Eagle, N., and Pentland, A., "Eigenbehaviors: Identifying Structure in Routine", Submitted to: *Ubicomp '05*. September 11-14, 2005. Tokyo, Japan.

Eigenbehaviors: Identifying Structure in Routine

Nathan Eagle and Alex (Sandy) Pentland

The MIT Media Laboratory
20 Ames St. Cambridge, MA 02139
{nathan, sandy}@media.mit.edu

Abstract. In this work we identify the structure inherent in daily human behavior with models that can accurately cluster, analyze and predict multimodal data from individuals and groups. This approach capitalizes on the large amount of rich data on human behavior which we have collected by using mobile phones to continuously log location, people in proximity, and communication of 100 subjects over the course of nine months. We show that it is possible to accurately model many people's lives with just a few parameters – thus allowing accurate prediction of their future behavior from limited observations of their current behavior – as well as to create a similarity metric between individuals and groups that allows accurate identification of group affiliation and behavioral 'style'. We conclude with a discussion of the potential ramifications of eigenbehaviors to the field of Ubiquitous Computing.

1 Introduction

Although human behavior can appear random, typically there are repeating and easily identifiable routines in every person's life. These can be found on a range of timescales: from the daily routines of getting out of bed, eating lunch, and driving home from work; to weekly patterns such as the Saturday afternoon softball games; to yearly patterns like seeing family during the holidays in December. While our ultimate goal is to create a predictive classifier that can perceive aspects of a user's life more accurately than a human observer (including the actual user), we begin by building simple mechanisms that can recognize many of the common structures in the user's routine.

Building models of long-term human behavior has been difficult due to the lack of continuous, rich data, as well as the perceived complexity of an idiosyncratic individual. Additionally, traditional Markov models work well for specific set of behaviors, but have difficulty incorporating temporal patterns across different timescales [5]. We present a new methodology for identifying the repeating structures underlying typical daily human behavior. These structures are represented by a set of vectors of characteristic behaviors called *eigenbehaviors*, principle components of the complete behavioral dataset.

An individual's behavior over a specific day can be approximated by a weighted sum of his or her primary eigenbehaviors. When these weights are calculated halfway through a day, they can be used to predict subsequent behaviors with accuracies for some users of over 90%. This is not only useful as a predictive tool, but also as a method of filling in gaps in the data set when the user turned the phone off. Additionally, groups of interacting people can be clustered into different "behavior spaces" spanned by a set of their aggregate eigenbehaviors. We will show that these behavior spaces can be used to reliably identify the group affiliations of an individual through a simple mathematical transformation described in section 4.2.

1.1 Background work

While behavior is perhaps not as characteristic a signature of an individual as a face, many analogies hold between analysis of an individual's behavior and facial features. Just as digital imaging created a wealth of data to train and test facial analysis tools, the explosive growth of location-aware devices, such as mobile phones, is beginning to enable much more comprehensive computational models of complex human behavior.

Location-aware Behavior Tracking Devices. Although primarily used for location-based applications, electronic badges can also generate rich data on individual behavior within a workplace. The exposed manner in which they are worn allows line-of-sight sensors, such as infrared (IR), to detect face-to-face interactions. Some of the earlier badge work to sense human behavior was done in the 80's and early 90's at EUROPARC and Olivetti Labs [18]. Recent developments in ultrasound tracking have greatly improved the ability to localize the badge, enabling a wide range of just-in-time information applications [16,1].

Outside the office, GPS has been used for location detection and classification [2,11,19], but the line-of-sight requirements prohibit it from working indoors. As an alternate approach, there has been a significant amount of literature regarding correlating cell tower ID with a user's location [3,4,9]. Laasonen et al. describe a method of inferring the significant locations from the cell towers by calculating graph metrics from the adjacency matrix formed by proximate towers. They were able to show reasonable route recognition rates, and most importantly succeeded in running their algorithms directly on the mobile phone [10].

It is a challenge to get accurate location with only the ID of the user's current tower, particularly since towers have a wide range and in urban areas it is not uncommon to be within range of over a dozen. If we were able to get information about all the current visible towers and their respective signal strengths, the location classification problem would become easier, although multipath propagation still makes it difficult to accurately estimate location. We have therefore incorporated use of static Bluetooth device ID as an additional indicator of location that can be used in the same manner as cell tower ID. We have found that use of the BTIDs provides a very significant improvement in user localization, especially within office environments.

This fusion of data is particularly appropriate due to the fact that cellular signals tend to disappear in the middle of large buildings - exactly the place where there are static Bluetooth devices such as desktop computers. On average, the subjects in our study were without reception 6% of the time. During this time with no signal, they spent 21% of it within range of a static Bluetooth device, and 29% near another mobile phone. We expect coverage by Bluetooth devices to increase dramatically in the near future as Bluetooth devices become more common in computers and electronic equipment, so that use of Bluetooth ID may become as important as cell tower mapping for estimation of user location.

Eigendecomposition for Machine Understanding. In machine vision and computer graphics, eigenrepresentations have become one of the standard techniques for many tasks. They are used in face and object recognition [17], shape and motion description [12], data interpolation [13], and computer animation [14]. More recently they have been used in a wide variety of robotic and control applications.

