

Increasing Trust in Personal Informatics Tools

Luis G. Jaimes, Tylar Murray, and Andrew Raij

Department of Electrical Engineering
University of South Florida, Tampa, Florida
{ljaimes,tylarmurray}@mail.usf.edu, raij@usf.edu

Abstract. Personal Informatics (PI) systems help individuals collect and reflect on personal physiological, behavioral and/or contextual data. Typically, these systems offer users interactive visualizations that allow meaningful exploration of the data. Through this exploration, PI systems have great potential to facilitate self-reflection and encourage behavior change.

One of the challenges facing PI systems is a general lack of transparency about the uncertainty, noise or measurement error in the information they display. Data acquisition, processing/inference, and wireless transmission, can each inject errors into the data. However, most PI systems do not provide a way to help users understand what types of errors could be in the data, where these errors come from, and to what extent they can trust the data they see is correct. This paper describes how errors can affect the perception and use of PI data and discusses ways to integrate more transparency into PI systems.

Keywords: personal informatics, data provenance, information visualization, transparent user interfaces.

1 Introduction

The combination of wearable sensors and ubiquitous computation (e.g., smart-phones) is creating a self-surveillance society, where individuals can track a variety of personal behaviors and states to improve their mood[1], health[2], productivity[15], and ecological footprint[6], among others (Figure 1). Complementing these sensing and computation tools are user interfaces that support reflection on such behaviors and states. Taken together, these components - self-tracking sensors, self-report tools, computation, and self-reflection user interfaces - are often called personal informatics (PI) systems [8]. PI systems have emerged as powerful tools in facilitating reflection on behaviors and ultimately, changes in behavior.

While PI systems are powerful tools for behavior change, they generally present opaque user interfaces. More specifically, they process, filter, and aggregate data before it is presented to the user, with the goal of hiding details that may not facilitate reflection on a behavior of interest. This is generally a good practice, as it reduces the complexity of the data and makes it more accessible to the end-user. However, this paper argues there is one category of data hidden by PI user interfaces that should be made more transparent: uncertainty.

Uncertainty (noise or measurement error) can be injected into PI data at several places, including data acquisition, processing, and display. This error or uncertainty, if left hidden from the user, can present two important challenges to the long-term adoption of PI systems. First, if the behaviors presented by a system do not match the user's memory, the user may feel the system is too inaccurate to be useful, and thus may lose trust in the system, and perhaps abandon its use. Second, errors in the data can skew the story the data is supposed to tell users. If users are not aware that errors, and by extension skew, exist in the data, then they could make inappropriate behavior changes based on the data.

This paper aims to address these challenges by proposing two forms of transparency for PI systems: data acquisition transparency (DAT) and data uncertainty transparency (DUT). DAT involves presenting users with the provenance, or history of the data, including the source of the data, and any processes that alter or transform the data until it is presented to the user. By presenting provenance data, we aim to help users identify sources of data errors and better understand the weaknesses of their PI systems. DUT involves modifying PI user interfaces to juxtapose known error or uncertainty with the data itself, allowing end-users to evaluate the quality of their PI data in context. By adding transparency to PI systems, we aim to increase user trust in PI systems and improve decision-making based on observations of noisy PI data.

