home computers to portable laptops, PDAs, mobile phones, keyboards, mice and headphones. Bluetooth, although originally devised to support content exchange, can be used for much more than this, for example, Bluetooth familiarity. By monitoring the presence of nearby devices Nicolai [11,12] popularised the concept of familiarity within the Bluetooth space and the work demonstrated that social context could be drawn from general encounters with devices.

He suggests that there are the main types of devices that are encountered: '*familiar'* devices, '*familiar stranger'* devices and '*strangers*'. A *familiar* device is one belonging to a familiar individual who typically form the core social group within a person's life, friends, family, work colleagues, etc.. A *stranger* device typically belongs to an unknown, or a rarely encountered individual, who is outside the friends/family/work-colleague group. Finally *familiar strangers* (a category proposed originally by Milgram [10]) are devices owned by individuals that you encounter on a somewhat regular basis, for example, people who eat in the same restaurant at lunchtime. Typically you will never have interacted with these people.

In this paper, we discuss the collection of Bluetooth sensor data to determine co-present device familiarity. First, we extend Nicolai's work and outline a more robust mechanism for automatically calculating a measure of familiarity for encountered Bluetooth-enabled devices. We demonstrate this technique to be effective at rating familiarity for encountered devices through experimental means. Bluetooth familiarity has a wide range of applications, relevant to real world mobile interactions including: context-aware information retrieval and content delivery; social networking, and privacy and security for Bluetooth interaction. We explore these applications and as well as some of the limitations of familiarity in the final sections.

2. FAMILIARITY DETERMINATION

In order to automatically assign each encountered device into one of the 3 established familiarity categories we calculate a cumulative score. This score represents a device's presence relative to the other encountered devices and is determined based on duration of presence during short periods each day.

Previous work on familiarity determination [11] only employed a basic metric to determine familiarity. In Nicolai's approach a device becomes familiar after only 5 encounters. While such an approach works well for short periods, it does scale nor does it provide a realistic determination of familiarity in periods longer than a week. Employing this approach in the long-term allows 'strangers' to become erroneously classified as familiar after a small number of infrequent encounters. Our approach to determination seeks to overcome these scalability issues, be more robust and to provide a more accurate representation of device familiarity, useful in more operationally realistic timeframes.

In our technique, each day is divided into short intervals for which a presence score of each encountered device is calculated. Short intervals are used to provide a more representative and comparative measure and to allow for differences in recording span of days or any gaps in data where a logger might be disabled which could affect. The use of intervals will also allow for the calculation to account for varied device discovery frequencies. (e.g. polling every minute would use a larger interval size compared with 10 second polling). The resulting interval scores for a day are then summed and added to previous day's scores, providing a cumulative measure of a device's overall presence. An advantage of using this approach is that the scores do not have to be recalculated for all data in the set but only additions to the set.

2.1 Familiarity Calculation Technique

With the availability of a presence score a determination on familiarity may be made. For this, we use dynamic thresholds to set a point at which a device transition between familiarity categories. Initially a device will begin as a 'stranger' and then as a result of increases in their cumulative score can become 'familiar' or a 'familiar stranger'.

Using a static threshold, as in [11], presents problems as a stranger may become familiar after a number of short infrequent encounters over a long enough timeframe. To prevent this, we utilise a small dynamic cumulative threshold in combination with a base threshold to prevent strangers becoming familiar through longevity of encounters. With each additional encounter the cumulative threshold is incremented with a small value. As such to become considered 'familiar' a device must pass the sum of the static and cumulative thresholds. The static threshold marks the point at which a 'stranger' becomes a 'familiar stranger'.