2 Reality Mining: Complex Human Behavior from Mobile Phones

For over a century, social scientists have conducted surveys to learn about human behavior. Surveys, however, are susceptible to issues such as bias, sparsity of data, and lack of continuity between discrete questionnaires. It is this absence of dense, continuous data that also hinders the machine learning and agent-based modeling communities from constructing more comprehensive predictive models of human dynamics. Over the last two decades there has been a significant amount of research attempting to address these issues by building location-aware devices capable of collecting rich behavioral data. While these projects were relatively successful, by depending on a limited supply of custom hardware, they were unable to scale due to limitations on the supply of custom hardware. However, with the rapid technology adoption of mobile phones comes an opportunity to collect a much larger dataset on human behavior [8]. The very nature of mobile phones makes them an ideal vehicle to study both individuals and organizations: people habitually carry their mobile phones with them and use them as a medium for much of their communication. In this paper, we capture all the information to which the phone has access (with the exception of content from phone calls or text messages) and describe how it can be used to provide insight into both individual and collective behaviors.

2.1 The Dataset

Our study consists of one hundred Nokia 6600 smart phones pre-installed with several pieces of software we have developed as well as a version of the Context application from the University of Helsinki [15]. Seventy-five users are either students or faculty in the same laboratory, while the remaining twenty-five are incoming students at the business school adjacent to the laboratory. Of the seventy-five users at the lab, twenty are incoming masters students and five are incoming freshman. The information we are collecting, which includes call logs, Bluetooth devices in proximity, cell tower IDs, application usage, and phone status (such as charging and idle), is primarily obtained from the Context application. The study will generate data collected by one hundred human subjects over the course of nine months and represents approximately 500,000 hours of data on users' location, communication and device usage behavior.[1] Upon completion of the study, we plan to release a public, anonymous version of the data set for other researchers to use [6].

While our users' behavior is quite diverse, one universal similarity is the critical role time plays. To model human behavior over time, we have developed a Hidden Markov Model conditioned on both the hour of day as well as weekday or weekend. A straightforward Expectation-Maximization inference engine was used to learn the parameters in the model, and to perform clustering in which we defined the dimensionality of the state space. After training our model with one month of data from several subjects, we were able to provide a good separation of ({*office*}, {*home*}, {*elsewhere*}) clusters, typically with greater than 95% accuracy. Examination of the data suggests that non-linear techniques will be required to obtain significantly higher accuracy. However, for the purposes of this paper, this accuracy has proven sufficient.

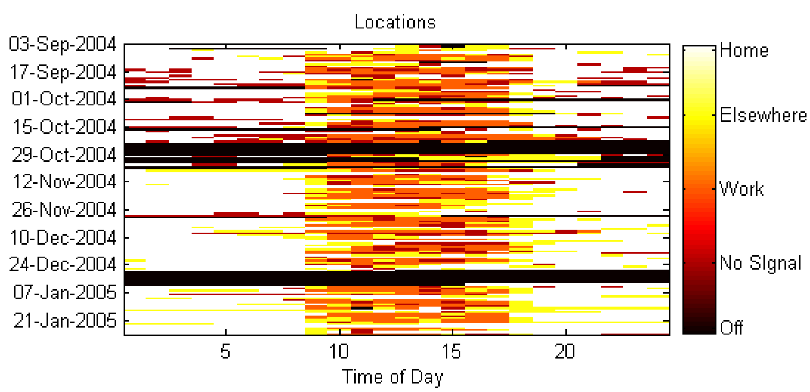

Fig 1. A set of days of a 'low entropy' subject based on celltower ID. The 'hot spot' in mid-day is when the subject is at the workplace.

We attempt to quantify the amount of predictable structure in an individual's life by using an entropy metric. People who live high-entropy lives tend to be more variable and harder to predict, while low-entropy lives are characterized by strong patterns across all time scales. As we will show, low entropy subjects also require few parameters to reasonably model their daily lives. Figure 1 depicts the patterns in cell tower transitions encountered each hour for a 'low entropy' subject. It is clear that the subject is typically at home during the evening and night until 8:00, when he com-

[1] At the time of submission one hundred human subjects have been participating in the study for time periods ranging from two to seven months, representing over 250,000 hours of data. The total duration of the study will be for nine months, and all users will have been enrolled for at least six months.

mutes in to work, and then stays at work until 17:00 when he returns home. Figure 2 displays the number of Bluetooth devices the subject has encountered each hour of the day. For many subjects this number is indicative of people being proximate at work, while for others it is also a sign of socializing after hours.