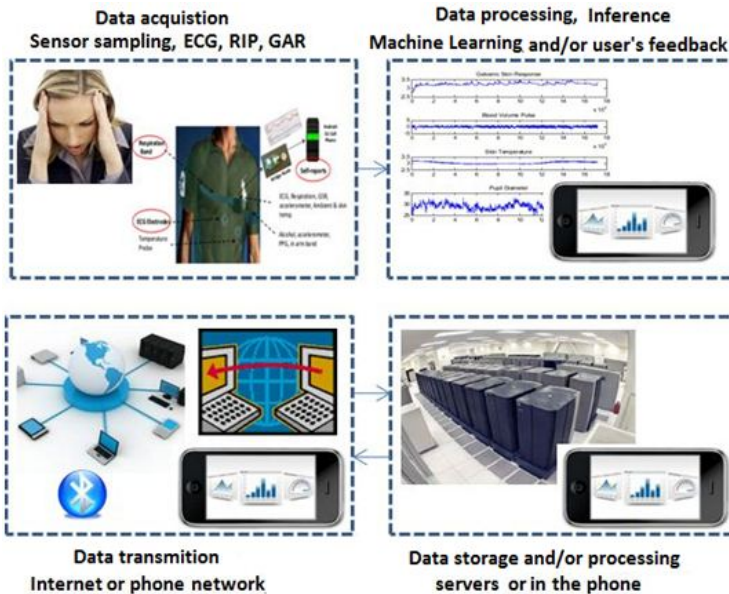


Fig. 1. A generic architecture for personal informatics systems

2 Sources of Error in PI Systems

The path from the acquisition of raw data to the presentation of the data in a PI tool requires several transformations and interpretations of the data, and thus presents several opportunities for error to be injected into the PI data. To demonstrate this problem, we present *Corredor*, a PI system designed to provide feedback on physical activity.

Corredor processes smartphone accelerometry to identify one of three user activities: standing, walking, and running. First, data is sampled from an accelerometer built into the user's smartphone. Second, before passing this data to *Corredor*, the smartphone may apply additional filtering algorithms to reduce the effects of noise. The filtered data is then passed on to *Corredor* for classification. Third, *Corredor* starts the classification process by applying data reduction and smoothing algorithms, eliminating outliers, and imputing missing values. Fourth, frequency-based features are extracted from the signal. Fifth, these features are then passed to a J48 decision tree classifier, which transforms input features to the specific activity the user is performing.

Error can be injected into *Corredor's* data processing pipeline at several places. Amplifiers and ambient temperature changes add noise when the accelerometer is sampled. Filtering and smoothing is critical to making the data useable, but valuable features of the data could be lost in the process. Overzealous outlier elimination could remove good data, while overly conservative outlier elimination might leave too many outliers in the data, both of which could skew classification later in the pipeline. Lastly, machine learning based classifiers (e.g., J48 decision trees) are also prone to error, especially if not enough data of each class was used to train the classifier.

Thus, opportunities for error and uncertainty to play a role in PI system outputs are numerous. Furthermore, as in Li et al's stage-based model of personal informatics [8], errors in earlier stages cascade through the data processing pipelining, affecting the quality of the output at later stages in the pipeline. All of these potential sources of error occur before the user has any interaction with the system, and thus the error and its sources are often hidden from view.

3 Consequences of Error in PI Systems

Figure 2 illustrates scenarios that occur when users interpret and compare the displayed PI data to their memory of their activities. When the user's perception of history differs from what the PI system displays (Figure 2 - left side in red), trust in the system falls (independent of whether the system is correct). On the other hand, when the PI system display matches the user's perceptions (Figure 2 - right side in green), trust in the system is maintained.

In the latter case, we note that trust in the system is maintained even when the data presented by the PI system is inaccurate, because the user has no reason to suspect its inaccuracy. Adjusting behavior based on inaccurate data could have significant negative consequences. For example, if a PI system incorrectly




Fig. 2. User Perceptions vs. PI System Display vs. PI System Accuracy

tells a user that his/her stress level is generally low, and the user believes this assessment is correct, the user will not take action to reduce daily stress, even though he/she may need to do so to avoid heart disease later in life.

In previous work, a lack of trust in a PI system led users to quickly lose interest in the system [9]. Dzindolet et al [4] and Kulesza [7] point out that the frequent occurrence of data errors tends to erode the confidence of users and eventually users abandon the system. Users are often skeptical of black-box systems [3] and may often discount important data, rather than adapt their conceptions and change behaviors. Authors such as Lim and Dey [3] [10] and Kulesza [7] hypothesize that explaining the system behavior to end users will enable them to form better judgments and increase their trust in the system. Their proposals include self-explanatory algorithms to give users a better understanding of the system outputs.

Building on these ideas, we propose two approaches to increasing the transparency of, and trust in, PI systems: 1) data acquisition transparency and data uncertainty transparency. In data acquisition transparency, users are shown provenance models describing how the data is collected and derived. In data uncertainty transparency, information about the accuracy of the data is integrated into the PI system's visualizations.

