Mobile Netw Appl
DOI 10.1007/s11036-013-0484-5

BeWell: Sensing Sleep, Physical Activities and Social Interactions to Promote Wellbeing

Nicholas D. Lane · Mu Lin · Mashfiqui Mohammod · Xiaochao Yang · Hong Lu · Giuseppe Cardone · Shahid Ali · Afsaneh Doryab · Ethan Berke · Andrew T. Campbell · Tanzeem Choudhury

Abstract Smartphone sensing and persuasive feedback design is enabling a new generation of wellbeing apps capable of automatically monitoring multiple aspects of physical and mental health. In this article, we present *BeWell+* the next generation of the BeWell smartphone wellbeing app, which monitors user behavior along three health dimensions, namely sleep, physical activity, and social interaction. BeWell promotes improved behavioral patterns via feedback rendered as an ambient display on the smartphone's wallpaper. With BeWell+, we introduce new mechanisms to address key limitations of the original BeWell app; specifically, (1) *community adaptive wellbeing feedback*, which generalizes to diverse user communities (e.g., elderly, children) by promoting better behavior yet remains realistic to the user's lifestyle; and, (2) *wellbeing adaptive energy allocation*, which prioritizes monitoring fidelity and feedback responsiveness on specific health dimensions (e.g., sleep) where the user needs additional help. We evaluate BeWell+ with a 27 person, 19 day field trial. Our findings show that not only can BeWell+ operate successfully on consumer smartphones; but also users understand feedback and respond by taking steps towards leading healthier lifestyles.

Keywords Smartphone sensing · mHealth · Wellbeing apps

N. D. Lane (✉)
Microsoft Research Asia, Beijing, China
e-mail: niclane@microsoft.com

M. Lin · X. Yang · S. Ali · E. Berke · A. T. Campbell
Dartmouth College, Hanover, NH 03755, USA

H. Lu
Intel Lab, Santa Clara, CA 95054, USA

G. Cardone
University of Bologna, Bologna, Italy

A. Doryab
Carnegie Mellon University, Pittsburgh, PA 15213, USA

M. Mohammod · T. Choudhury
Cornell University, Ithaca, NY 14850, USA

1 Introduction

Our lifestyle choices have a deep impact on our personal health. For example, our sleep, socialization and exercise patterns are connected to the presence of a wide range of health related problems such as, high-blood pressure, stress [45], anxiety, diabetes and depression [22, 27]. Positive health effects can be observed when these wellbeing indicators (e.g., sleep, physical activity) are kept in healthy ranges. However, people are typically not exposed to these health indicators as they go about their daily lives. As a result, unbalanced unhealthy lifestyles are present in the general population. People demonstrate concern for some aspects of their wellbeing, such as fitness or diet, yet neglect the wellbeing implications of other behaviors, such as, poor sleep, hygiene or prolonged social isolation. We believe this situation is caused by an absence of adequate tools for effective self-management of overall wellbeing and health.

We envision a new class of *personal wellbeing apps* for smartphones capable of monitoring multiple dimensions of human behavior, encompassing physical, mental and social dimensions of wellbeing. An important enabler of this vision are the recent advances in smartphones, which are equipped with powerful embedded sensors, such as

Published online: 09 January 2014

an accelerometer, digital compass, gyroscope, GPS, microphone, and camera. Smartphones present a programable platform for monitoring wellbeing as people go about their lives [49]. It is now possible to infer a range of behaviors on the phone in real-time, allowing users to receive feedback in response to everyday lifestyle choices that enables them to better manage their health. In addition, the popularity of smartphone app stores (e.g., the Apple App Store, Android Market) has opened an effective software delivery channel whereby a wellbeing app can be installed in seconds, further lowering the barrier to user adoption. We believe production-quality wellbeing apps will gain rapid adoption globally, driven by: (1) near zero user effort, due to automated sensor based activity inference and (2) universal access, only requiring a single download from a mobile phone app store and installation on an off-the-shelf smartphone.

In [34] we introduced BeWell, a *wellbeing app* that runs on off-the-shelf sensor-enabled smartphones. BeWell coarsely tracks the physical, social and sleep dimensions of wellbeing by monitoring several key behavioral patterns and providing feedback to the user. Ideally feedback would allow users to easily understand the consequences of their actions, enabling them to make appropriate changes in their behavior and more informed choices going forward. We evaluated BeWell through lab-based single phone experiments that measured the resource requirements of our design. In addition, a small five person experiment was conducted to investigate the robustness of the activity inferences BeWell performs [34]. These benchmark experiments, although small-scale, highlighted key barriers to wider-scale deployments. First, we found despite careful engineering and even when using a high-capacity battery, BeWell exhausted the battery life of the smartphone after only 8-12 hours - forcing users to recharge multiple times per day. Second, even with a small number of users we encountered significant diversity, with users exhibiting a wide range of behavioral patterns. As a result of this initial finding, it became clear using a single static expectation of ideal "healthy" behavioral patterns would lead to feedback that was not consistently realistic for everyone; limiting the ability of the system to scale to a larger population of users with diverse health needs.

Guided by these insights gained in developing the original BeWell app we propose *BeWell+*,[1] which incorporates a set of new wellbeing scaling techniques for generating community-guided wellbeing feedback and overcoming energy constraints evident in resource limited smartphones:

Community adaptive wellbeing feedback We design a wellbeing feedback mechanism in which the expectations of healthy behavioral patterns are adjusted to remain realistic for what is possible in the near-term for certain user communities. Instead of only relying on generally ideal (i.e., one size fits all) behavioral feedback (e.g., suggesting 8 hours of sleep) our new BeWell+ design is based on the community the user is associated with – their peer group. For example, it is unrealistic to expect an elderly person to meet the same goals for physical activity as a young adult; or for that matter a doctor that is on call having the same goal for hours of sleep as a high schooler. For each user a "wellbeing network" is identified in the user population, based on shared behavioral traits. Within the network positive and negative "role-models" are identified, and their behavior – along with established ideal behavioral goals – determine user wellbeing feedback. As a result, users are not provided with unrealistic expectations of behavior change, since they are compared to role-models/groups of peers. Here health goals are tailored to peer group norms as opposed to a generalized population wide norm. As the health of a user improves they join progressively more healthy communities of users with more challenging wellbeing feedback. We describe this new scalable mechanism as community adaptive wellbeing feedback.

Wellbeing adaptive energy allocation We design an energy allocation scheme that prioritizes resources so those dimensions of behavior that the individual is currently struggling with (e.g., physical activity) are: (1) more accurately assessed and (2) provided with immediate feedback, helping to create awareness and promote change in individuals. User behaviors that consistently trend close to healthy norms are monitored less closely, with feedback provided on a slower time-scale – therefore, less system resources (e.g., energy) are required in this case. Using this approach, the more problematic user wellbeing behaviors still receive the attention they demand, while key elements of smartphone usability (e.g., stand-by time) also can remain within ranges acceptable to users. We describe this new scalable mechanism as wellbeing adaptive energy allocation.

Both of these techniques are implemented as part of the BeWell+ app deployment described in this article. We present results from the first user study of any BeWell app in the wild, which includes 27 people using BeWell+ over a 19 day field trial. Findings from our study show that (1) BeWell+ can coarsely assess wellbeing automatically, which we validate using established medical self-report surveys of physical and mental wellbeing; (2) community-adaptive wellbeing feedback can promote realistic personalized health goals for each user; (3) despite the complexity of multi-dimensional wellbeing feedback, users understand BeWell+ feedback and are able to identify appropriate

[1]BeWell+ is available for download and use with any off-the-shelf Android Smartphone. Please download BeWell+ from: http://www.bewellapp.org

corrective actions to take; (4) significant increases in energy efficiency result from wellbeing adaptive energy allocation; and, (5) users react positively to their overall experience and even show improvements in their ability to link everyday actions to wellbeing outcomes.

This article is an extended version of a paper [41] presented at Wireless Health 2012. We provide an overview of the BeWell+ app in Section 2. Section 3 details the design of community adaptive wellbeing feedback. We describe wellbeing adaptive energy allocation in Section 4. We present the evaluation of the system and results from a user study in Section 5. Finally, we discuss related work (Section 6) and then make some concluding remarks in Section 7.