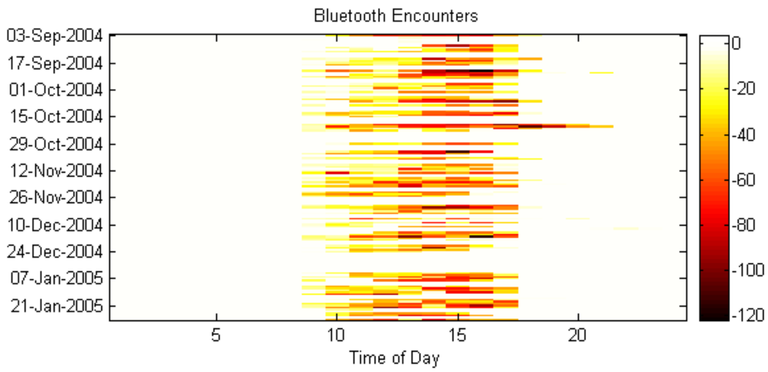

Fig 2. The frequency of Bluetooth encounters for a 'low entropy' subject.

2.2 Privacy Implications

Mining the reality of our one hundred users raises justifiable concerns over privacy. However, the work in this paper is a social science experiment, conducted with human subject approval and consent of the users. Outside the lab we envision a future where phones will have greater computation power and will be able to make relevant inferences using only data available to the user's phone. In this future scenario, the inferences are done in real-time on the local device, making it unnecessary for private information to be taken off the handset. However, the computational models we are currently using cannot be implemented on today's phones. Thus, our results aim to show the potential of the information that can be gleaned from the phone, rather than to present a system that can be deployed today outside the realm of research.

3 Eigenbehavior Analysis

Human life is inherently imbued with routine across all temporal scales, from minute-to-minute actions to monthly or yearly patterns. Many of these patterns in behavior are easy to recognize, however some are more subtle. Although many of life's patterns can be modeled as a Markov process, where the future state depends only on the current state and observational data, these types of models have difficulty capturing correlations that span beyond several time slices. For example, as shown in the second eigenbehavior of Figure 4, sleeping late in the morning for many users appears to be correlated with going out that evening. To capture these characteristic behaviors, we compute the principle components of behavioral data over a set of days and people. We find that these principle components are a set of vectors that span a 'behavior space' and have commonalities with other similar subjects. These vectors are essentially the eigenvectors of the covariance matrix of behavior data and represent a set of features that characterize the variation between people. Each person's behavior data (such as the type shown in Figure 1) contributes in some way to these eigenvectors; and when they are plotted, it is clear that the largest ones are correlated with a type of behavior, such as sleeping in late and going out on the town. It is for this reason we have termed them *eigenbehaviors*.

A linear combination of the eigenbehaviors of a group of people can accurately reconstruct the behavior of each individual in the group. However, the behavior of most people (especially if they work in a collocated group), can be approximated by using only the 'top' eigenbehaviors – the ones that have the largest eigenvalues and account for largest amount of variance in the set of people's behaviors. How well these top eigenbehaviors can approximate an individual's behavior depends on how similar the individual's behavior is to the collective.

3.1 Computing Eigenbehaviors

We initially characterize person *I* by location data shown in Figure 3 as *B(x,y)*, a two-dimensional *D* by 24 array of location information, where *D* is the total number of days in person *I* has been in the study. *B* contains *n* labels corresponding to behavior, where in our case these labels are *{Home, Elsewhere, Work, No Signal, Off}*. To perform the analysis, we transform *B* into *B'*, a *D by 24*n* array of binary values, shown in Figure 3.

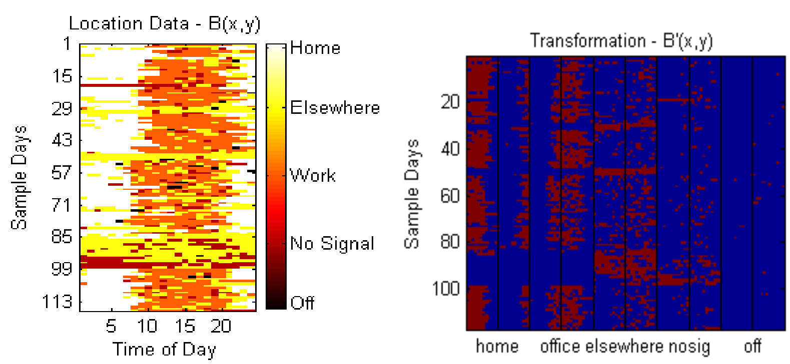

Fig 3. Transformation from *B* to *B′* for data from Subject 4

For these experiments we use *D*=100 days and *n=5*, so that the dimensionality of vector *B′* is 500. This vector represents an individual’s behavior over a single day and can be represented by a point in a 500-dimensional space. A set of *D* days can then be described as a collection of points in this large space.

Due to the significant amount of similar structure in most people's lives, days are not distributed randomly though this large space. Rather, they are clustered, allowing the group to be described by a relatively low dimensional 'behavior space'. This space is defined by a set of vectors of dimension *24*n* than can best characterize the distribution of people's behaviors within the behavior space and are referred to as *eigenbehaviors*. The top three eigenbehaviors that characterize the individual shown in Figure 3, are plotted in Figure 4. The first eigenbehavior corresponds to either a normal day or a day spent traveling (depending on whether the associated eigenvalue is positive or negative). The second eigenbehavior has an eigenvalue that is typically positive on weekends and negative on weekdays, corresponding to the characteristic behavior that sleeping in is correlated with spending that night out somewhere besides home or work. The third eigenbehavior is emphasized when the user is in locations with poor phone reception.