Our technique is summarised in the below formula:

$$CS_d = \sum_{i=0}^{I} \frac{F_d}{AVG(F)} \times \frac{T_i}{T_d}$$

$$Familiar = CS_d > \alpha + (\beta \times I_d)$$

$$FamiliarStranger = \alpha \ < CS_d < \alpha + (\beta \times I_d)$$

where CS_d = Cumulative Score of Device d; I = a given interval (a specified unit of time at which device presence is determined e.g. 5 mins); F_d = Frequency of encounters of device d within given interval; AVG(F) = Average of all encounter frequencies within given interval; T_i = time in seconds for i: T_d = time in seconds for day d; α = Static Baseline threshold ; β = Cumulative Dynamic Threshold; I_d = Total intervals where device d is present.

2.2 Experiment

In order to validate our approach to determining familiarity among Bluetooth enabled devices, we collected Bluetooth sensor data using a proprietary Java ME Bluetooth logging application which employs the JSR-82 API for device discovery. During operation, the application polls for other Bluetooth enabled devices every 10 seconds. For each encountered device the following data is gathered; the device's hardware address (acting as a unique device identifier), the current '*friendly-name*' and a timestamp of the encounter. The hardware address is used to map a particular device to a user. Periodically the application automatically uploads data via a GPRS connection to a database. By running this application and through the Bluetooth sensor on the mobile device, it was possible to gather the device interactions for an individual device over an extended period of time.

Once the Bluetooth sensor data was collected and using the hardware address as a unique identifier, the encounters for each device were extracted and segmented into appropriate intervals. The familiarity score for each device was then calculated using this data as explained above.

2.2.1 Experimental Procedure

In total, six participants were involved in the experiment, each of which ran the logging application on their phones at all times the phone was operational for a period of 24 days. Typically these people worked in the university environment (regular hours) and interacted with people at work and outside of work during the course of a day. No incentive to participation was made, however any costs for upload of data incurred by the participants were covered. Participants of the experiment were asked to start the logging software when they turned on their mobile devices and to go about their daily routines keeping the phone with them as normal.

On completion of the data gathering process, the participants were presented with an exhaustive web-based list of all the (unique) devices they had encountered, ordered from most encountered to least frequent. Users were asked to select those devices they considered to be familiar. A list of all friendly-names for each device was provided to help in the judgment. From this, we formulated a ground-truth data set for each user, against which we could evaluate the results of our familiarity calculation technique.

2.2.2 Statistical Analysis

From the six participants 210,529 individual device encounters were recorded over a 24 day period. On average, each participant had 35,088 device encounters (min. 6,162, max. 62,803) recorded for 1204 unique devices (min: 243, max: 2233) and recorded for 15 days (min. 10, max. 24). Interestingly, we found the number of encounters for each unique device follows a power-law natural logarithmic distribution as shown in Figure 1.

Power-laws are used in mathematics when one wishes to relate one quantity to the power of another. A power-law implies that small occurrences are extremely common whereas large occurrences are extremely rare [9]. Our initial findings show this is coherent with the familiarity scores of users found for each participant's social network. This finding is in line with what we expected, as typically a user will know only a relatively small group of people very well and the remainder of people encountered over a day will be largely unknown. This natural distribution of devices is also important as it verifies that suitable thresholds can be dynamically calculated as the Bluetooth data set expands. It is also possible that we could use this distribution to tune thresholds to individual differences. This dynamic elicitation of thresholds will become part of our future work.

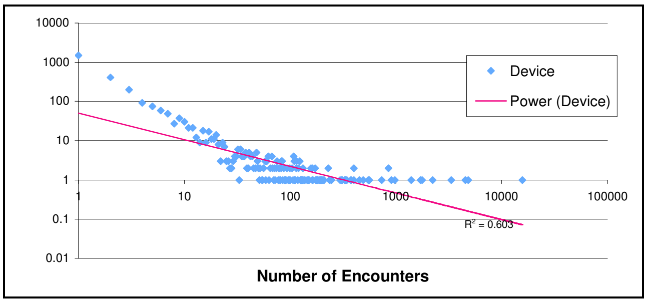

Figure 1. Natural distribution of devices. Only a small number of devices have a high number of encounters while a large number of devices have a small number of encounters.

