

Salud!: An Open Infrastructure for Developing and Deploying Health Self-management Applications

Yevgeniy Medynskiy, Elizabeth Mynatt

GVU Center

Georgia Institute of Technology

Atlanta, Georgia, USA

eugenem@gatech.edu

Abstract—Health self-management strategies have been shown to improve health outcomes and decrease the cost of healthcare for individuals living with chronic illness. Personal health informatics applications, which allow individuals to track and review personal health-related information have, in turn, been shown to improve health self-management outcomes by supporting learning and reflection. We identify the need for a software infrastructure that facilitates the creation and deployment of such applications and describe our progress in developing and making available such an infrastructure, which we call Salud! We discuss how the design of Salud! follows from key strategies and results in the health self-management literature, and how its application programming interface (API) and related services can be leveraged by other researchers wishing to build a personal health informatics applications. Finally, we describe our own health self-management application that we are building using Salud!

Keywords—health self-management; health informatics; personal informatics; web services

I. INTRODUCTION

Health self-management education seeks to empower individuals living with chronic illness to improve their quality of life and health outcomes by making informed decisions about actions and behaviors that affect their health. Individuals taking part in such educational programs are taught problem-solving skills which, combined with the disease-specific information and technical skills taught in traditional patient education, enable them to identify health problems and take the actions necessary to overcome them [1]. Health self-management strategies have been shown to be effective in improving the health outcomes of individuals living with diabetes, asthma, arthritis and other chronic illnesses [1, 2, 3].

Various health self-management strategies help individuals understand relationships between their health status and behavioral and/or environmental factors: diet, physical activity, air quality, and so forth. As such, personal informatics software that helps individuals track and visualize this type of information can help improve their health outcomes and quality of life. Such tools, which include network-aware medical devices, online meal diaries, digital pedometers and a wide variety of other hardware and software, have received a great deal of attention from academia and industry in recent years. Research has shown that using personal health informatics

applications can improve health outcomes and provide educational benefits to individuals participating in health self-management education programs [4, 5].

While there are general health self-management strategies which are effective in improving health outcomes for individuals living with a range of chronic conditions [3], these strategies must nonetheless be made relevant to specific chronic conditions. Similarly, personal health informatics applications based on health self-management principles are generally targeted toward a specific chronic condition, such as diabetes or asthma. We propose that an open, freely available infrastructure for personal health informatics applications would benefit researchers in this space by providing a ready-made backend to facilitate the development and deployment of a range of end-user applications. Our team is developing such an infrastructure, titled Salud!, which we are making available free of charge and with documentation and support, at <http://salud.cc.gatech.edu>.

In this paper, we identify several key strategies and results from the health self-management literature that have guided the design of the Salud! infrastructure and describe how these results are integrated into the services which Salud! provides. We also present an end-user personal health informatics application built on top of Salud!'s application programming interface (API). This application allows individuals to record personal health and wellness data using a variety of input methods, as well as review and analyze the collected data using simple visualization tools. Finally, we discuss how we are further extending the Salud! infrastructure and API through ongoing design and development.

II. HEALTH SELF-MANAGEMENT STRATEGIES

Individuals living with chronic illness can achieve better quality of life if they are well-informed about available trade-offs between health outcomes and desired lifestyle, and have the means necessary to execute those actions which best satisfy their desired life and health goals [1]. One way to work towards this outcome is to establish the individual as a stakeholder in her or his healthcare. This way, individuals and their healthcare providers become partners and both take responsibility for the individual's wellbeing. Both individuals and their healthcare providers can identify problems, hypothesize solutions, and share responsibility for outcomes. Health self-management education is one strategy that can be used to achieve such a

This work is sponsored by NSF Grant #0915934 and a Google Research Award. Some equipment provided by Georgia Tech Research Network Operations Center (GT-RNOC).

patient-doctor relationship, by empowering individuals to take responsibility for their health. In health self-management education programs, individuals living with chronic illness are taught problem-solving skills, as well as the disease-specific information and technical skills taught in traditional patient education programs. Together, these of skill sets empower patients to recognize personal health problems and make changes in everyday behaviors, such as diet and exercise, that are necessary to mitigate them. The physician and other healthcare providers play a key role in this arrangement, but the patient may perform some of the day-to-day management of her or his chronic illness.