4 Data Acquisition Transparency

Data provenance refers to metadata describing the derivation history or chronology of a data product [11]. Providing the user with a data provenance model would make the data acquisition history, and by extension the sources of error, more transparent. Users could then determine the causes of error and take appropriate action. The appropriate action could be to abandon the PI system because the error is unacceptable, to continue using the system as is, or perhaps make changes to the system or his/her behaviors to reduce errors (based on their understanding of the model).

The challenge then is choosing the right model and model visualization that can convey the data flow and possible sources of error clearly. We use an extended Open Provenance Model (OPM) proposed by Riboni and Bettini [14] to model

the provenance of *Corredor* (Figure 3). This extension is specifically designed to model components of context-aware and ubiquitous systems, such as sensors, context, smartphones, and inference algorithms.

OPM represents data provenance with directed acyclic graphs (DAGs). Graph nodes represent entities and edges represent causal dependencies. The Riboni and Bettini OPM model incorporates three types of nodes: *artifacts* (i.e., data) represented by ovals; *processes* (i.e., sensors and inference algorithms) represented by rectangles; and *agents*, entities that control processes (e.g., smartphones controlling onboard sensors). Agents are represented by hexagons. Entities are connected to each other by edges describing causal relationships, such as *Was-Controlled-By*, *Used*, *Was-Generated-By*, and *Was-Derived-By*.

Several of these relationships are shown in Figure 3: a *Was-Controlled* edge from the smartphone (agent) to sensors (processes); a *Used* edge from segmentation algorithms (processes) to accelerometer data streams (artifacts); a *Was-Generated* edge from activity predictions (artifacts) to inference algorithms (processes); and a *Was-Derived* edge that relates acceleration features (artifacts) to raw streams of acceleration data (artifacts).

Figure 3 represents a specific moment in the use of *Corredor*, when an individual named John Doe takes a morning run. The model shows that the prediction of John Doe's 7:00am activity (running) was generated by the J48 activity inference process. That prediction was derived from features generated by the Fast Fourier Transform (FFT) process. The FFT process received as input a set of accelerometer magnitudes computed in the time domain by a segmentation process. The segmentation process derived those magnitudes from streams of 3-axis acceleration data provided by the accelerometer sensors on John Doe's smartphone.

While we believe the extended OPM is a useful model to describe provenance, it is not clear if a visual representation of the model (similar to Figure 3) would help end-users understand the sources of error in PI systems. Perhaps collapsing nodes of the graph to reduce its complexity would provide a simpler, more useful view of the data. For instance, we could collapse the FFT and segmentation processes into one process to create a more direct presentation of the relationship between streams of raw data and features.

5 Data Uncertainty Transparency

Data uncertainty transparency refers to integrating uncertainty or error measurements directly into the PI system user interface. This allows users to visually assess the quality of the data they rely on to make behavior change decisions. A critical assumption is that the PI system is capable of assessing the quality of every data point it produces. In the rest of this section, we present several example uncertainty visualizations that could be integrated into PI systems to provide transparency.