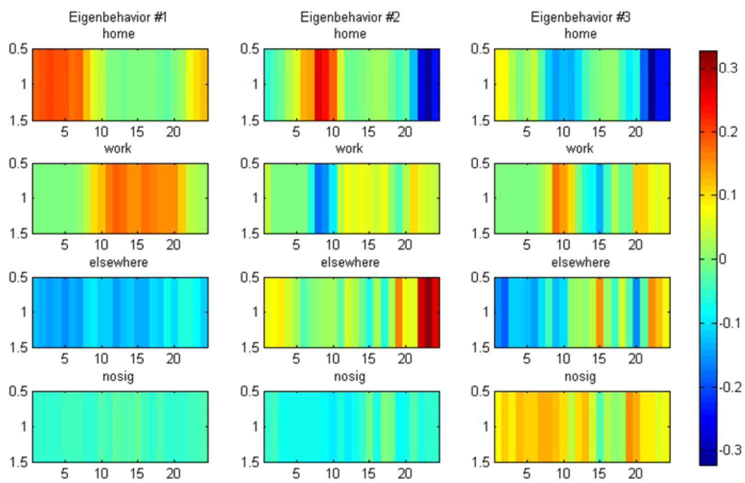

Fig 4. The top three eigenbehaviors for an individual, Subject 4.

Over the course of the Reality Mining study, we have generated a large set of behaviors, $\Gamma_1, \Gamma_2, \Gamma_3 \ldots \ \Gamma_M$, for a group of M people, where M is approximately 100 and individual *i*'s behavior vector, Γ_i, is *D* by *n* by *24*. Following the same notation as Turk and Pentland [17], the average behavior of the group is $\Psi = \frac{1}{M}\sum_{n=1}^{M}\Gamma_n$. And $\Phi_i = \Gamma_i - \Psi$ is the deviation of an individual *i*'s behavior from the mean. Figure 8 shows the different averages for Bluetooth device encounters. Principle components analysis is subsequently performed on these vectors generating a set *M* orthonormal vectors, u_n, which best describes the distribution of the set of behavior data when linearly combined with their respective scalar values, λ_n. These vectors and their corresponding scalars are the eigenvectors and eigenvalues of the covariance matrix of Φ, the set's deviation from the mean.

$$C = \frac{1}{M}\sum_{n=1}^{M} \Phi_n \Phi_n^T$$
$$= AA^T$$

where the matrix $A = [\Phi_1, \Phi_2, \Phi_3, \ldots \Phi_M]$. Each eigenbehavior can be ranked by the total amount of variance in the data for which it accounts, which is essentially the associated eigenvalue. Figure 5 and 6 shows how an individual's behavior can be reconstructed from the top eigenbehaviors. As shown on Figure 7, for 'low entropy' individuals, over 75% of the data can be accounted by simply the first eigenbehavior. Additionally, if the classes of "No Signal" and "Off" are ignored, over 85% of the variance in the behavior of low entropy subjects can be accounted for.

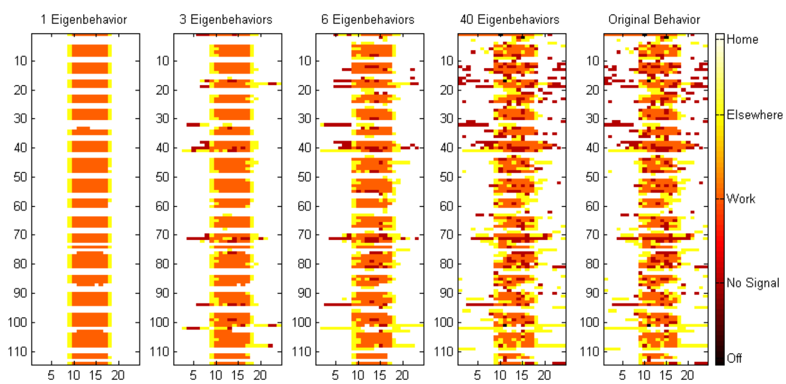

Fig 5. Approximation of a set of locations from Subject 9, a 'low entropy' subject. (120 total eigenbehaviors)

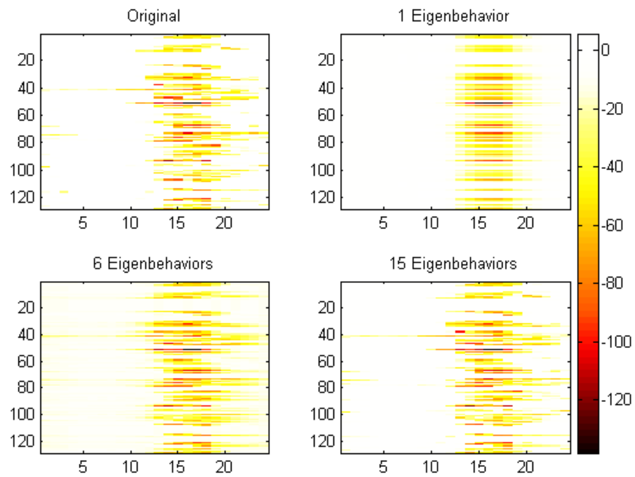

Fig 6. Approximation of a set of Bluetooth devices from Subject 23, a 'high entropy' subject. (24 total eigenbehaviors)