In a literature survey, Bodenheimer and colleagues found many studies of different chronic conditions that showed improved outcomes for patients receiving additional health self-management education, compared to control groups receiving only traditional patient education [1]. These included eleven out of twenty three studies of asthma and twelve of eighteen studies of arthritis. A separate review of adult asthma literature done by Gibson and colleagues [2] concludes that health self-management education also improves outcomes for patients who receive it, including lowering the frequency of hospitalization, unscheduled physician visits, days off work, and nocturnal asthma. The same authors also conducted a review of information-only (traditional) asthma education studies, and concluded that this type of educational curriculum has little or no effect on health outcomes of adults with asthma [6]. These studies suggest that self-management education can be a beneficial component of an educational curriculum for adults with asthma.

The Chronic Disease Self-Management Program (CDSMP), developed and tested in a series of large-scale studies by Lorig and colleagues [3, 7, 8] makeup of a successful self-management education program. The CDSMP is a community-based patient self-management education course that is taught to heterogeneous groups of chronically ill individuals, many with comorbid chronic conditions, over the course of several weeks. One of the underlying assumptions of the CDSMP is that “patients with different chronic diseases have similar self-management problems and disease-related tasks” [3:6] therefore the program is designed to cover topics that apply to a wide range of chronic conditions: exercise, use of cognitive symptom management techniques, nutrition, use of medications, problem-solving, decision-making, etc. To test the effectiveness of such a generic educational program, the CDSMP was initially tested in a large-scale, controlled trial. Participants were individuals 40 years of age or older who had been diagnosed with a chronic lung disease, heart disease, stroke, and/or arthritis (the average participant had been diagnosed with 2.2 of these chronic conditions). The treatment group showed significant improvement on a range of measures of positive health behaviors (including exercise and cognitive symptom management), health status (including self-rated health, energy/fatigue and disability), and health services [7]. Importantly, the authors found that the four disease categories showed no significant interactions for any of the outcome variables. The changes in outcome scores were similar in all four diagnostic groups—implying that the program was equally

effective for participants regardless of their underlying condition.

The overall success of the Chronic Disease Self-Management Program regardless of chronic condition [3, 7], medium of delivery (in-person or online) [8], or group facilitator background (lay or professional) [7] suggests that the structure and content of the curriculum is the key to its effectiveness. This curriculum is based on strategies for enhancing *self-efficacy*—an individual’s expectation of her ability to produce a desired result [9]. Lorig et al. cite weekly action planning, behavior modeling, problem-solving, reinterpretation of symptoms and brainstorming of different management techniques—all both individually and with the group—as activities undertaken by program participants and facilitated by group leaders in order to enhance participants’ self-efficacy [7].

Bodenheimer and colleagues also identify the use of an *action plan* as one common factor shared by the successful educational interventions in the studies they surveyed [1]. An action plan is a short-term goal agreed upon by an individual and a healthcare provider, and which the patient feels confident she or he can achieve. Seven out of the eleven asthma studies surveyed that made use of an action plan showed improved health outcomes, while only four of twelve of the asthma studies that did not use an action plan showed improved outcomes. In the treatment of arthritis, the effectiveness of an action plan is even more pronounced: all ten studies that offered self-management education and made use of an action plan showed improvement over a control group, while only two of the eight studies in which an action plan was not used showed improvement. Gibson and colleagues also cite the inclusion of an action plan into a self-management educational program for adult asthmatics as increasing the overall impact of the program [2].

By focusing on problem-solving and goal-setting, successful health self-management programs enable individuals living with chronic illness to identify and correct behaviors or habits which they judge to be inconsistent with their life and health ambitions. Through sound problem-solving, individuals can judge their condition and set personal goals that more optimally balance their quality of life with their desired state of health. The goal-setting portion of the self-management curriculum then provides a guide for how goals can be met and revised, until desired outcomes are achieved.

III. RELATED WORK IN PERSONAL HEALTH INFORMATICS

Personal informatics applications—systems that allow their users to track and review personal information such as diet, exercise, spending habits and so forth—are also valuable tools for health self-management. Health self-management programs usually ask participants to keep records of their activities, physiological variables and other health-related data, and personal informatics applications can make this process simpler and easier. There have been a number of threads of previous research into improving health outcomes or supporting health self-management through health-focused personal informatics applications.