2.3 Results & Evaluation

As stated, using the time, duration and number of encounters during each day, a familiarity score for each device was calculated. The allocated score was relative to the other devices encountered at given intervals over the course of the day. Precision and recall figures were then calculated against the user specified ground-truth and the effectiveness of our approach measured, though only precision figures are presented here.

To assess the impact that different granularities of time had on the results, familiarity scores were calculated for each device at a range of intervals (including 1 day, 6 hours, 2 hours, 1 hour, 30 minutes, 15 minutes and 5 minutes) to examine the affect of interval size on the accuracy of the resulting familiarity scores. The resulting scores were then sorted in descending order. Precision performance at 20 devices is shown below in Figure 2. Interestingly we found the interval size used had only a marginal effect on the performance of the familiarity scoring for the most familiar devices. Additionally we noted the 5 and 15 minutes offer marginally optimised performance.

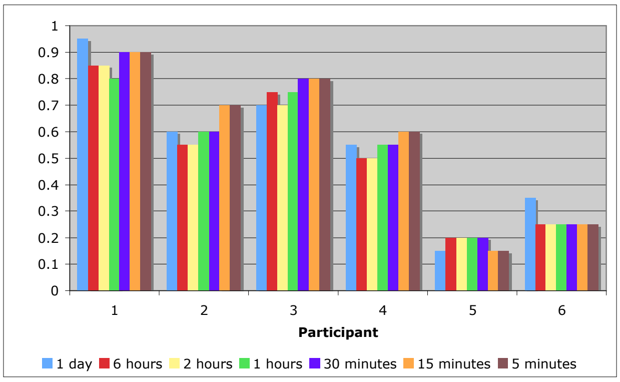

Figure 2: Precision@ top 20 devices

The precision figures for participants 5 and 6 were found to be considerably lower. This can be attributed to the fact that the volume of data recorded was considerably less than the other participants. After analysing their data, we found that they logged infrequently and as a result data was captured sporadically. Based on the analysis of our data we would estimate that between 15,000 - 20,000 device encounters are required before we can reasonably identify friendly devices. The results of this evaluation demonstrate that our methods offer a robust and accurate to determining familiarity of Bluetooth-enabled devices within more operationally realistic timeframes.

3. APPLICATIONS OF BLUETOOTH FAMILIARITY

3.1.1 Security & Privacy

Bluetooth has proven effective as a means of exchange of information, however, it is not without issues of security and privacy. We believe that Bluetooth familiarity can be employed to effectively mitigate against a number of these.

We have previously demonstrated that there is a large overlap in the Bluetooth '*friendly names*' used to label devices, with as much as 25% of names overlapping [7]. Typically when exchanging information between Bluetooth enabled devices, the device discovery is socially mediated so confusion or problems in device identification resulting from name overlap can easily be overcome One can, however, envisage times where it would be helpful to quickly identify how well the person is known to you in order to facilitate faster and more accurate identification of known devices. In Figure 3 we provide an example of how Bluetooth familiarity might be used to aid in this process.

Another issue where familiarity may be of use is '*bluejacking*', a simple exploitation used to send unsolicited messages or files to Bluetooth-enabled devices [14]. This may provide a route by which mobile phone viruses can propagate or by which malicious users can gain full access to all content on the recipient's device [2]. However, this can easily be thwarted through familiarity. Presenting the familiarity score of the sender when a new Bluetooth file transfer request or message is received would enable the device owner to make a better determination of the appropriate action (see Fig. 3). It could also enable the automatic rejection of messages from unfamiliar sources and thereby prevent Bluetooth spamming, which is increasingly being employed in advertising [5] or unwitting access to malicious content [2,14].

Familiarity could similarly be employed within social applications to restrict access to personal phone-based content. For example, BlueTuna [1], an application that enables music recommendations between collocated Bluetooth devices mobile phones, could allow users to restrict access of the music and recommendations to only the most familiar devices .