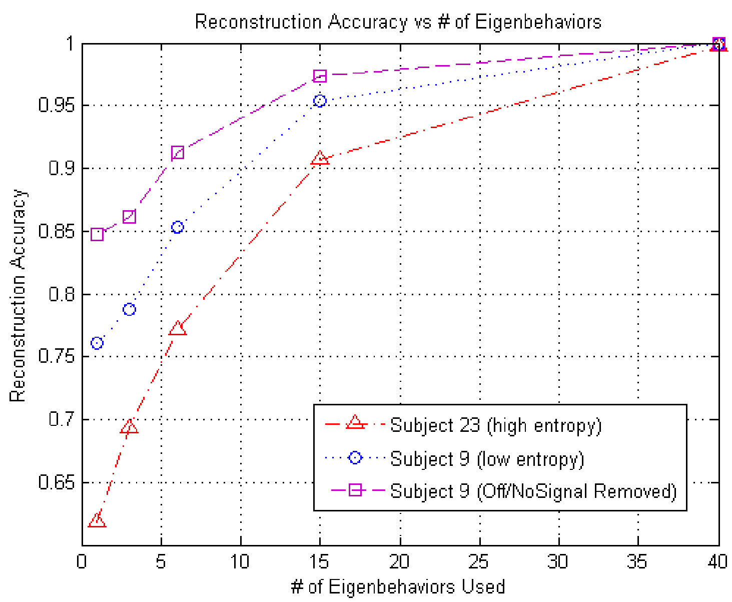

Fig 7. Approximation error for a 'low entropy' subject vs. a 'high entropy' subject

4 Eigenbehaviors of Complex Social Systems

In the previous section we have demonstrated that we can use data from Bluetooth-enabled mobile phones to discover a great deal about a user's patterns of activities by reducing these complex behaviors to a set of principle components characteristic of the individual. In this section we will extend this base of user modeling to modeling complex social systems. By continually logging and time-stamping information about activity, location, and proximity for 100 individuals at an academic institution, the large-scale dynamics of collective human behavior can be analyzed. The eigendecomposition process we have implemented supports a variety of data including a user's trained transition probability matrix from our conditioned Hidden Markov Model, proximity patterns, daily communication activity, motion energy and biometric signals (three of the subjects have been wearing BodyMedia units, collecting galvanic skin response (GSR), acceleration, and heat-flux). For representation purposes, we will show data related to solely Bluetooth proximity events for 3 groups of individuals: incoming business school students, incoming lab students, and senior lab students. Figure 8 shows the mean behaviors for each group, Ψ_j, while Figure 9 depicts the top three eigenbehaviors $[u_1^j, u_2^j, u_3^j]$ of each group.

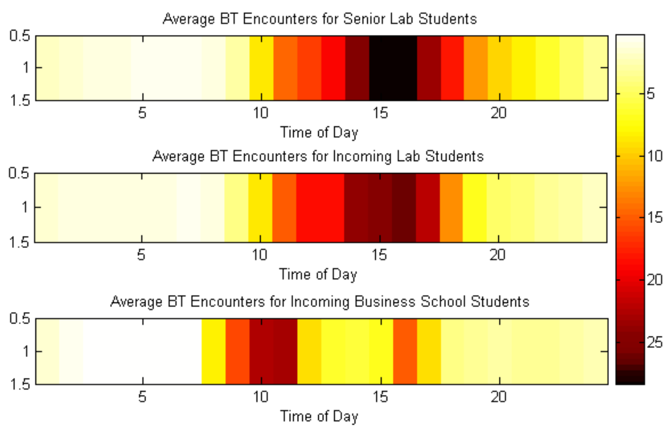

Fig 8. The mean behaviors, Ψ_j, for each group. These values correspond to the number of total encounters with Bluetooth devices over the course of an hour from 12 scans (1 scan/5 minutes).

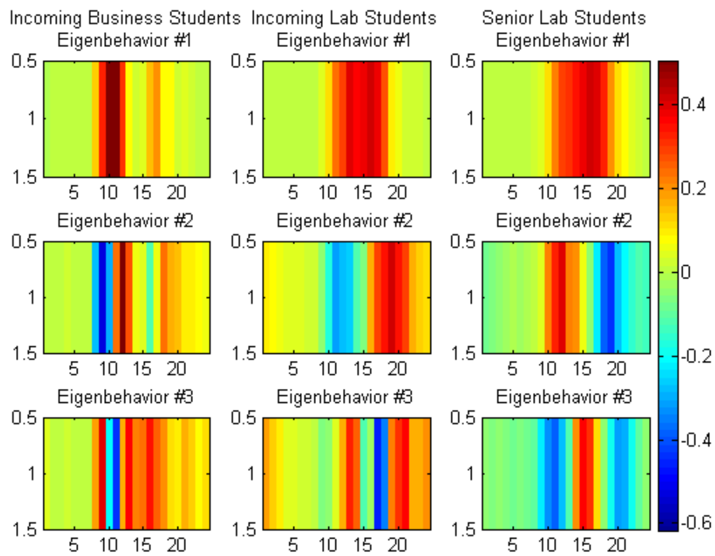

Fig 9. The top three eigenbehaviors $[u_1^j, u_2^j, u_3^j]$ for each group, *j,* comprised of the incoming business school students, incoming lab students and senior lab students