Mamykina and colleagues have shown that collecting records of diet, exercise and blood glucose readings, and sharing these with a diabetes educator results in improved health outcomes and a move toward internal locus of control among individuals newly-diagnosed with diabetes [5]. Individuals who used Mamykina's systems were more reflective about their condition, which led them to improved self-understanding and a greater sense of personal control over health outcomes [5, 10]. This work supports the claim that an available history of one's own physiological data can support health self-management. Frost, Smith and colleagues have also developed applications to help individuals living with diabetes to capture records of health-related activities [4, 11]. These systems were designed to help individuals reflect on their actions and experience, and were effective at generating reflection and discussion in communal settings (e.g. wellness classes and support groups).

There is also a great deal of commercial interest in personal informatics applications for health. Several companies offer physical devices that allow individuals to track physiological data, which is then made available for review online. The most common type of device is a pedometer, which may be stand-alone (e.g. Fitbit¹) or integrated into a multi-purpose device like the iPhone (e.g. iTreadmill²). Devices to track a variety of other physiological variables are available as well (e.g. bodybugg³). Other companies provide desktop and/or web-based interfaces to track health-related activities and metrics. The majority of these tend to focus on a particular health or wellness issue such as diet or exercise (e.g. Weight Watchers Online⁴), though several open-ended applications that can be used to track a variety of variables do exist (e.g. TheCarrot.com, Grafitter⁵). While these applications promote their value for increasing motivation and helping users make sense of complicated health/wellness conditions, at this time only anecdotal evidence for these claims is available.

Another group of related commercial projects is the personal health record (PHR) platforms currently being developed by companies such as Google⁶ and Microsoft.⁷ These platforms seek to centralize an individual's health record and make it available to a host of third-party services that run on the platform. Our infrastructure differs from these systems in several key ways. First, we are specifically looking to support health self-management applications, rather than the much broader range of applications supported by PHR platforms; this allows us to provide a more targeted and effective design for this class of applications. Second, in addition to an API-accessible data management infrastructure, we provide a suite of ready-made data entry services for end-users. However, we consider the Salud! infrastructure to be complementary to existing PHR platforms and are exploring options to more tightly integrate our system with them.

1 <http://www.fitbit.com>
2 <http://www.itreadmill.net>
3 <http://www.bodybugg.com>
4 <http://www.weightwatchers.com>
5 <http://www.grafitter.com>
6 <http://health.google.com>
7 <http://www.healthvault.com>

IV. THE SALUD! INFRASTRUCTURE

The burgeoning interest in personal health informatics applications led us to design and develop an open web services infrastructure that would allow interested researchers to more easily develop and deploy such applications. We strive for two complementary goals with this infrastructure:

1. To provide a set of easy-to-use web services that incorporate best practices and common paradigms for personal health informatics applications.
2. To provide a flexible, extensible set of personal informatics tools that can be customized by researchers to support a broad range of health self-management interventions.

To achieve these goals, we have designed a software infrastructure, which we call Salud!, in line with design principles culled from the literature about health self-management and personal informatics applications discussed in the previous sections. We are making Salud! available to the broader research community as a set of web service APIs. In this section, we first describe the data architecture made available to users of the Salud! infrastructure and then discuss the design principles which guided our work.

In the remainder of the paper, we use the terms "researchers" and "developers" to refer to individuals who create applications that use the Salud! API and the term "end-users" to refer to the users of those applications.

A. Overview of Salud! Architecture

The Salud! infrastructure and associated API allow researchers to manage end-users and end-users' data for a custom personal informatics application. To better illustrate how the infrastructure works, we will describe how a fictitious, simple diabetes-oriented personal informatics application, DiabetesApp, could integrate with it. This application will allow an individual to register for an account, record her blood glucose reading, record one other variable she believes may be affecting her blood glucose levels (e.g. physical exercise), and review the historic values of the two variables overlaid on a graph.