2 BeWell+ overview

In this section, we describe the BeWell+ app and architecture. The BeWell+ app was developed for current smartphones as a proof-of-concept system for monitoring and promoting holistic wellbeing. In our prior work [34] we evaluated an earlier implementation of BeWell+, testing the accuracy of the human activity inferences that it builds upon and its ability to meet a series of system requirements (e.g., battery, computation). But we did not deploy the app or evaluate the system's ability to monitor and provide feedback along different health dimensions.

Overview As shown in Fig. 1, BeWell+ consists of two software components: (1) BeWell+ phone app and (2) Cloud infrastructure. The BeWell+ phone app automatically monitors user's everyday activities using the accelerometer and microphone sensors on phone. Inference results from the classifiers on the phone are then transmitted to the BeWell+ Cloud infrastructure. The Cloud infrastructure stores all the data and computes wellbeing scores. Wellbeing scores summarize the impact on overall health based of the inferred behavioral patterns. BeWell+ computes wellbeing scores for each health dimension it tracks. In the current prototype these are: physical activity, sleep patterns and social interaction. The BeWell+ phone app presents these scores back to users on the phone, using an ambient display rendered on the wallpaper of the device (see Fig. 2).

2.1 Monitoring behavior

To track physical activity and social interaction BeWell+ relies on already developed activity recognition techniques (e.g., [14, 37, 42, 43]). The physical activities of users are classified into one of the following classes: $\{stationary, running, walking\}$ using an accelerometer. Similarly, the microphone is used to classify ambient audio as either: $\{non - voicing, voicing\}$. Our implementation adopts the set of audio and accelerometer features detailed in [43], which includes a range of time and frequency domain features. Classification is done using a boosted ensemble of naive Bayes classifiers [15]. Temporal smoothing is applied to classifier results using a simple Markov model.

Monitoring sleep behavior requires a completely different approach. We developed in [34] a simple logistic model that estimates the amount of hours slept based on a collection of phone usage pattern. Our sleep model uses statistics (e.g., duration, frequency) of everyday events correlated with the amount of sleep the user receives. The model uses occurrences of mobile phone recharging along with other events detected by our activity inference model, namely, periods of near silence or the phone being stationary.

2.2 Wellbeing scores

Wellbeing scores range between 0 and 100 and are calculated for each of the three dimensions (viz. physical activity, social interaction and sleep patterns). A score of 100 indicates the person is matching or exceeding recommended guidelines (e.g., averaging 8 hours sleep per day is represented by a score of 100). Our wellbeing scoring functions are a result of a careful design process which leveraged: the existing literature, guidelines from institutions (e.g., CDC), collaboration with medical researchers and short

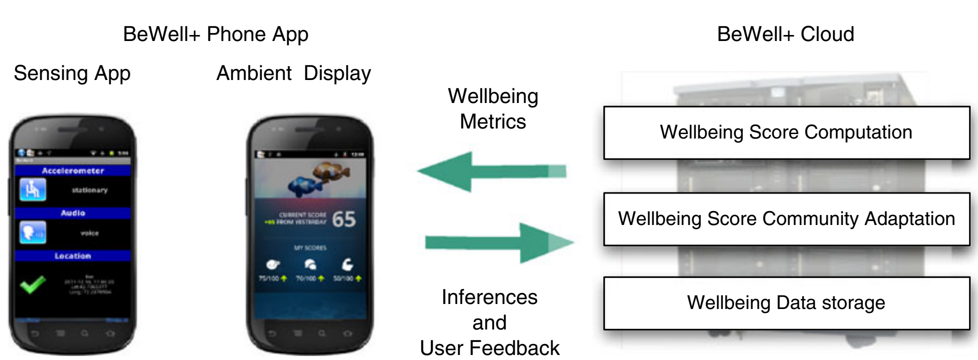

Fig. 1 BeWell+ app implementation, including smartphone components supported by a scalable cloud system

Fig. 2 Multiple wellbeing dimensions are displayed on the smartphone wallpaper. An animated aquatic ecosystem is shown with two different type of fish whose behavior are affected by changes in wellbeing (i.e., activity and social interaction); in addition, the ocean ambient lighting conditions reflect the users sleep duration (*shown on Nexus S*)

field experiments. For a comprehensive description of these functions consult [34]; in what follows we provide a brief summary:

Sleep Research exploring sleep health effects focus both on the quantity and quality of sleep [39]. Although both of these facets are important we focus solely on monitoring sleep duration. Studies show oversleeping ("long sleep") carries similar negative health consequences to insufficient sleep [11], thus, we penalize both behaviors. BeWell+ computes a wellbeing score for sleep behavior within a single day using a gaussian function,

$$sleep_{day}(\mathrm{HR}_{act}) = Ae^{-\frac{(\mathrm{HR}_{act}-\mathrm{HR}_{ideal})^2}{2(\mathrm{HR}_{hi}-\mathrm{HR}_{lo})^2}} \quad (1)$$

in which HR_{act} is the total quantity of sleep over a 24 hour period, HR_{ideal} is the ideal hours asleep with HR_{hi} and HR_{lo} being the upper and lower limits of acceptable sleep duration. Our sleep function is parameterized using a HR_{ideal} of 7 hours with a HR_{hi} of 9 hours and HR_{lo} of 5 hours, these values are consistent with existing sleep studies [11].

Physical activity Our assessment of physical activity is based on the duration the user is recognized as performing common physical activities (viz. walking, stationary, and running). These inferences are used to estimate a daily Metabolic Equivalent of Task (MET) value [12]. Being definitive as to the ideal MET levels for an individual is difficult as mental and physical health benefits occur at different levels of activity. These MET levels are also sensitive to user characteristics, such as, existing physical fitness or particular genetic determinants. Initially, we currently rely on generic guidelines established by the Centers for Disease Control and Prevention [1] (CDC). Our daily scores of physical activity are simply a linear regression,

$$physical_{day}(\mathrm{MET}_{act}) = (\mathrm{MET}_{hi} - \mathrm{MET}_{lo})\mathrm{MET}_{act} + \mathrm{MET}_{lo} \quad (2)$$

where MET_{act} is the actual MET value for a user during that day, with MET_{hi} and MET_{lo} being calibrated by the high-end and minimum guidelines for adult aerobic activity set by the CDC. These values range between 300 and 150 minutes of moderate-intensity per week. Such aerobic activity should be accompanied by muscle-strengthening programs, ideal behavioral patterns for these programs are also available from the CDC and are included within the existing physical activity guidelines. However, we neglect this aspect of physical activity due to the inaccuracy in monitoring muscle-strengthening programs without on-body sensors.

Social interaction Medical studies use a variety of measures to capture the social environment of a person. The development of these measures are still an active area of research. BeWell+ focuses on one of these metrics, social isolation, as it is more easily captured with sensors available in smartphones today. Studies of particular high-risk communities show social isolation is correlated with basic forms of human contact. For example, health deterioration exhibited in the elderly is linked with, amongst others, a decline in the frequency of human interaction (e.g., phone calls and visits with friends and relatives) [29]. In the general population, those with profound acquired hearing loss have been seen to suffer a deterioration of psychological wellbeing due to the associated communication difficulties [28]. We measure social isolation based on the total duration of ambient conversations, which are detected by inferences made using the mobile phone microphone. Insufficient medical evidence exists to parameterize this relationship. At this time we again use a wellbeing score for social interaction with a linear regression,

$$social_{day}(\mathrm{DUR}_{act}) = (\mathrm{DUR}_{hi} - \mathrm{DUR}_{lo})\mathrm{DUR}_{act} + \mathrm{DUR}_{lo} \quad (3)$$

where DUR_{act} is the duration of conversation detected relative to the total time the microphone is active during a single day. We determine empirically a value for DUR_{hi}, 0.35, using the mean conversation ratio of a small 10 person experiment; we also utilize this group to train our classifiers (see Section 5). As we lack a population in which poor wellbeing has caused atypical conversation patterns our DUR_{lo} ratio is simply set to zero.

2.3 Ambient display

The ambient display is an animation that is rendered on the phone's lock-screen and wallpaper, making it visible to the user whenever the user glances or interacts with their smartphone. The display provides passive feedback to the user of their current wellbeing scores. Prior examples of successful persuasive systems [16] have found that wallpaper can effectively promote changes in user activity. These studies show phone wallpaper, when used as a glanceable display, can keep user goals "persistently activated" [32] in the mind of the user.