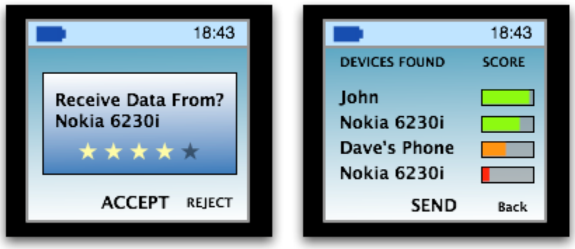

Receive Data From?
Nokia 6230i
★★★★★

John
Nokia 6230i
Dave's Phone
Nokia 6230i

Figure 3. Examples of how familiarity can be employed in current Bluetooth device discovery and message sending.

3.1.2 Social Networking

Using mobile phone context data to infer social relationships is an established concept. Eagle et al. [4] have demonstrated that mobile device data such as application usage, cell tower IDs (i.e. current location), and proximal Bluetooth devices can be used to determine social patterns in daily activity. Bluetooth familiarity is particularly useful in determining social networks of users as our familiarity calculation offers the ability to determine a devices' relative importance based on their frequency and duration of encounters. This can then be used to map the relationships between devices, potentially across multiple users and devices. The resulting social network could be used in a variety of mobile applications, particularly Bluetooth socialising and dating software for example Nokia Sensor [13].

3.1.3 Context Awareness

Typically applications are unaware of the context in which they operate, however, the use of context data in information access promises increases in the effectiveness, ease and relevance of information retrieval. Applications can be made 'context-aware' "*by exploiting the wide range context data available describing the environment, the searcher and the information itself*" [6]. These applications seek to automatically and continuously mine sources of context data for information access, examples of these might include: the location of the user; sensors such as for ambient temperature; biometric or body media devices; and application use and activity. With the widespread availability of mobile devices, wireless access to the internet and the availability of contextual cues from mobile devices, context-aware applications are particularly relevant to the mobile space [3].

Realising context-aware retrieval is however not without its challenges, but we believe that Bluetooth combined with familiarity solves a piece of the puzzle: namely, people-awareness. By continually monitoring Bluetooth activity and associating proximal devices with a familiarity score in real time, the relevant people (determined by familiarity) and the social situation of the owner could be inferred at any given time (e.g. home alone or out with friends). Using this information it will be possible to adapt, alter and adjust the type and amount of information presented to the user in real time. Not only can system output then be tailored but the presentation and interface can also be adapted to the context of use. Combining a real-time understanding of Bluetooth co-presence and familiarity with other sources of context information such as GSM or GPS location or biometric context, will further increase the power, scope and usefulness of Bluetooth-enabled context-aware applications and particularly so in domains such as human digital memory and ubiquitous computing.

4. LIMITATIONS OF BLUETOOTH FAMILIARITY

Bluetooth signals can easily penetrate walls and ceilings without degrading and they also operate in a 3-dimensional plane. This means Bluetooth-enabled devices can be detected in rooms above, below and adjacent to the owner and this presents a challenge for Bluetooth familiarity. From the human perspective, in order for a person to become familiar to us, they must spend time physically proximal to us. Someone sitting in an adjacent room while we are in another will not become familiar to us. Unfortunately due to the nature of Bluetooth, devices in nearby rooms can quickly gain familiarity and there is no easy means to negate against this. This subtle difference between the human and the machine, in what is understood to be proximal has a mild impact for familiarity. Consequently, it can be expected that a small number of devices not personally known by the device owner may be considered extremely familiar when using Bluetooth interaction and co-presence data. These devices typically belong to those working in the same building but not the same physical space as the owner. This difference in the concepts of personal and Bluetooth familiarity should be born in mind when applying our work.