All data stored on the Salud! infrastructure is associated with a particular end-user's account. Accounts can be created by end-users themselves through a simple registration page, or can be created programmatically by making an API call to the user management interface. Data in an account is organized into structures called Logbooks. A Logbook is a named collection of timestamped data points. In the DiabetesApp example above, a user's account would contain a Logbook for blood glucose measurements and a Logbook for measurements of the other tracked variable. Each data point in a Logbook is called an entry. The simplest kind of Logbook consists of entries which only contain timestamps—for example, the approximate times for when the account's owner went to the gym.

In addition to timestamps, entries in a Logbook can contain other values, including numbers, text, and pictures. The number and type of these additional values are defined by a Logbook's

columns. A Logbook’s columns are defined at the time of its creation and can also be changed later. At its core, the Salud! API provides a set of web services for managing Logbooks and their entries.

The Salud! infrastructure also provides several services that facilitate data entry by end-users. Currently, all end-users can enter data into their own accounts via text messages, email, a smartphone application, and a web-based interface. Developers and researchers using the Salud! infrastructure can choose which services they make available to their group of end-users, and can also extend the data entry interface to fit their needs. In the remainder of this section, we will describe how various aspects of the design and architecture of the Salud! infrastructure can support personal health informatics applications.

Full documentation and sample source code to utilize the Salud! API are provided on our developer wiki at <http://wiki.cc.gatech.edu/salud>. Researchers and developers interested in integrating with the infrastructure are encouraged to contact the primary author.

TABLE I. FEATURES PROVIDED BY THE SALUD! INFRASTRUCTURE

<i>Core Web Services API</i>
Creating and managing user accounts
Creating and managing Logbooks
Creating and managing data entries
Managing images associated with entries
<i>Additional Salud! Services</i>
Managing and sending reminders for creating Logbook entries to end-users
SMS- and MMS ^a -based services for end-users to create data entries
Email service for end-users to create data entries
Android and iPhone ^a mobile applications for end-users to create data entries
Web-based data entry and analytics interface

a. Currently under development

B. Primacy of Time-series Data

All data stored in a user’s account on the Salud! infrastructure is associated with a timestamp. This design decision leads from the importance of temporal relationships between variables in health self-management education. Helping individuals understand how physiological metrics are affected by behavioral and environmental factors is a key goal of health self-management education. For this reason, individuals participating in such programs are usually encouraged to keep logs of significant events and activities that may affect their health; for example, diabetics may keep track of their blood glucose readings and influencing factors such as diet composition and physical exercise [10]. Some educational interventions, such as Blood Glucose Awareness Training (BGAT), specifically encourage individuals to monitor temporal effects of stimuli, such as meals or exercise routines, on blood glucose levels over time [12].

Associating timestamps with all data also prevents problems that may arise from outdated information. End-users may only sporadically update their accounts with current data, in which case old measurements would only be valid for historical comparison and not for in-the-moment decisions (either by the user herself, or by any decision-support system which may also be accessing the data). Much physiological or

behavioral data is outdated within days or even hours after its measurement (e.g. blood glucose readings, medication intake) and so it is important to be able to account for the relevance of a data point before factoring it into an in-the-moment decision process.

C. Flexible Data Representation

The Salud! infrastructure provides a flexible data representation mechanism, which allows both developers and end-users to create Logbooks which track a wide range of data. Columns in Logbooks can hold text and numeric data, photos, and data which comes from sets (list data). Columns of the list data type allow for a middle ground between quantitative data and free-form text. For example, if the developers of DiabetesApp wanted to allow end-users to track their food intake, they could use list columns to allow users to specify the meal (“breakfast”, “lunch”, “dinner”, or “snack”) as well as describe the contents of a particular meal, entered as a comma-separated list (e.g. “lettuce, tomatoes, croutons, low-calorie vinaigrette”). List columns work akin to tag lists on social bookmarking services, in that a vocabulary is built up from data values entered into that column.