BeWell+ displays multiple wellbeing dimensions as an aquatic ecosystem, as illustrated in Fig. 2. The animated activities of a clown fish (which mirrors the user's activity), the ocean ambient lighting conditions (which mirrors the user's sleep duration) and a school of small fish (which mirror the user's level of social interaction) provide a quick summary of the current wellbeing to the user. The relationship between the ambient display and the wellbeing scores is described below:

Clown fish The clown fish represents the physical activity of the user. The score modifies the speed at which the clown fish swims. At low levels of physical activity the fish moves slowly from left to right lethargically. As the user's physical activity increases, the fish swims more vigorously, even performing summersaults and backflips at high levels of activity.

School of blue fish A school of fish swims with the clown fish and represents the users social activity. The closeness of the school of fish to the clown fish and its size grows proportionally to the amount of social interaction of the user.

Lighting of ocean Sleep patterns are captured by the light of the ocean. The ocean gets darker when the user lacks sleep and has a low sleep score. As the user sleep level increases, the ocean gradually become brighter.

The aquatic ecosystem represents a single point in the design space of the ambient display for BeWell+. Before selecting this visualization we performed small scale informal surveys of people from the target population. We found strong preference from people for the aquatic ecosystem, which is in agreement with examples from the literature where animated animals are effective at motivating behavior change (e.g., [40].)

3 Community adaptive wellbeing feedback

In this section, we describe BeWell+'s data-driven community-adaptive approach to wellbeing feedback. The behavior goals that underpin wellbeing feedback are based on a combination of observations from the user population and ideal "healthy" behavioral patterns. This allows feedback to automatically tune itself to the population in which the system is deployed.

Implications of community diversity Wellbeing problems and solutions can be highly personal. Each individual has their own challenges to wellbeing shaped by factors including, personal characteristics and behavioral tendencies. We first observed this problem even in our initial deployments of the original BeWell system, where we observe large differences even within small groups.

To help quantify this problem further we turn to data from our 27-person field trial (see Section 5.1 for further details). To measure differences in wellbeing we use the three dimensions of wellbeing scores previously defined in Section 2 (prior to any adaption). Figure 3 shows the distribution of all wellbeing scores for each user, irrespective of the particular health dimension. Surprisingly, even within a relatively small and homogenous group of people significant diversity is present. From this figure we see that the value and the variance of the wellbeing scores vary significantly across users. Although not visible in the figure, we also find the user behavior is also diverse within each separate dimension. For example, there are larges difference between the upper and lower quartile wellbeing scores of subjects for each dimension. Specifically, these differences are 61 % for physical wellbeing, 83.1 % for social wellbeing and 75.5 % for sleep dimension.

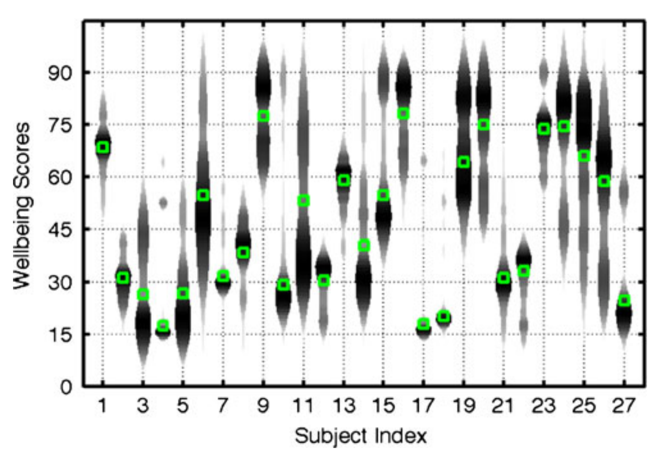

Fig. 3 A high diversity of wellbeing behavioral patterns exists among our study population. A score of 100 refers to a "healthy" behavioral pattern

Feedback from wellbeing apps (along with most mobile health systems) is typically goal based, implying an ideal behavioral pattern to promote. However, high levels of community diversity prevent a single goal behavioral pattern being applied to an entire user population. For example, it is unrealistic to expect an elderly person to meet the same goals for physical activity as a young adult. Similarly, doctors or students often can not conform to a “normal” sleep pattern – but would still benefit from appropriate feedback indicating how their sleeping habits could be made more healthy. Without adjusting the expectations that underpin this feedback it will be ineffective to many users while also damaging the confidence they have in the system. This is the main goal of our new community feedback paradigm – it presents a mechanism to potentially provide effective feedback for very large populations of diverse users in a scalable manner.

Adaptive wellbeing feedback BeWell+ adapts generic wellbeing score functions based on the overall behavior similarity within the user population along with the similarity of users to ideal wellbeing behaviors. This novel process within BeWell+ allows feedback to adapt to the differences between user communities. Without adaptation improvements in behavioral patterns are not considered within the correct context. Another example is as follows: a shift worker who is able to increase her average quantity of sleep from 4 to 5 hours, but this improvement may still score poorly if compared to the general expectation applied to the general public. However, if compared to other shift workers this change could well place the individual in a high performing percentile of that community or peer group. Therefore, wellbeing feedback should recognize this as a substantial positive change, even if the change required to achieve “normal” sleep hygiene remains large.

Adaptation is a data-driven process which relies on activity inferences, along with a trace of periodic GPS estimates, being transferred to the cloud from the BeWell+ app. Figure 4 illustrates each phase of the adaptation process, all computational stages of adaptation are performed by the cloud. Although only a single dimension is shown this process is repeated for all three dimensions. The detailed wellbeing score functions for these three dimensions can be found in [34]. Each function takes a specific statistic related to a user behavioral pattern. Adapted score functions maintain the same functional form, but with parameters being revised to accommodate user diversity. At the conclusion of this process a personalized set of wellbeing score functions are generated for all BeWell+ users.

Guiding the adaptation process is a behavioral similarity network, a weighted graph in which nodes correspond to users and edge-weights quantify the level of similarity. This network attempts to identify people with related

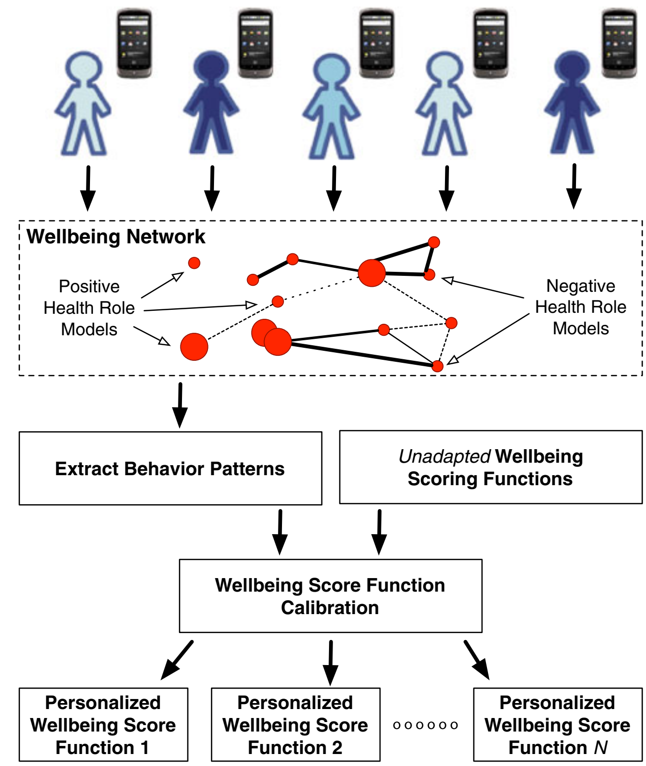

Fig. 4 BeWell+ community adaptive wellbeing feedback

lifestyles and behavior constraints. BeWell+ computes similarity by adopting the lifestyle similarity definition used in [35, 36]. Specifically, we compute lifestyle similarity using three types of information: mobility and diurnal patterns in combination with the distribution of activities performed by users. Mobility patterns are based on GPS location estimates, which are tessellated into m distinct square tiles of equal size. Diurnal patterns are captured as a series of timestamps that are recorded whenever the user is inferred to be non-stationary by the classification pipeline. These timestamps are rounded, and are represented as the particular hour in the week in which they occur (e.g., they range between hour 0 at the start of the week to hour 167 on the final hour of the final day). The distribution of activities is based on the duration users are inferred to be performing each activity classes (e.g, walking, socializing) detected by the classification pipeline. We construct three histograms for each of these types of lifestyle information for every user, normalizing the frequencies across all histograms. For each pair of users (i, j), we compute the lifestyle based similarity by the following equation:

$$sim(i, j)^{life} = \sum_{f \in \mathcal{F}} \mathbf{T}_f(i)^{\top} \mathbf{T}_f(j) \quad (4)$$

where $\mathbf{T}_f(i)$ is a histogram vector for user i of type f and $\mathcal{F}$ contains each type of lifestyle histogram. Lifestyle similarity between two users is the sum of the inner product of the histograms for each type of lifestyle information used by BeWell+. Our use of a similarity network allows us to avoid making a hard assignment of users to specific communities. Hard assignment may degenerate wellbeing scoring functions if the community assignment is incorrect.