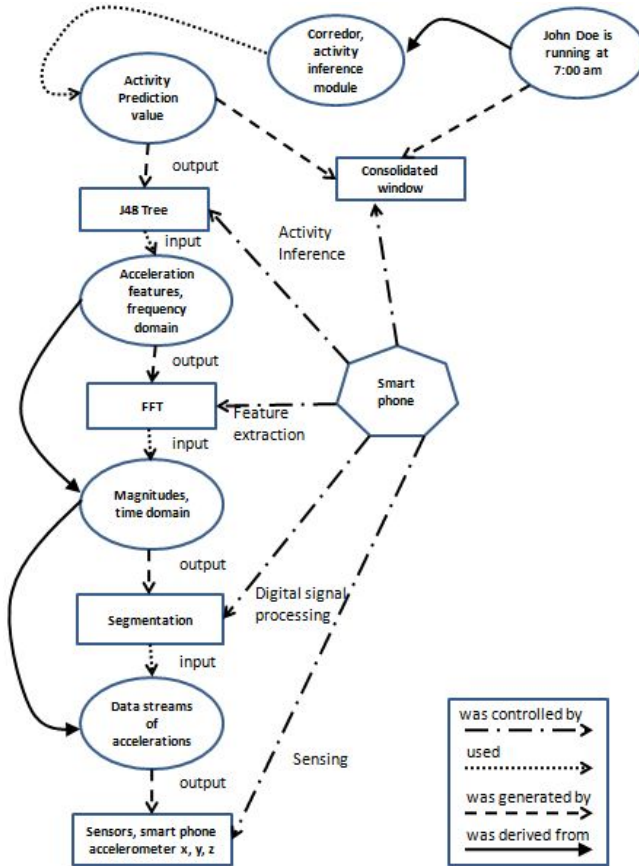


Fig. 3. *Corredor's* OPM extension graph describing detection of John Doe's 7:00am running episode

Figure 4 is a time-series visualization of sample stress predictions from AutoSense [12] [5], a suite of sensors and a smartphone that infers a user's stress level from physiological data. A value of 1 indicates the user is not stressed, 2 indicates the user is stressed, and 3 indicates there is not enough information to make a prediction. However, in the intervals in which predictions are uncertain, the time series is plotted with a dashed line. The size of the dash is inversely proportional to the uncertainty in the data. Larger dashes mean the uncertainty is low, smaller dashes mean the uncertainty is high, and no dashes (a continuous line) mean there is no known error in the data. The user can then interact with this visualization to selectively query the data. Figure 4 shows the result of querying an interval (shown in red) for precise accuracy data.

We present another time-series visualization in Figure 5, this time using data collected by *Corredor*. As in Figure 4, the horizontal axis is a time span (800 seconds), and the vertical axis represents different activities *Corredor* can classify: standing, walking, and running. The three activities are shown in red, green,




Fig. 4. Visualizing data uncertainty in a PI system for stress monitoring and reflection




Fig. 5. 800 seconds of activity data produced by *Corredor*. Uncertain activity classifications are shown in gray, and are positioned between the two classifications most likely to be correct.

and blue, respectively, and are also separated spatially. In addition to the three activities, *Corredor* assesses its confidence in each activity classification. Only those classifications above a confidence threshold σ are shown in red, green, or blue. The remaining uncertain classifications are shown in gray. They are positioned spatially between the two activities that are most likely to be correct. By reviewing Figure 5, a user can see the frequency of uncertain classifications for each activity, and thus better understand what types of activity are likely to confuse the system.

Figure 6 uses a pie chart to present an aggregated version of the *Corredor* data in Figure 5. The pie chart shows the percentage of data classified into each activity and the percentage of data for which confidence in the classifications was low. The latter are spatially positioned between two activity classes to indicate which two activity classes are most likely to be correct in these low-confidence situations. One piece of information present here that is not in Figure 5 is the distribution of low-confidence classifications among the two possible correct classes. These distributions are demonstrated visually by subdividing the low-confidence sections of the pie chart proportionally to the percent that fall in each class. For example, in Figure 5, 12% of the classifications are considered low-confidence. Of those low-confidence classifications, 8% could either be classified as walking or standing. Furthermore, 5% is more likely to be walking and the remaining 3% is more likely to be standing.




Fig. 6. A pie chart summarizing the percentages of activities classified by *Corredor* with high and low confidence




Fig. 7. An avatar-based iconographic visualization of the activities detected by *Corredor*

Lastly, Figure 7 presents a third approach to presenting *Corredor's* data to the end-user, using avatars as icons that demonstrate the three activity classes. Percentages of each activity detected are shown below the activity's corresponding avatar icon. Visual aggregations of low-confidence classifications are positioned between the two corresponding activity icons and are depicted as the two activity icons layered on top of each other. Transparency and color blending are used to indicate the extent to which each low confidence classification could be one class or the other.

6 Conclusion and Future Work

In this work, we presented an approach to increasing the level of transparency, and by extension trust, in personal informatics (PI) systems. We introduced data acquisition transparency as the use of provenance model visualizations to explain how data is produced by a PI system and help users pinpoint sources of error in PI data. We also introduced data uncertainty transparency as the integration of error and uncertainty measures directly into the UI elements already presented in many PI systems. This integration of error and uncertainty into the UI leverages

the juxtaposition of data with error to help users understand what types of errors can occur and whether they should be a concern.