Logbooks can be created and populated with data both by developers of Salud!-based applications and by their end-users. We presume three types of ways in which Logbooks may be created for a particular application:

- Starter Logbooks created by developers for all end-users of an application. These may include “Blood Glucose Readings” for users managing diabetes, or a Logbook to track use of rescue inhalers for users with asthma. These Logbooks may also be automatically populated with data (more on this below) and users can be prevented from editing them.
- Optional Logbooks, which are made available to end-users of an application to add to their account if and as they want. For example, in DiabetesApp the second variable to be tracked may be selected by each user from a list of available options.
- Personal Logbooks, which are created by end-users to track personally relevant information for which Logbooks are not provided *a priori*. This may include Logbooks which track personal metrics like stress, time spent with family or friends, particular recreational activities, or any other type of activity or behavior in which the end-user is interested. Because Logbooks can contain a variety of columns (or none at all), end-users are able to track variables in a way that is personally meaningful to them, instead of necessarily relying on pre-set templates.

In most cases, end-users can also edit existing Logbooks to better serve their needs. For example, a user who is automatically provided with a Logbook for tracking meals can add additional, personally-relevant columns, such as where the meal was eaten, how it tasted, etc. When necessary, however, researchers or developers can prevent end-users from editing specific Logbooks—for example, if the Logbook’s data comes from an automated service.

The Salud! infrastructure also explicitly allows incomplete data entries. Only a timestamp needs to be initially provided to create a new entry in a Logbook—all other columns (if any) can be populated at a later time. This functionality is meant to encourage prompt and accurate logging. At the time of an event or action that is being recorded (e.g. a meal, or an occurrence of chronic pain), an end-user may not have the time or desire to fill out a complete Logbook entry. In this case, she may choose to create only a minimal entry, noting only the time and possibly adding a picture (e.g. for a meal). Such entries can serve as placeholders and reminders until the end-user has more time to review them and input additional data.

D. Support for Images

Multiple studies have shown that photos can serve as important educational and reflective tools for individuals learning health self-management. For example, a photo of a meal can help an individual better recall the composition of the meal, as well as reflect about whether and why the meal was a healthy or unhealthy choice [5]. Photos of meals can also be shared with educators and peers to spur important discussions about eating patterns and dietary choices [4, 11].

Additionally, photos make it easier to remember the context of a particular situation and may be more engaging than simple text for some users. Taking a photo of a meal with a cameraphone and sending it to a dedicated phone number as a MMS message is an easy way to track eating habits. The details of the meal—calories, contents, etc.—can be filled out at a later time by using the photo as a memory aid.

For these reasons, the Salud! infrastructure allows Logbooks to contain image columns, which permits end-users to associate photos with individual Logbook entries. The infrastructure handles the storage, resizing and retrieval of the images, so developers and researchers can incorporate photos or other images into their personal health informatics applications with less effort.

E. Extensible Data Entry Interface

The final aspect of the Salud! infrastructure that we wish to highlight is the range of ways in which end-users can add data to their accounts, and how these may be extended by developers and researchers. We have developed several data entry interfaces to try to accommodate the preferences of a variety of users. Currently, individuals with Salud! accounts can create entries in their accounts via email, SMS/MMS messages, a mobile application, and a Flash-based online interface. In order to use the email and text message services, an email address and mobile phone number must be associated with the user's account, respectively. We additionally provide an API for managing and sending data entry reminders to end-users. Currently, reminders are sent as text messages and can be scheduled to be sent daily or weekly at specific times of the day. Researchers building applications on top of the Salud! infrastructure can choose to inform their end-users of these data entry options, or opt to develop custom data-entry mechanisms that suit their needs.

The availability of data entry management web services in the Salud! infrastructure makes it straightforward to develop

both automated and interactive data entry services for specific needs. Researchers can also request special system-level authentication tokens, which allow their applications to add or retrieve data from their end-users' accounts without requiring an interactive login. This functionality enables the creation of services which automatically populate data into end-users' Logbooks. For example, one research team who used the Salud! infrastructure to develop a personal health informatics application for families with children with asthma developed a service that regularly updated a Logbook with outdoor air quality data from sensors in end-users' neighborhoods. Similarly, a service could be developed that allows end-users to enter personal data via, e.g., Twitter posts, instant messages, etc. We believe that this flexibility makes the Salud! infrastructure a viable backend for a wide range of potential personal health informatics applications.