It is critical to keep our personalized wellbeing score functions grounded with respect to "healthy" behavioral norms. Consequently, we use an *unadapted* wellbeing score function to balance the need to identify similar people with the need to recognize which of these people are either positive or negative wellbeing role-models. We begin by using the unadapted wellbeing score function once to score all previously collected user behavior. For every observation of user behavior (e.g., physical activity across a day) a data tuple is formed containing, (1) the wellbeing score, and (2) the relevant statistic concerning the wellbeing behavior (e.g., a daily MET value). Each tuple is weighted based on two factors: first, the similarity network edge-weight; and second, how "healthy" users are compared to an ideal behavioral pattern – determined by their average wellbeing score (using an unadapted scoring function). Tuple weight is a linear combination of these values with a parameter, `sim_strength` that determines how much influence user similarity has over the final weight used.

Finally, the adapted scoring function is generated by applying a weighted smoothing over the collection of tuples and fitting the wellbeing functions to the smoothed tuples. Adapted scoring functions set the underlying goals associated with high scores as realistic near-term objectives, rewarding improvement relative to people within their own community or peer group. Of course over the longer term, the ultimate goal of reaching a more ideal pattern remains important. Our adaptive scoring strategy incorporates this requirement by the process repeating as new data accumulates. Each time the adaptation process repeats it incrementally selects higher performing people as a frame of reference for the user while still emphasizing the need for these people to be relatively similar to the target user, with `sim_strength` controlling this trade-off.

4 Wellbeing adaptive resource allocation

In this section, we discuss the design of our wellbeing adaptive resource allocation strategy. The novelty of this approach is to prioritize the resource allocation based on how well the user is coping with each individual health dimension. BeWell+ dynamically shifts resources between wellbeing dimensions (viz. physical activity, social interaction and sleep patterns) as the behavior of the user changes – dimensions with low wellbeing scores receiving more resources than those with high scores. As a result, the accuracy and responsiveness of the BeWell+ app are optimized within resource constraints and with an awareness of the user's wellbeing needs.

Insufficient energy resources Monitoring wellbeing requires multiple aspects of daily life to be constantly monitored. This puts undue load on the battery of smartphones as this requires sensing and inference to be performed continuously across a range of sensor modalities. Figure 5 shows the battery life of five subjects using the original BeWell system, as reported in [34]. Even though each Android smartphone is equipped with a large-capacity battery (3200 mAh) battery life varies between 12 and 21 hours. If we assume the use of a factory standard battery (1400 mAh) then these lifetimes will be reduced to between 7 and 10 hours. At this level users will have to recharge their phone multiple times per day, otherwise BeWell will only be able to monitor them for the fraction of the day when the phone is active. This problem is more broadly applicable to the growing number of mobile health apps that consider multiple dimensions of behavior; and even further, is known to impact a variety of mobile sensing apps [47] and smartphone platforms [44].

Adaptive energy allocation BeWell+ conserves smartphone energy usage by dynamically tuning the duty cycle of core system components based on the wellbeing score of the user. Figure 6 illustrates the control-loop used by BeWell+ to intelligently allocate the energy consumption, and highlights which component duty cycle parameters are tuned.

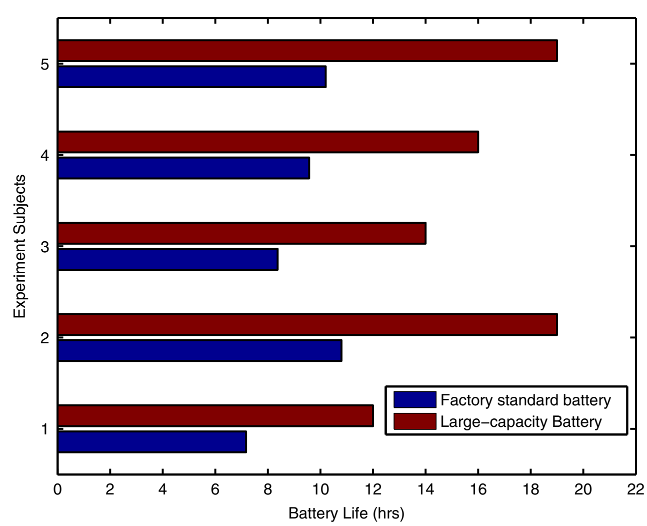

Fig. 5 The spread of battery lifetime for the original BeWell system. For each person we show battery life using a large capacity battery (3200 mAh) and an estimate of battery life with a standard factory battery (1400 mAh)

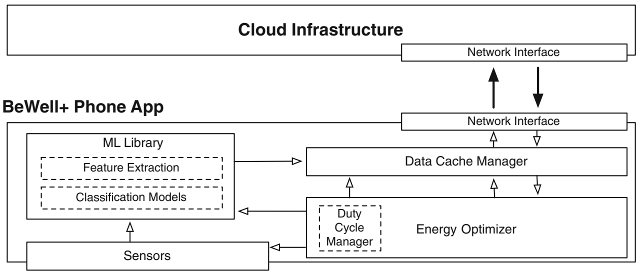

Fig. 6 BeWell+ Energy Management Subsystem

Specifically, these parameters are: the rate at which sampling, feature extraction and activity inference routines are performed; along with how often BeWell+ interacts with the cloud to either upload user-specific statistics or collect revised wellbeing scores – both of which require community interaction and so necessitate the cloud to be involved.

Our energy management strategy is based on a simple yet effective optimization, which we will now describe:

Let $\mathtt{duty}_\mathtt{i}$ denote the $\mathtt{i}th$ duty cycle parameter in $\mathtt{duty.all}$, the set of all duty cycle parameters in the BeWell+ app. Let function $\mathtt{acc}_\mathtt{j}(\mathtt{duty.all})$ estimate the increase in error for specific dimensions of wellbeing scores (indicated by $\mathtt{j}$) due to increasing levels of duty cycling. Further, let $\mathtt{bat}(\mathtt{duty.all})$ estimate the per day smartphone energy consumption due to BeWell+ operation, relative to potential $\mathtt{duty}_\mathtt{i}$ values. The functions of $\mathtt{bat}$ and $\mathtt{acc}_\mathtt{j}$ use a polynomial regression fitted with data by profiling the BeWell+ app running with different $\mathtt{duty}_\mathtt{i}$ values, in addition to data from user experiments, which enables accuracy to be assessed. The values for each $\mathtt{duty}_\mathtt{i}$ parameter is found by optimizing the following objective function:

$$\arg \min_{\mathtt{duty.all}} bat(\mathtt{duty.all}) + \sum \alpha_j \cdot acc_j(\mathtt{duty.all}) \tag{5}$$

where, α_j is a weighting term allowing the accuracy of certain dimensions of wellbeing scores to be emphasized over others. Specifically, α_j is simply:

$$\frac{1}{\mathtt{z}} \cdot (\mathtt{score}_{\mathtt{max}} - \mathtt{score}_{\mathtt{actual}})/\mathtt{score}_{\mathtt{max}} \tag{6}$$

where $\mathtt{score}_{\mathtt{max}}$ is the maximum wellbeing score, $\mathtt{score}_{\mathtt{actual}}$ is the present value for the $\mathtt{j}$th dimension of wellbeing and $\mathtt{z}$ is the term used to normalize weights across all wellbeing dimensions.

The adaptive energy allocation component, shown in Fig. 6, performs this optimization each time there is a change in the wellbeing scores. As the wellbeing of the user shifts (e.g., an unhealthy behavior improves significantly), BeWell+ can automatically re-allocate energy to provide more accurate monitoring and more responsive feedback for the new wellbeing dimension of highest concern.

5 Evaluation

In this section, we study the performance of BeWell+ with a 27 person field trial conducted over 19 days. We find that: (1) BeWell+ can accurately track wellbeing across multiple behavioral dimensions; (2) our community adaptive wellbeing feedback mechanism can reconcile health norms with the practical restrictions that limit near-term user lifestyle changes; (3) users can digest multi-dimensional BeWell+ feedback and are seen to make positive changes in their behavior; (4) wellbeing adaptive energy allocation is able to intelligently allocate resources to underperforming aspects of user wellbeing, while also adjusting to lifestyle changes; and, (5) users report an overall positive experience from the BeWell+ field trial.