We realise that the collection of Bluetooth activity and presence information raises concerns over privacy. Our work is academically motivated with consent provided by all participants and is designed only to demonstrate the potential of such information. Suitable consideration for privacy should be made when employing Bluetooth presence and familiarity, particularly, if information is not processed or stored on the mobile device, which carries out the logging.

5. CONCLUSIONS AND FUTURE WORK

In this paper we have presented an approach for utilising a mobile device's Bluetooth sensor to automatically identify social interactions and relationships between individuals in the real world. By means of a user experiment, we have shown that by simply running a Bluetooth logger on a mobile phone and analysing interactions with other Bluetooth enabled devices, that an accuracy rate of 90 percent can be achieved in detecting the devices of familiar individuals. This finding has implications for mobile device security, social networking and the application of context awareness.

The focus of our current evaluation has been on establishing the effectiveness of the familiarity scoring for encountered devices and we recognise that additional work is required to determine appropriate thresholds. In our future research, we plan to examine the most appropriate means to establish and optimise the thresholds for familiarity determination. In addition we hope to explore if thresholds need to be tailored to an individual or if a single threshold will generally apply to all devices.

6. ACKNOWLEDGMENTS

We would like thank all members of the CDVP who participated in the experiments. This work was partially supported by the Irish Research Council for Science Engineering and Technology and by Science Foundation Ireland under grant 03/IN.3/I361.

7. REFERENCES

[1] Baumann, S., Jung, B., Bassoli, A. and Wisniowski, M. BluetunA: let your neighbour know what music you like.

[2] Bialoglowy, M. Bluetooth Security Review, Part 1. (Apr. 2005) Retrieved from: http://www.securityfocus.com/infocus/1830

[3] Brown, P.J. & Jones, G.J.F., Exploiting contextual change in context-aware retrieval, Proceedings of the 2002 ACM symposium on Applied computing, Madrid, Spain, pp. 650 - 656, 2002

[4] Eagle, N. and Pentland, A., Reality Mining: Sensing Complex Social Systems, Personal and Ubiquitous Computing, Volume 10, Issue 4, 2006

[5] Graham-Rowe, D. Billboards beam adverts to passing cellphones. Aug 2005. Retrieved from: http://www.newscientist.com/article.ns?id=dn7883

[6] Jones, G.J.F. Challenges and Opportunities of Context-Aware Information Access, International Workshop on Ubiquitous Data Management, pp. 53-62, 2005

[7] Lavelle B., Byrne D., Jones G.J.F., Smeaton A.F. Bluetooth Friendly Names: Bringing Classic HCI Questions into the Mobile Space, British HCI 2007, Lancaster, UK, September 2-5 2007.

[8] Marcus, A., Ferrante, J., Kinnunen, T., Kuutti, K. and Sparre, E. Baby faces: user-interface design for small displays. ACM SIGCHI, Los Angeles, CA, Apr 18-23, 1998.

[9] Mitzenmacher, M. A Brief History of Generative Models for Power Law and Lognormal Distributions, Retrieved from: ftp://ftp.deas.harvard.edu/techreports/tr-2001.html, 2001

[10] Milgram, S. The Familiar Stranger: An Aspect of Urban Anonymity. The individual in a social world, pp. 51-53. Reading, MA: Addison-Wesley, 1977

[11] Nicolai, T., Behrens, N. and Yoneki, E. Wireless Rope: Experiment in Social Proximity Sensing with Bluetooth, PerCom 2006, Pisa, Italy

[12] Nicolai, T., Yoneki, E., Behrens, N. & Kenn, H. Exploring Social Context with the Wireless Rope. In Proceedings of MONET'06, Montpellier, France.

[13] Nokia Europe. Nokia Sensor. Retrieved from: http://europe.nokia.com/A4144923, 2005

[14] Thom-Santelli, J., Ainslie, A. and Gay, G. Location, location, location: a study of bluejacking practices. In Extended Abstracts of CHI 2007. (San Jose, USA, April 2007). ACM Press, New York, NY, 2693 - 2