F. Privacy and Security

The security and privacy of the data stored in the Salud! infrastructure is of paramount concern to us. All data is stored on a secured machine on our university's campus and backups are created regularly. Communication with the web services that make up Salud!'s API happens over secure (SSL) connections. End-user logins are handled securely by an open-source, industry-standard infrastructure (Drupal). Only individuals directly affiliated with on-going research are provided with authentication tokens that allow them access to appropriate end-user data. This research has been vetted by our university's Institutional Review Board, and a privacy policy is available on the Salud! homepage.

G. Projects Using the Salud! Infrastructure

Currently, we are working with two groups of researchers who are actively using the Salud! platform. The first group of researchers is deploying a personal informatics application for families with children with asthma. To better suit their end-users' needs, the researchers working on this project have developed a set of custom data-entry interfaces better suited for asthma. They have also developed a Salud! service which automatically populates a Logbook in their end-users' accounts

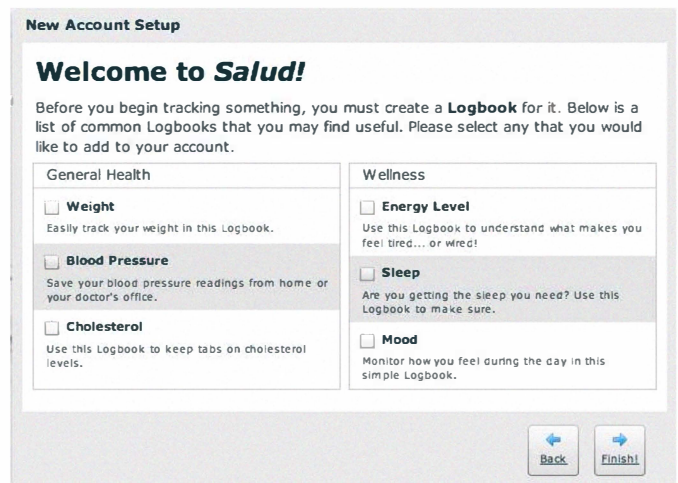


Figure 1. Portion of the Salud! application's account creation process, showing starter Logbooks offered to end-users.

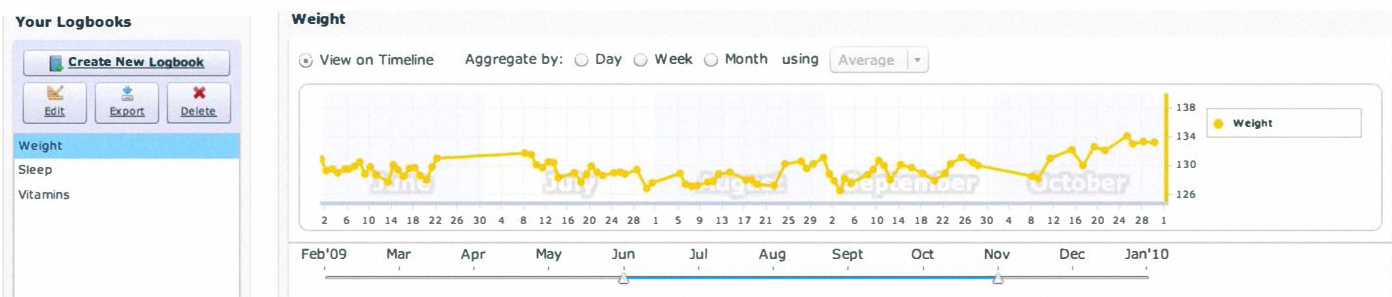


Figure 2. A portion of the Salud! application's interface, showing a user's self-reported weight measurements plotted on a timeline.

with outdoor air quality data obtained from municipal sensors.

The second group of researchers currently working with the Salud! infrastructure are developing a kiosk-based nutritional advice system, Dr. J SAYS.

V. THE SALUD! APPLICATION

We are concurrently developing an application targeted at end-users, which implements the functionality offered by the Salud! infrastructure. The Salud! application is a web-based personal health informatics application aimed at a broad range of end-users. This application is also available as a default user interface (UI) for researchers wishing to use the Salud! infrastructure without creating a custom UI. In the future, we expect to integrate novel health self-management interfaces into the Salud! application, to evaluate their effectiveness. We are also exploring options for making this application more customizable and platform-like so that other researchers could use it as a foundation for their own applications.