5.1 Study methodology

Our study population contains 16 men and 11 women aged between 21 and 37. Of these subjects, 9 % are faculty or graduate students in a computer science department, 34 % are doctors or medical researchers and the remaining 57 % are students in the arts and life sciences graduate program. Each volunteer agrees to carry a phone with the BeWell+ app installed. The subjects either move their mobile phone SIM card into the Nexus One or use call forwarding so they can use the study phone as their primary phone. We provide each user with a holster to clip the phone on to their belt or clothing. Users agree to keep the phone with them at all times.

To verify the effectiveness of presenting multi-dimensional feedback using the ambient display, the participants are randomly and uniformly split into two groups: *multi-dimensional group* and *baseline group*. All subjects have the core BeWell+ software installed that tracks sleep, physical activity and social interaction. However, the baseline group did not have the ambient display and could only view the collected information via a web portal that summarizes the time spent in each activity as a fraction of the day. The multi-dimensional group has the ambient display.

5.2 Multi-dimensional wellbeing monitoring

Our first series of experiments show that BeWell+ is able to automatically monitor three dimensions of wellbeing (viz. physical activity, social interaction and sleep patterns). We find that (1) our automated wellbeing assessments compare well with commonly used paper-based medical wellbeing

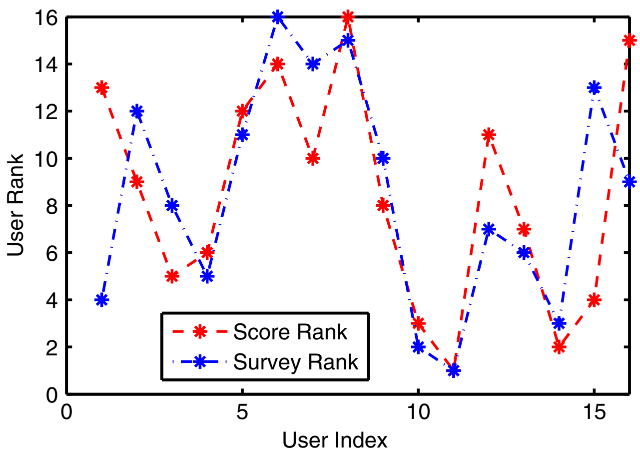

Fig. 7 Subjects are ranked using the BeWell+ physical activity score and the YPAS self-report survey. Both set of ranks agree fairly well and have a levenshtein similarity of 81.3 %

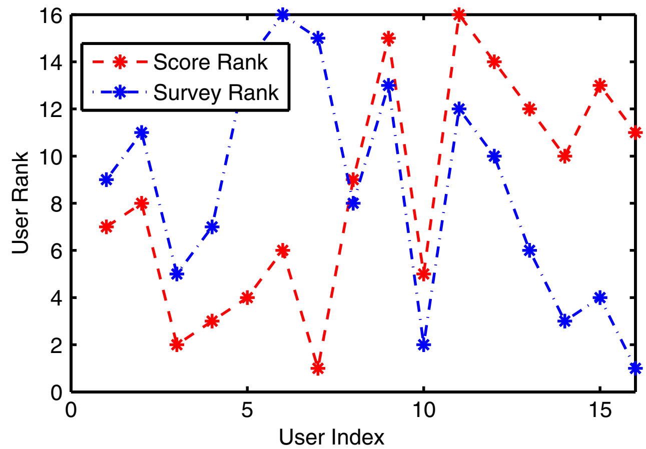

Fig. 8 Subjects are ranked using the BeWell+ social interaction score and the SF-36 self-report survey. The levenshtein similarity between the two sets of subject rankings is 56.3 %

surveys and (2) sensor-based estimates of sleep duration closely follow reported ground-truth provided by study participants.

Comparison to standard medical instruments The Yale Physical Activity Survey (YPAS) [19] and SF-36 [6] are standard ways of measuring physical and mental wellbeing. To measure the agreement between these surveys with the wellbeing scores produced by BeWell+ we rank all participants within the experiment, first by their aggregate wellbeing score and next by their survey results. Figure 7 visually shows the agreement between the two rankings for each subject when using either physical activity wellbeing scores or YPAS. The levenshtein similarity metric, shown in Table 1 is 81.3 %, indicating these rankings correlate well. We repeat this experiment for SF-36 and the social interaction wellbeing score, as shown in Fig. 8, and again find agreement - although not as strong as the last experiment. When comparing these two ranks the levenshtein similarity metric falls to 56.3 %. This result is understandable given mental health can have a greater variety of external factors that influence the outcome, as opposed to case of physical activity scores.

Sleep duration validation Finally, we perform an experiment to assess the accuracy of sleep duration during the field trial. Each subject self-reports the duration of their sleep once a day via an online survey. As shown in Table 2 our sleep model is able to estimate self-reported duration within $\pm$ 1.7 hours. This level of accuracy is inline with our finding in [34] that considers only 5 people. We are unable to examine the other inferences made by BeWell+ (viz. voicing, walking, stationary, running) because the collection of ground-truth necessary would have been too disruptive to the user experience of our study subjects. [34] reports the accuracy of these additional inferences which range between 85 % and 98 % (but again only for 5 individuals).

Table 1 Levenshtein similarity between user ranks based on medical surveys and BeWell+ wellbeing scores

	YPAS	SF36
Physical activity score	81.3 %	N/A
Social interaction score	N/A	56.3 %

5.3 Community adaptive wellbeing feedback

Our next experiments investigate two key aspects of wellbeing feedback: (1) the effectiveness of adapting feedback to keep implied healthy goals within realistic ranges for all users; and, (2) the benefit of multi-dimensional feedback, as observed in the behavioral decisions of our study population.

Adaptive wellbeing scoring in action To better understand how our adaptive wellbeing feedback can compensate for per user differences (e.g., lifestyles, occupation) we compare the use of adaptive and non-adaptive feedback on representative users from our study. Table 3 shows both forms of wellbeing scoring compared to different behavior changes which occur over the span of two days. In this table we examine two groups selected from the top and

Table 2 Error of sleep duration estimation

	RMSE	MAE
Duration error	2.2 hrs	1.6 hrs

Table 3 Under adaptive wellbeing feedback users continue to receive feedback even if they are well above (or well below) the ideal expectations of healthy behavior

High performance group (Sleep)			
	User A	User B	User C
Behavior change	-4 %	5 %	10 %
Baseline score	88	90	100
Adaptive score	36	67	100
Low performance group (Social)			
Behavior change	-2 %	5 %	10 %
Baseline score	10	55	58
Adaptive score	10	77	83

bottom 20 % of study subjects in sleep and social dimensions respectively. We refer to users in the top 20 % as *high performance group*, and those in the bottom 20 % as *low performance group*. In Table 3 user A from high performance group in the sleep dimension declines in performance by 4 % (≈ 0.3 hours). However, when using an unadaptive wellbeing scoring scheme her score remains high as she continues to far outperform the expectation of this unadaptive scheme. User B has increased performance by 5 % (≈ 0.4 hours), so she gets a high score under the baseline scoring. But within high performance group, her performance is in the middle, higher than user A but lower than user C - this fact is only reflected in the adaptive version of the wellbeing feedback. These users only receive personalized feedback when using an adaptive scoring system that understands their performance relative to their peers. Low performance group illustrates an identical scenario. These users from the bottom 20 % generally have low scores as they are far behind the performance of the overall user population. But if only compared with their low performance group counterparts, they will have significant changes in the scores, depending on their relative performance inside this group. Finally, Fig. 9 presents a time-series view of wellbeing scores (12 days) for three different users from our field trial. For each user we show their performance within a single wellbeing dimension. From this figure one can see that these users hardly receive informative feedback (e.g., their scores remain at 100) without adaptive wellbeing scoring. This is again caused by their behavior exceeding (negatively or positively) the expected norms of the unadaptive wellbeing score system. In Fig. 9, we also plot the relative percentile ranking of these users within their own group (the green curve). Clearly, the adaptive scores correspond much more closely to the users' actual peer performance compared to unadaptive scoring.