As expected, the top eigenvector in each of the groups corresponds to the mean. For business school students, there is particular emphasis during the school's coffee breaks at 10:30. Besides this emphasis, the other pattern is simply reflective of the standard course times (9 until noon, a lunch break and the subsequently afternoon courses). The lab students have less of an enforced structure on their day. While the entire group of incoming lab students is taking courses, along with approximately half of the senior students, these courses can be selected by the students from anywhere in the institution and typically are not attended by many other subjects. However, each of the lab students has an office within the lab and typically works from there when not in class. While the two groups of lab student share virtually identical principle eigenbehavior, the secondary eigenbehaviors are more telling about the differences. It is common knowledge around the lab that incoming students tend to get overwhelmed by over-commitments to coursework and research leading to late nights at the work-place. This characteristic is emphasized from the group's second and third eigenbehaviors with an emphasis from 20:00 to 2:00.

4.1 Comparing Members of a Group

When the eigenbehaviors are created from the aggregate behavior of a group of individuals, it becomes possible to determine how similar group members are to the mean behavior by just seeing how closely their behavior can be approximated by the groups top M' eigenvectors. Because the Reality Mining dataset contains data for both incoming and senior students, it is possible to verify the onset of concordance between the incoming lab students and the rest of the laboratory. Likewise it is possible to distinguish between different groups of behavior, such as business school students and engineering students. An individual's behavior (Γ) can be projected onto the j group's "behavior space" through the following transformation into the group's eigenbehavior components ($[u_1^j, ..., u_{M'}^j]$) shown in Figure 9.

$$\omega_k^j = u_k^j \left(\Gamma - \Psi_j\right)$$

for $k = 1, ..., M'$ and Ψ_j corresponds to the mean behavior of the group. Ψ_j for Bluetooth encounters of senior lab students, incoming lab students, and business school students is shown in Figure 8.

These weights form a vector $\Omega_j^T = \left[\omega_1^j, \omega_2^j, \omega_3^j, ...\omega_{M'}^j\right]$ which is the optimal weighting scheme to get the new behavior as close as possible to the "behavior space". Each element in the vector gives a scalar value corresponding to the amount of emphasis to place on its respective eigenbehavior when reconstructing the original behavior Γ. By treating the eigenbehaviors as a set of basis behaviors, the vector Ω^T, can be used to determine which person k the individual is most similar to in a particular group, j. We follow the method of Turk and Pentland [17] by using Euclideian distance as our metric for describing similarity.

$$\varepsilon_{j_k}^2 = \left\|\Omega^j - \Omega_k^j\right\|^2$$

where Ω_k^j are the reconstruction weights for the kth person in group j. Figure 10 shows values for ε_j, the distance between one business school student and his peers. Preliminary results show that this distance is correlated with survey responses about the relationship between the users, although further analysis is necessary to ensure statistical significance. [2]

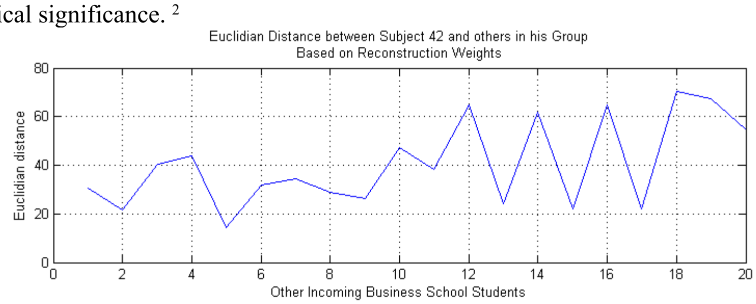

Fig 10. Values corresponding to ε_j, the Euclidian distance between Subject 42 and other incoming business school students.

[2] This method can also be applied to data from a single individual to determine which days are most like the ongoing one. We are starting to use this Euclidian distance metric ε to help predict the subsequent actions of the user.

4.2 Identifying Group Affiliation

Instead of comparing an individual to people within a group, it is also possible to determine how much an individual 'fits in' with the group as a whole by determining the distance ε as the difference between the projection of the individual onto the 'behavior space' of a group, *j,* and the original behavior. We again use Euclidian distance to calculate the difference between the mean-adjusted behavior, $\Phi^j = \Gamma - \Psi^j$ and its projection onto the group's behavior space $\Phi_b^j = \sum_{i=1}^{M_j'} \omega_i^j u_i^j$.

$$\varepsilon_j^2 = \left\| \Phi^j - \Phi_b^j \right\|^2$$

When determining the affiliation of an individual, there can be four possible outcomes, as shown on Figure 11. The dark gray plane represents the group behavior space, containing any set of behaviors that would constitute being part of the group. The first option has the input behavior near on the behavior space as well as proximate to other individuals, Ω_{j_3} , within the behavior space. The second example can be approximated accurately by the behavior space, but there are no other individuals in the same area of the space. Input three appears to have something in common with some members in the group's behavior space, however contains behavioral elements that cannot be reconciled within the behavior space. Lastly, four is a disparate input neither near the behavior space nor any individual in the space.