The Salud! application allows end-users to track daily activities, events and measurements, as well as to review and analyze the recorded data. Upon registering for an account, new users can select from a set of standardized Logbooks to add to their account. These Logbooks are designed for tracking common health and wellness metrics, such as blood pressure,

sleep, and weight (see Figure 1). After this step, users can create custom Logbooks in which to track other variables of interest to them (see Figure 3). The online application provides a data-entry interface, but users of the Salud! application can also create entries in their Logbooks via text messages, email, and a mobile phone application.

Once recorded, entries are displayed on a visual timeline. If a Logbook contains numerical columns, these are plotted along the y-axis (see Figure 2). Otherwise, entries in the Logbook appear as points on the timeline. This allows users to view temporal fluctuations in numeric variables or spot patterns in the frequency of events. Additionally, entries can be aggregated by summing or averaging over daily, weekly or monthly periods.

The application also provides an analytics interface where data from different Logbooks can be viewed together. Users can drag columns or Logbooks from a list into the analysis pane to overlay data in a single graph or stack several graphs (see Figure 4). All of the data is plotted against the same timeline on the x-axis, which makes it easy to look for relationships between data in different Logbooks.

We have not yet formally evaluated the Salud! application with users and have not marketed it to end-users. However, at the time of this writing, over thirty individuals unaffiliated with our team or any research projects currently using the Salud! infrastructure have registered for an account, and approximately half of these are active users. Many of these users learned about the application by first downloading our Android mobile data entry application from the Android Market.

VI. ONGOING AND FUTURE WORK

The Salud! infrastructure is under active development, and in this section we describe some of the features and services we expect to launch within the next several months. We expect that these additions will make the Salud! infrastructure more effective and desirable for other researchers and developers to use when creating and deploying health self-management interventions.

A. Goal Management API

Action plans—short-term goals agreed upon by an individual and a healthcare provider, and which the individual feels confident she can achieve—are an effective health self-management strategy [1, 2]. For this reason, we are designing a

Name	Type	Description
Date	Date	The date of an entry.
Time	Time	The time of an entry.
Meal pic	Picture	
Contents	List	What was in the meal.
Meal type	List	Breakfast, lunch, snack, etc.
Calories	Number	

Figure 3. Example of a Logbook created by a user of the Salud! application.



Figure 4. A portion of the Salud! application's analytics interface. Two columns from different Logbooks are shown on the same graph. Graph values are weekly averages of the two columns.

goal management system into the Salud! infrastructure. The API for this system will allow interested researchers and developers to easily integrate goal management functionality into their applications. We are looking to literature on behavior change and self-efficacy to understand effective strategies for goal setting and achievement, and expect to integrate these into the goal management system. As such, the functionality made available through the goal management API will encapsulate effective strategies and best practices for goal management. We are currently focusing on two specific health self-management goal management strategies: actionable vs. inactionable goals, and experiences of personal mastery.

An actionable goal can be achieved through a specific behavior or activity, whereas an inactionable goal cannot be addressed directly. For example, a weight goal is an inactionable goal because no specific activity leads directly to loss of body mass. Rather, weight loss is achieved by activities such as reducing calorie intake or increasing physical activity. As such, individuals who wish to set a weight goal should be additionally encouraged to set a related, actionable goal (such as reducing weekly calorie intake by a specific amount). Actionable goals are a necessary component of action plans, and as such we wish to integrate this concept into Salud!'s goal management API.

The notion of experiences of personal mastery [9] is a goal management strategy concerned with how to iteratively adjust goals in the face of success or failure. Though a full discussion of the concept is outside of the scope of this paper, a key takeaway is that goals should be set within reach of an individual's current abilities; starting with an ambitious goal has a high chance of leading to failure, which is often demotivating. Instead, individuals should move in small steps from their current state to their desired state via a series of successful goal achievements. This process results in increased self-efficacy and is more likely to result in lasting change [9]. The goal management infrastructure we are designing will include logic for goal adjustment based on these principles.