Multi-dimensional wellbeing feedback We measure the quantitative benefit of providing feedback along multiple

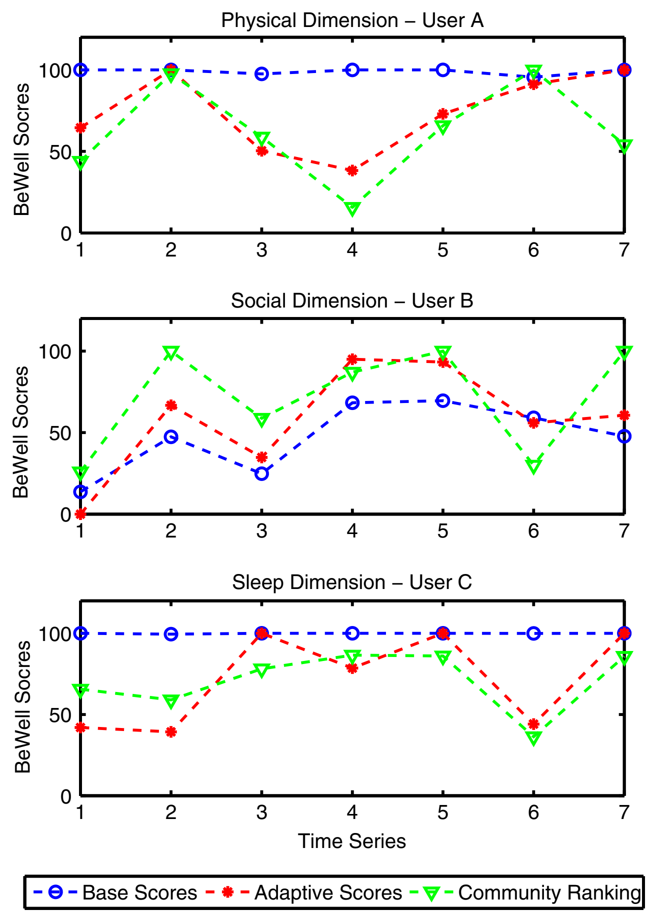

Fig. 9 By adapting wellbeing score functions users receive feedback that considers their relative position with peers who have similar lifestyles

dimensions by comparing the changes in wellbeing scores between our two study populations (viz. *multi-dimensional group* and *baseline group*). To compensate for individual variation that could bias results (i.e., participants that have abnormally high or low wellbeing scores) we compare any changes during the study relative to a baseline average score for each person along each dimension. The baseline score is calculated from data collected during the calibration phase just before the start of the study – none of the subjects had feedback or ambient display during the calibration phase. Figure 10 shows the average difference in the daily score for each person during the study period, relative to their personal baseline. This figure shows a significantly greater increase in score for *multi-dimensional group* compared to *baseline group*. Specifically, this outperformance is 105 % for physical activity, 88 % for social interaction and 507 % for sleep. Two-sample t-tests at the 95 % significance level indicate that these differences between *multi-dimensional group* and *baseline group* are all statistically significant ($p = 0.049$, $p < 0.01$ and $p = 0.04$ for the physical, social and sleep dimensions respectively).

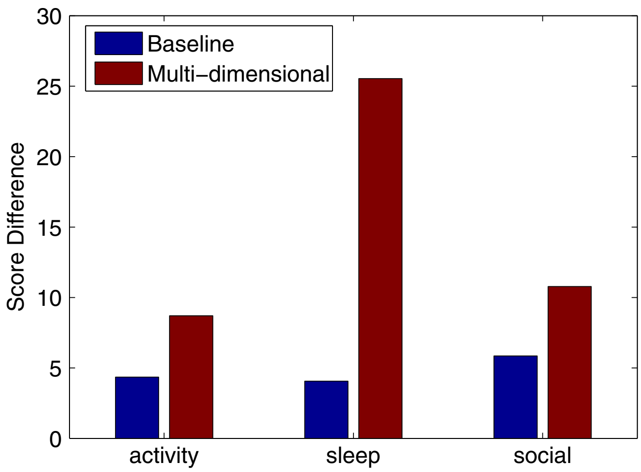

Fig. 10 Increases in user wellbeing are largest for those subjects who receive multi-dimensional wellbeing feedback

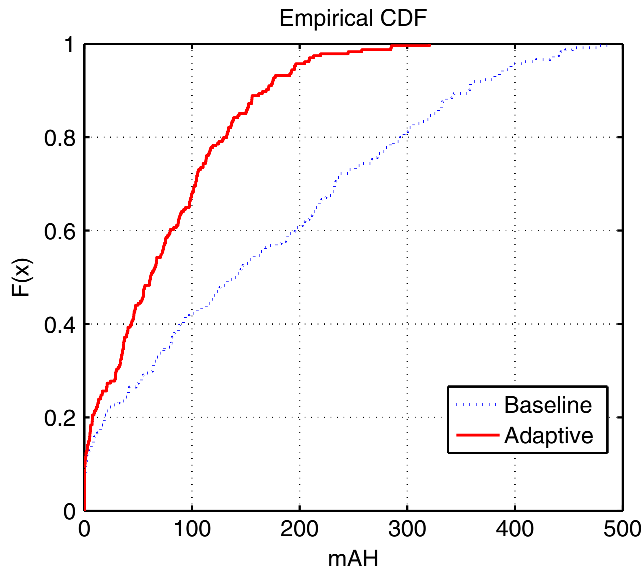

Fig. 11 CDF of per user daily energy consumption under wellbeing adaptive energy allocation compared to a hand-tuned baseline

5.4 Wellbeing adaptive resource allocation

In the following set of experiments, we investigate how efficiently BeWell+ manages smartphone energy while still closely monitoring user wellbeing.

Energy efficiency In this experiment we compare BeWell+'s adaptive resource management to a baseline in which BeWell+ performs no duty cycling. This baseline represents the upper bound accuracy of wellbeing scores with respect to errors that are caused by duty cycling. To compare these two schemes within identical experiment conditions we perform a trace based experiment. We begin by profiling the energy consumption of key energy consuming stages of our BeWell+ prototype when using both the adaptive and baseline approaches. We replay all 19 days of raw data sensor data for each participant, which we collect during our field trial. For each day of each participant we estimate the energy consumed, in addition to computing wellbeing scores.

Figure 11 shows a CDF of the average energy consumption for each day in this experiment. This figure shows our adaptive scheme is able to reduce average energy consumption by more than 50 % for 80 % of the days, which is approximately a 3-hour increase in battery life. Reductions in energy consumption should be considered in comparison to Fig. 12 which shows the impact to wellbeing score accuracy. For example, lowering energy consumption by 50 % results in approximately 18 points of error in the wellbeing score across all three dimensions. We consider this score difference, which corresponds to 5 % error in voice fraction measurement or a 1.4 hour error in sleep duration, tolerable given the large increases in energy efficiency that result.

Adaption to user wellbeing profile Figure 13 provides some further insight into the findings of the prior experiment. This figure illustrates the energy consumed for four representative subjects, and the relative allocation of energy to each sensor (and associated computation). For example, user B consumes the most energy as this subject has uniformly poor wellbeing scores across all dimensions, making it difficult to conserve energy from any one dimension. As expected in this case the allocation of energy between dimensions is evenly split. In contrast, user C uses significantly less energy as she has comparatively high wellbeing scores, allowing the adaptive scheme to lower energy used for these dimensions. The reason why the accelerometer is allocated a larger proportion of the energy budget for user

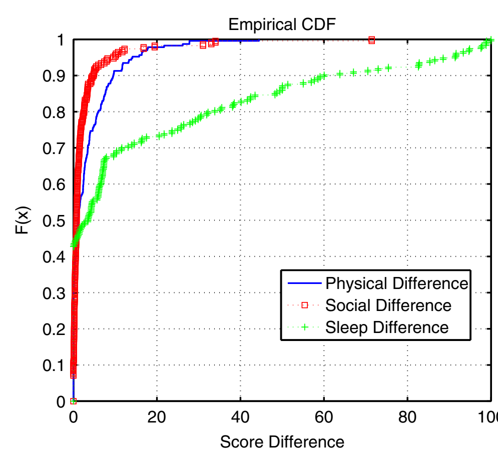

Fig. 12 CDF of per wellbeing dimension score error (i.e., score difference) under wellbeing adaptive energy allocation

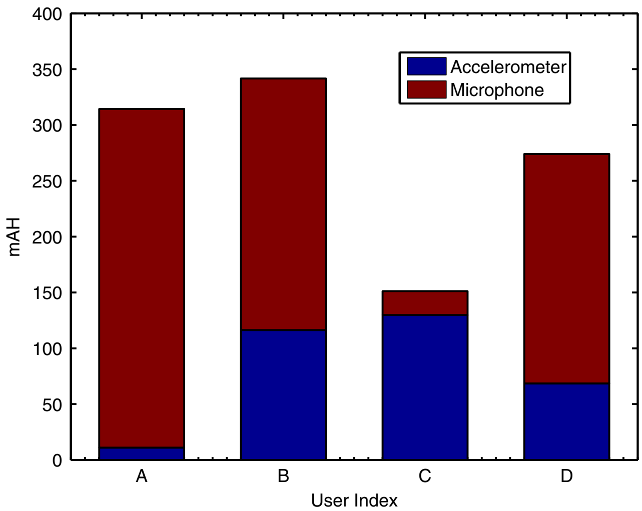

Fig. 13 Breakdown of daily energy consumption (by sensor) for four different BeWell+ users

C is that it is still the weakest dimension (in comparison to other dimensions).