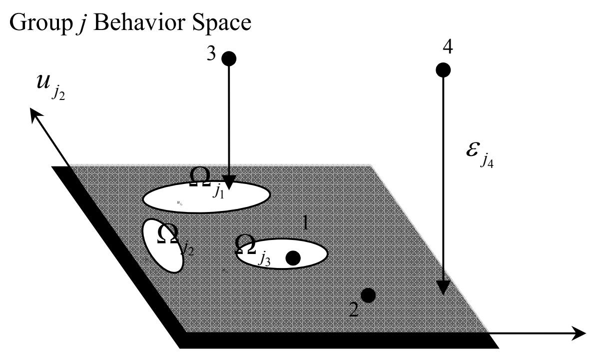

Fig 11. A toy example of group behavior space. Individuals 1 and 2 are on the behavior space and can be affiliated with the group. Projections of Individuals 3 and 4 onto the behavior space do not yield an accurate representation of the two people and therefore are not affiliated with the group.

When classifying users into groups based solely on Bluetooth frequency data shown in Figure 2, this approach works reasonably well. Using six eigenbehaviors to define the business school behavior space, all twenty-five of the business school students are quite proximate to the behavior space. However, as shown in Figure 12, projections of laboratory students are an average of three times further from the business school behavior space. This yields a classification accuracy of 92%. When the behavior space is defined only by the top eigenbehavior, classification accuracy remains a respectable 81%.

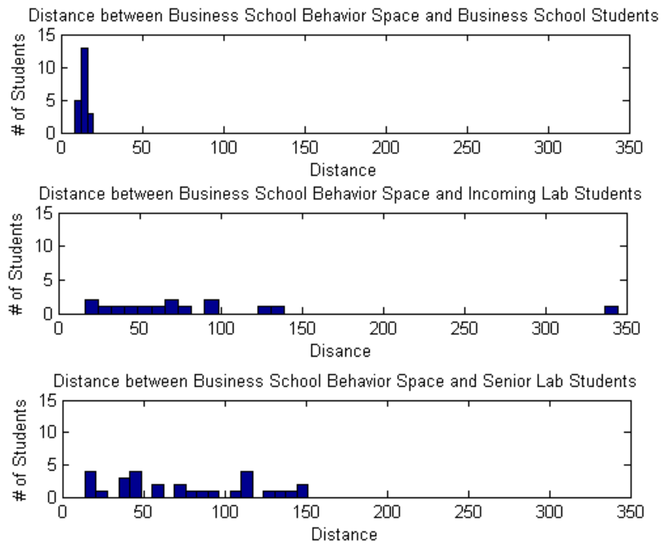

Fig 12. The distance ε_j between the three groups of students and the business school behavior space as defined by its top six eigenbehaviors

5 Eigenbehaviors and Ubiquitous Computing

While we have shown that eigenbehaviors can be used effectively for extracting the underlying structure in the daily patterns of individuals and groups, they also enable a variety of potential applications:

5.1 Usage and Behavior-based Clustering

Currently handset manufacturers sell the same mobile phone to every demographic, from pre-teen to power-executive, to grandmother. If the phones came with a preset behavior spaces corresponding to different demographics, with only a limited amount of usage data, the phone would have the ability to approximate the distance from the user to a given behavior space. By classifying the user into a particular space such as "texting teenager", the phone can harness a much greater set of knowledge than what could have been gleaned from only a few days of standalone behavioral analysis, no matter how sophisticated. With this type of information about the user, the phone should be able to adjust its interface and functionality accordingly [20]. Likewise, these types of clustering can be used in a matchmaking algorithm that incorporates both explicit profile information about a user, as well as implicit behavioral data to identify proximate individuals the user doesn't know, but probably should [7].

5.2 Eigenbehaviors as Biometrics

Just as the eigenvalues associated with a set of eigenfaces are somewhat unique signatures of an individual's face, likewise, the eigenbehaviors can be used to recognize a specific user by characteristic behaviors. Detecting incidents that are far from the user's behavior space could be useful in warning system for the elderly who have boarded the wrong bus, or an automotive alarm that detect can when the owner isn't behind the wheel.

5.3 Data Interpolation

A significant problem that occurs when building models from many human subjects is missing data. On average we have logs accounting for approximately 85.3% of the time since the phones have been deployed. Approximately 5% of this is due to data corruption, while the majority of the missing 14.7% is due to the phones being turned off. However, with a set of these characteristic eigenbehaviors defined for each user, it now becomes possible to generate a rich synthetic dataset from the approximations of the user's eigenvalues over a particular time window of interest. We have shown in initial experiments over 80% accuracy when attempting to generate 5-hour chunks of location data for low entropy individuals. Similarly, this type of interpolation works equally well for behavior prediction.