B. User Analytics Interface

One of the key services currently under development for Salud! is an analytics interface for researchers. This interface will provide simple business analytics tools for researchers using the Salud! infrastructure to evaluate the adoption and use

of their particular application. Specifically, it would show how a particular application's end-users are using it: how many Logbooks and entries they are creating, and via which data entry methods; show temporal patterns of use, such as how frequently end-users log in; and show breakdowns of use patterns by filters such as highly-active and less-active users, which optional Logbooks users selected, and so forth. We are actively soliciting input from researchers using or expecting to use the Salud! infrastructure to obtain and prioritize specifications for this feature.

VII. CONCLUSION

In this paper, we presented Salud!, an open infrastructure for creating and deploying personal health informatics applications. We described how the design of the platform follows from effective health self-management strategies and how we are working to extend the infrastructure with new features and services. We also presented an end-user targeted personal health informatics application that we are currently implementing using the functionality provided by the Salud! infrastructure and its associated API.

The Salud! infrastructure is currently being used by two research groups, but as the Salud! project grows we are looking for additional researchers to partner with in the design and development of this infrastructure. We are also actively soliciting feedback from current and potential Salud! users about how to prioritize development. Researchers and developers interested in working with Salud! should contact the primary author and visit our homepage at <http://salud.cc.gatech.edu>.

ACKNOWLEDGMENT

We thank past and present Salud! developers and designers: Ruchi Goswami, Akshit Proothi, Aparna Raman, Catherine Grevet, Dev Priya, Nabanita Ghosal, Zack Brunson, Vasudhara Kantroo, Mebaa Kidane, Joy Buolamwini, Dhawal Shah and Arun Kutty. Thanks to Ashwin Ram, Keith Edwards, John Stasko and Anind Dey for valuable discussions. Thanks also to Jiten Chhabra, Tae-Jung Yun and Andrew Miller for their feedback.

REFERENCES

- [1] T. Bodenheimer, K. Lorig, H. Holman and K. Grumbach. Patient Self-management of Chronic Disease in Primary Care. *JAMA* 288, 19 (2002), 2469-2475.
- [2] P. G. Gibson, J. Coughlan and M. Abramson. Self-management education for adults with asthma improves health outcomes. *Western Journal of Medicine* 170, 5 (1999), 266.
- [3] K. R. Lorig, et al. Evidence Suggesting That a Chronic Disease Self-Management Program Can Improve Health Status While Reducing Hospitalization: A Randomized Trial. *Medical Care* 37, 1 (1999), 5-14.
- [4] J. Frost and B. K. Smith. Visualizing health: imagery in diabetes education. Proceedings of the 2003 conference on Designing for user experiences, ACM (2003), 1-14.
- [5] L. Mamykina, E. D. Mynatt, P. A. Davidson and D. Greenblatt. MAHI: Investigation of Social Scaffolding for Reflective Thinking in Diabetes Management. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM Press (2008).
- [6] P. G. Gibson, et al. Limited (information only) patient education programs for adults with asthma. *Cochrane Database of Systematic Reviews* (Online), 2 (2002), CD001005.
- [7] K. R. Lorig, D. S. Sobel, P. L. Ritter, D. Laurent and M. Hobbs. Effect of a self-management program on patients with chronic disease. *Effective Clinical Practice* 4, 6 (2001), 256-62.
- [8] K. R. Lorig, P. L. Ritter, D. D. Laurent, and K. Plant. Internet-based chronic disease self-management: a randomized trial. *Medical Care* 44, 11 (2006), 964-71.
- [9] A. Bandura. Self-efficacy: toward a unifying theory of behavioral change. *Psychological Review* 84, 2 (1977), 191-215.
- [10] L. Mamykina, E. D. Mynatt, and D. R. Kaufman. Investigating Health Management Practices of Individuals with Diabetes. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM Press (2006), 927-936.
- [11] B. K. Smith, J. Frost, M. Albayrak, and R. Sudhakar. Integrating glucometers and digital photography as experience capture tools to enhance patient understanding and communication of diabetes self-management practices. *Personal and Ubiquitous Computing* 11, 4 (2007), 273-286.
- [12] D. J. Cox, L. Gonder-Frederick, W. Polonsky, D. Schlundt, B. Kovatchev and W. Clarke. Blood glucose awareness training (BGAT-2): long-term benefits. *Diabetes Care* 24, 4 (2001), 637-642.