5.5 Exit interview

In the remainder of this section we explore user reactions to: (1) the multi-dimensional ambient display; and, (2) subject attitudes and preferences to general usage of the BeWell+ app.

Reactions to ambient display Table 4 summarizes exit interview questions related to the ambient display. Participant responses indicate they have a positive reaction to the phone wallpaper as a means to visualize multi-dimensional well-being scores. A natural concern is that the use of multiple dimensions will overwhelm the user and they will not be able to easily digest the information. However, for example, question 2 in Table 4 shows that people overall had little difficulty in interpreting the ambient display.

Table 4 Ambient display results from exit interview

Survey questions	Answers
1. User would prefer different wallpaper	-1.00
2. Multi-dimensional Display easy to interpret	1.50
3. Multi-dimensional Scores helped keep balance	1.56
[-2: Strongly disagree, -1: Disagree, 0: Neutral, 1: Agree, 2: Strongly agree]	
4. I showed others my wallpaper	83.5 %
5. Animation was annoying	0.00 %
[Percentage of person choose]	

During exit interviews we discover friends and co-workers often casually ask *how is your fish today?* Many of the participants mention that they compare scores with other participants; 83.5 % of *multi-dimensional group* report that they show the display to their friends and colleagues. Exit surveys highlight an unexpected amount of social activity attributable to the ambient display in only a few weeks. Still, this enthusiasm may be due to a potentially short-lived novelty effect among subjects, this observation requires further testing as part of a long-term followup study.

From Table 4 we find very few subjects prefer an alternative wallpaper – we believe this number may rise when deployed in a broader population. During discussion we find that participants commonly turn off the phone screen when in more formal settings (e.g., meetings or while giving presentations) because of concerns it may be mistaken for a game or lead to them not being taken seriously by their peers. The ability to temporarily hide the display seems to be a necessary feature. Still, none of the subjects describe the visualization or the frequent animation as annoying (see question 4 in Table 4).

BeWell+ app experience We find 70 % of subjects believe that BeWell+ is a helpful and enjoyable app. A common theme with subjects is that they are surprised by what they learn from the study about their lifestyles. They report they find themselves motivated to actively change their daily behavior.

Encouraged by some early interview responses we decide to investigate some of the reasons for improved behavioral patterns during the study. We are curious if such increases are partially due to an improved ability within *multi-dimensional group* to connect everyday actions to wellbeing outcomes. To test this we perform a simple recall test. We show a timeline of participant wellbeing scores along different dimensions (viz. sleep, activity, social) and ask the participant to annotate and explain the variations seen in the timeline. Our findings show that the subjects that have access to multi-dimensional feedback on the phone are better able to connect life events to fluctuations in well-being. On average *multi-dimensional group* recalls 4.28 events per week compared to just 1.8 events for *baseline group*. Similarly, *multi-dimensional group* is able to recall a larger number of unique events as well. Common annotated events included: friends visiting for the weekend, change of (hospital) rotation, or pressure from work.

6 Related work

Recently, encouraging progress has been made towards mobile systems that can monitor and improve specific health goals – such as, the development of an open software architecture for mHealth [52]. Furthermore, various research prototypes have been demonstrated to reliably track a wide variety of key heath factors (e.g., sleep [7], stress [20], diet [46], smoking [10], mood [38]). Similarly, a number of persuasive systems [21] have been designed to assist people in making desired behavior changes and to motivate them to become, for example, more physically active [16]. Commercial activity is also increasing, with products such as, Nike+ [4] and DirectLife [5] becoming more prevalent as mobile health gains mainstream consumer acceptance.

However, a person's wellbeing is shaped by a diverse combination of health and lifestyle factors. Effective personal management of wellbeing requires apps that address a large variety of daily behaviors which have broad health related consequences. As a result, there is a growing interest in building mobile systems that take a broader health perspective. Some approaches rely on developing a software suite of separate mobile apps that manage multiple aspects of wellbeing (e.g., [8, 9]). In contrast, [33] and [46] take a more integrated approach to wellbeing management but rely on manual data entry in the form of a diary to collect information. AndWellness [30] utilizes a mixture of sensor-based activity inferences and manual data entry to provide a general monitoring platform for a range of wellbeing concerns. However, AndWellness is designed to monitor the user rather than promote behavior change. BALANCE [18] also combines user and sensor input to closely monitor multiple wellbeing factors (diet and physical activity), but it neglects other important wellbeing dimensions including emotional and social wellbeing. Finally, purpose-built sensor systems (e.g., Fitbit [2]) can automatically monitor multiple wellbeing relevant behaviors, such as sleep and physical activity while also providing user feedback; but – unlike BeWell+ – these solutions require the user to carry an additional sensor at all times.

Wellbeing feedback Prior research has also investigated how ambient displays, different types of goal settings, classifier accuracy, and user interaction affect mobile system's ability to encourage positive behavior changes (e.g., [16, 24, 40]). Ubifit Garden [16], one of the first mobile persuasion system for improving physical wellbeing uses the wallpaper of mobile phones to dynamically provide feedback about the different types of physical exercise performed by the user. The system presented in [40] links a player's daily foot step count to the growth and activity of a fish in a fish tank. Although researchers have recognized certain groups within a user population will benefit from personalized persuasive feedback (e.g., [17, 23]), existing persuasive systems still typically provide the same type of feedback across all users. Under BeWell+, each user receives wellbeing feedback automatically tuned to match their particular lifestyle patterns.

Energy allocation One of the most significant practical challenges to the everyday usage of mobile health systems is the resource limitations of smartphones (e.g., battery lifetime). Continuously sensing wellbeing states and providing real-time feedback will consume a significant fraction of mobile device energy. Many proposed solutions consider the general form of this problem and apply resource optimization and/or adaptation techniques (e.g., [13]) to address smartphone energy constraints while executing resource-expensive tasks. Recent research (e.g. [47, 50, 51]) has focused on minimizing the energy cost directly related to mobile sensing apps. For example, [50] tunes sensing pipelines in real-time both on mobile devices and in the cloud based context, available resources and the requirements of the app (i.e., social science experiment). However, unlike BeWell+, none of these systems are specifically designed to take the user's wellbeing into account while attempting to optimize resource usage on the phone.

7 Conclusion

In this article, we presented the next generation of the BeWell app – *BeWell+*, a smartphone app for monitoring and providing feedback across multiple dimensions of wellbeing. The primary goal of our field trial was to deploy BeWell+ to mainstream users in a real-world setting. Our deployment allowed us to both investigate fundamental issues that may influence the design of future generations of wellbeing apps and validate some of the assumptions that underpin BeWell+. Due to the relative short duration of this study it is not possible to make any claims of long-term behavioral change. The behavioral changes we do observe we believe are positive indications of the ability of BeWell+ to convey information and increase awareness. A longer-term field study and a more diverse population of users are both desired to further study BeWell+ and in particular the novel wellbeing mechanisms it introduces, namely, community adaptive feedback and wellbeing adaptive energy allocation.