The next step in this research effort is the systematic study of both data acquisition and uncertainty transparency in real PI systems. Most importantly, these studies need to verify that transparency does indeed improve user trust in PI systems. If this is indeed true, then next steps should focus on identifying the most transparent, intuitive visualizations of data provenance and uncertainty. This is particularly critical for PI systems, where the end-user is often the 'average' person rather than a data scientist with significant experience interpreting visualizations [16,13].

References

1. Church, K., Hoggan, E., Oliver, N.: A study of mobile mood awareness and communication through mobimood. In: *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries*, pp. 128–137. ACM (2010)
2. Consolvo, S., McDonald, D.W., Toscos, T., Chen, M.Y., Froehlich, J., Harrison, B., Klasnja, P., LaMarca, A., LeGrand, L., Libby, R., et al.: Activity sensing in the wild: a field trial of ubifit garden. In: *Proceedings of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems*, pp. 1797–1806. ACM (2008)
3. Dey, A.K.: Understanding and using context. *Personal and Ubiquitous Computing* 5(1), 4–7 (2001)
4. Dzindolet, M.T., Peterson, S.A., Pomranky, R.A., Pierce, L.G., Beck, H.P.: The role of trust in automation reliance. *International Journal of Human-Computer Studies* 58(6), 697–718 (2003)
5. Ertin, E., Stohs, N., Kumar, S., Raij, A., Al' Absi, M., Shah, S.: Autosense: unobtrusively wearable sensor suite for inferring the onset, causality, and consequences of stress in the field. In: *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, pp. 274–287. ACM (2011)
6. Froehlich, J., Findlater, L., Landay, J.: The design of eco-feedback technology. In: *Proceedings of the 28th International Conference on Human Factors in Computing Systems*, pp. 1999–2008. ACM (2010)
7. Kulesza, T.: An explanation-centric approach for personalizing intelligent agents. In: *Proceedings of the 2012 ACM International Conference on Intelligent User Interfaces*, pp. 375–378. ACM (2012)
8. Li, I., Dey, A., Forlizzi, J.: A stage-based model of personal informatics systems. In: *Proceedings of the 28th International Conference on Human Factors in Computing Systems*, pp. 557–566. ACM (2010)
9. Lim, B.Y.: Improving trust in context-aware applications with intelligibility. In: *Proceedings of the 12th ACM International Conference Adjunct Papers on Ubiquitous Computing-Adjunct*, pp. 477–480. ACM (2010)
10. Lim, B.Y., Dey, A.K.: Investigating intelligibility for uncertain context-aware applications. In: *Proceedings of the 13th International Conference on Ubiquitous Computing*, pp. 415–424. ACM (2011)
11. Moreau, L., Clifford, B., Freire, J., Futrelle, J., Gil, Y., Groth, P., Kwasnikowska, N., Miles, S., Missier, P., Myers, J., et al.: The open provenance model core specification (v1. 1). *Future Generation Computer Systems* 27(6), 743–756 (2011)

12. Plarre, K., Raij, A., Hossain, S.M., Ali, A.A., Nakajima, M., Al' Aabsi, M., Ertin, E., Kamarck, T., Kumar, S., Scott, M., et al.: Continuous inference of psychological stress from sensory measurements collected in the natural environment. In: 10th International Conference on Information Processing in Sensor Networks (IPSN), pp. 97–108. IEEE (2011)
13. Pousman, Z., Stasko, J.T., Mateas, M.: Casual information visualization: Depictions of data in everyday life. *IEEE Transactions on Visualization and Computer Graphics* 13(6), 1145–1152 (2007)
14. Riboni, D., Bettini, C.: Context provenance to enhance the dependability of ambient intelligence systems. *Personal and Ubiquitous Computing*, 1–20.
15. Santos, J.L., Govaerts, S., Verbert, K., Duval, E.: Goal-oriented visualizations of activity tracking: a case study with engineering students. In: *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge*, pp. 143–152. ACM (2012)
16. Shneiderman, B.: Session 2 infovis for the masses. *IEEE Transactions on Visualization and Computer Graphics* 13(6)