6 Conclusion

It is inevitable that mobile devices of tomorrow will become both more powerful and more curious about their user and his or her context. We have distributed a fleet of one hundred curious mobile phones throughout an academic campus. We currently have hundreds of thousands of hours of continuous human activity data which requires fundamentally new techniques for analysis. To analyze data of such magnitude and depth, eigendecompositions are useful because they provide a low-dimensional characterization of complex phenomina. This is because the first few eigenvectors of the decomposition typically account for a very large percentage of the overall variance in the signal. As a consequence it becomes easier to characterize complex systems such as groups of people (since there are fewer parameters to learn), easier to analyze the individual and group behavior (since their projection onto the behavior space is low dimensional), and thus easier to predict the behavior of both the system as a whole and the behavior of the individual elements of the system.

These unique properties make eigenbehaviors ideal as a representation of peoples' daily movements, interactions, and their communication behaviors. The low dimensional representation provided by the eigendecomposition will allow us to quickly characterize people, match them to similar people, and predict their behavior in the near future. These capabilities will in turn allow us to build interfaces that can accurately guess the users' preferences, social connections, and their daily plans.

Acknowledgements

The authors would like to express our thanks to Mika Raento for his help modifying the Context application to suit our data logging needs. We also would like to thank Harri Pennanen and Nokia for their generous hardware donation.

References

1. Addlesee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward, A., and Hopper, A. Implementing a Sentient Computing System. IEEE Computer Magazine, Vol. 34, No. 8, August 2001, pp. 50-56.
2. Ashbrook D, Starner T, "Using GPS to learn significant locations and predict movement across multiple users", Personal & Ubiquitous Computing (2003) 7: 275-286.
3. Bar-Noy A, Kessler I, "Tracking mobile users in wireless communication networks," IEEE Transactions on Information Theory, 39(6): 1877-1886, November 1993.
4. Bhattacharya A, Das SK, "LeZi-update: an information-theoretic approach to track mobile users in PCS networks. In: Proceedings of the International Conference on Mobile Computing and Networking, Seattle, WA, August 1999.
5. Clarkson, B., "Life Patterns: structure from wearable sensors", Massachusetts Institute of Technology, September 2002.
6. Eagle, N., and Pentland, A., "Reality Mining: Sensing Complex Social Systems", To appear: J. of Personal and Ubiquitous Computing., (2005)
7. Eagle, N., and Pentland, A., "Mobile Matchmaking: Proximity Sensing and Cuing"', IEEE Pervasive Computing. To appear: April 2005.
8. Himberg J, Korpiaho K, Mannila H, Tikanmäki J, and Toivonen H.T.T. "Time series segmentation for context recognition in mobile devices". In: Proceedings of the IEEE International Conference on Data Mining (ICDM 2001), pp. 203-210, San José, California, USA, 2001.
9. Kim SJ, and Lee CY, "Modeling and analysis of the dynamic location registration and paging in microcellular systems," IEEE Transactions on Vehicular Technology, 45(1):82-90, February 1996.
10. Laasonen K, Raento M, Toivonen H, "Adaptive On-Device Location Recognition", In: Proceedings for Pervasive, pp 287-304, 2004.
11. Liao L, Fox D, Kautz H, "Learning and Inferring Transportation Routines" In: Proceedings for the National Conference on Artificial Intelligence (AAAI-04), San Jose, CA, July 2004.
12. Pentland, A. and Sclaroff, S., "Closed-Form Solutions for Physically Based Shape Modeling and Recognition", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. 13, No. 7., (1991) pp 715-730
13. Pentland, A., "Fast solutions to physical equilibrium and interpolation problems", The Visual Computer. Vol. 8, No. 5-6., (1992) pp. 303-314
14. Pentland, A., and Williams, J., "Good Vibrations: Modal Dynamics for Graphics and Animation", ACM Computer Graphics, Vol. 23, No. 4., (1989) pp 215-222
15. Raento, M., Oulasvirta, A., Petit, R., Toivonen, H., "ContextPhone – A prototyping platform for context-aware mobile applications". IEEE Pervasive Computer. To appear: April 2005.
16. Schilit, B., Adams, N., Gold, R., Tso, M., and Want, R.. "The ParcTab mobile computing system." In Proceedings of the Fourth Workshop on Workstation Operating Systems, pp. 34--39, October 1993

17. Turk, M., and Pentland, A., "Eigenfaces for Recognition", J. of Cognitive Neuroscience. Vol 3, Number 1., (1991) 71-86
18. Want, R., Hopper, A., Falcao, V., and Gibbons, J., "The active badge location system," ACM Transactions on Information Systems, vol. 10, pp. 91--102, Jan. 1992.
19. Wolf J, Guensler R., and Bachman W., "Elimination of the travel diary: an experiment to derive trip purpose from GPS travel data". In: Proceedings from the Transportation Research Board 80th annual meeting, Washington, DC, 7-11 January 2001
20. Weld D, Anderson C, Domingos P, Etzioni O, Gajos K, Lau T, Wolfman S, "Automatically Personalizing User Interfaces", In: Proceedings of the International Joint Conference on Artificial Intelligence, (IJCAI03), Acapulco, Mexcio, 2003.