References

1. Center for Disease Control and Prevention. http://www.cdc.gov
2. Fitbit. http://www.fitbit.com
3. National Sleep Foundation. http://www.sleepfoundation.org
4. Nike+. http://nikerunning.nike.com
5. Philips Directlife. http://www.directlife.philips.com

6. SF-36.org. A Community for Measuring Health Outcoming using SF tools. http://www.sf-36.org/tools/SF36.shtml
7. Shuteye. http://dub.washington.edu/projects/shuteye
8. bLife. http://www.myblife.com
9. Ahtinen A, Mattila E, Vaatanen A, Hynninen L, Salminen J, Koskinen E, Laine K (2009) User experiences of mobile wellness applications in health promotion: user study of wellness diary, mobile coach and selfrelax. In: Proceedings of Pervasive Health '09
10. Ali A, Hossain S, Hovsepian K, Rahman M, Plarre K, Kumar S (2012) mPuff: automated detection of cigarette smoking puffs from respiration measurements. In: Proceedings of IPSN '12
11. Alvarez G, Ayas N (2004) The impact of daily sleep duration on health: a review of the literature. Prog Cardiovasc Nurs 19(2):56
12. Ainsworth BE, Haskell WL, Whitt MC, Irwin ML, Swartz AM, Strath SJ, O'Brien WL, Bassett DR, Schmitz KH, Emplaincourt PO, Jacobs DR, Leon AS (2000) Compendium of physical activities: An update of activity codes and MET intensities. Med Sci Sports Exerc 32:9
13. Balan R, Satyanarayanan M, Park S, Okoshi T (2003) Tactics-based Remote execution for mobile computing. In: Proceedings of Mobisys '03
14. Bao L, Intille SS (2004) Activity recognition from user-annotated acceleration data. In: Proceedings of Pervasive '04
15. Bishop CM (2006) Pattern recognition and machine learning (Information Science and Statistics). Springer
16. Consolvo S, Klasnja P, McDonald DW, Avrahami D, Froehlich J, LeGrand L, Libby R, Mosher K, Landay JA (2008) Flowers or a Robot Army?: encouraging awareness and activity with personal, mobile displays. In: Proceedings of UbiComp '08
17. Consolvo S, Klasnja P, McDonald DW, Landay J (2010) Goal-setting considerations for persuasive technologies that encourage physical activity. In: Proceedings of Persuasive '09
18. Denning T, Andrew A, Chaudhri R, Hartung C, Lester J, Borriello G, Duncan G (2009) BALANCE: towards a usable pervasive wellness application with accurate activity inference. In: Proceedings of HotMobile '09
19. Dipietro L, Caspersen C, Ostfeld A, Nadel E (1993) A survey for assessing physical activity among older adults. Med Sci Sports Exerc 25:628–628
20. Ferreira P, Sanches P, Höök K, Jaensson T (2008) License to chill!: how to empower users to cope with stress. In: Proc of NordiCHI '08
21. Fogg BJ (2002) Persuasive technology: using Computers to change what we think and do. Ubiquity 2002 (December) 2
22. Fox K (1999) The influence of physical activity on mental well-being. Public Health Nutr 2(3a):411–418
23. Halko S, Kientz J (2010) Personality and persuasive technology: an exploratory study on health-promoting mobile applications. In: Proceedings of Persuasive '10
24. Hareva DH, Okada H, Kitawaki T, Oka H (2009) Supportive Intervention using a mobile phone in behavior modification. Acta Med Okayama 63(2):113–20
25. Penedo F, Dahn J (2005) Exercise and well-being: a review of mental and physical health benefits associated with physical activity. Curr Opin Psychiatr 18(2):189
26. Rozin P (2005) The meaning of food in our lives: a cross-cultural perspective on eating and well-being. J Nutr Educ Behav 37:S107–S112
27. George L, Blazer D, Hughes D, Fowler N (1989) Social support and the outcome of major depression. Brit J Psychiatry 154(4): 478
28. Knutson JF, Lansing CR (1990) The relationship between communication problems and psychological difficulties in persons with profound acquired hearing loss. J Speech Hear Disord 55(4):656
29. Blazer DG (1982) Social support and mortality in an elderly community population. Am J Epidemiol 115(5):684
30. Hicks J, Ramanathan N, Kim D, Monibi M, Selsky J, Hansen M, Estrin D (2010) Andwellness: an open mobile system for activity and experience sampling. In: Proceedings of wireless health '10
31. Paffenbarger R, Hyde R, Wing A, Hsieh C (1986) Physical activity, all-cause mortality, and longevity of college alumni. N Engl J Med 314(10):605–613
32. Klasnja P, Consolvo S, McDonald D, Landay J, Pratt W (2009) Using mobile and personal sensing technologies to support health behavior change in everyday life: lessons learned. In: Proceedings of AMIA 09
33. Lamminmaki E, Parkka J, Hermersdorf J, Kaasinen J, Samposalo K, Vainio J, Kolari J, Kulju M, Lappalainen R, Korhonen I (2005) Wellness diary for mobile phones. In: Proceedings of EMBEC '05
34. Lane ND, Mohammod M, Lin M, Yang X, Lu H, Ali S, Doryab A, Berke E, Choudhury T, Campbell A (2011) BeWell: a smartphone application to monitor, model and promote wellbeing. In: Proceedings of Pervasive Health '11
35. Lane ND, Xu Y, Lu H, Hu S, Choudhury T, Campbell AT, Zhao F (2011) Enabling large-scale human activity inference on smartphones using community similarity networks (CSN). In: Proceedings of Ubicomp '11
36. Lane ND, Xu Y, Lu H, Campbell A, Choudhury T, Eisenman S (2011) Exploiting social networks for large-scale human behavior modeling. Pervasive Comput IEEE 10(4):45–53
37. Lester J, Choudhury T, Kern N, Borriello G, Hannaford B (2005) A Hybrid discriminative/generative approach for modeling human activities. In: Proceedings of the IJCAI '05
38. LiKamWa R, Liu Y, Lane N, Zhong L (2011) Can Your Smartphone Infer Your Mood? In: Proceedings of PhoneSense '11
39. Pilcher J, Ginter D, Sadowsky B (1997) Sleep quality versus sleep quantity: relationships between sleep and measures of health, well-being and sleepiness in college students. J Psychosom Res 42(6):583–596
40. Lin JJ, Mamykina L, Lindtner S, Delajoux G, Strub HB (2006) Fish'n'steps: encouraging physical activity with an interactive computer game. In: Proc. of UbiComp '06
41. Lin M, Lane ND, Mohammod M, Yang X, Lu H, Cardone G, Ali S, Doryab A, Berke E, Campbell A, Choudhury T (2012) BeWell+: multi-dimensional wellbeing monitoring with community-guided user feedback and energy optimization. In: Proceedings of wireless health '12
42. Lu H, Pan W, Lane ND, Choudhury T, Campbell AT (2009) Soundsense: scalable sound sensing for people-centric applications on mobile phones. In: Proceedings of MobiSys '09
43. Lu H, Yang J, Liu Z, Lane ND, Choudhury T, Campbell AT (2010) The jigsaw continuous sensing engine for mobile phone applications. In: Proceedings of Sensys '10
44. Miluzzo E, Oakley J, Lu H, Lane N, Peterson R, Campbell A (2008) Evaluating the iphone as a mobile platform for people-centric sensing applications. In: Proceedings of UrbanSense '08
45. Norris R, Carroll D, Cochrane R (1992) The effects of physical activity and exercise training on psychological stress and well-being in an adolescent population. J Psychosom Res 36(1):55–65
46. Patrick K, Raab F, Adams M, Dillon L, Zabinski M, Rock C, Griswold W, Norman G (2009) A text message–based intervention for weight loss: randomized controlled trial. J Med Internet Res 11(1)

47. Wang Y, Lin J, Annavaram M, Jacobson Q, Hong J, Krishnamachari B, Sadeh N (2009) A framework of energy efficient mobile sensing for automatic user state recognition. In: Proceedings of Mobisys ’09
48. Yaggi H, Araujo A, McKinlay J (2006) Sleep duration as a risk factor for the development of type 2 diabetes. Diabetes Care 29(3):657
49. Lane ND, Miluzzo E, Lu H, Peebles D, Choudhury T, Campbell A (2010) A survey of mobile phone sensing. IEEE Commun Mag 48:140–150
50. Rachuri KK, Mascolo C, Musolesi M, Rentfrow P (2011) SociableSense: exploring the trade-offs of adaptive sampling and computation offloading for social sensing. In: Proceedings of Mobicom ’11
51. Chu D, Lane ND, Lai TT, Pang C, Meng X, Guo Q, Li F, Zhao F (2011) Balancing energy, latency and accuracy for mobile sensor data classification. In: Proceedings of Sensys ’
52. Estrin D, Sim I (2011) Open mHealth architecture: an engine for health care innovation. Science