Taylor & Francis
Taylor & Francis Group

IISE Transactions on Healthcare Systems Engineering

ISSN: 2472-5579 (Print) 2472-5587 (Online) Journal homepage: https://www.tandfonline.com/loi/uhse21

Longitudinal planning for personalized health management using daily behavioral data

Cao Xiao, Shupeng Gui, Ji Liu, Yu Cheng, Xiaoning Qian & Shuai Huang

To cite this article: Cao Xiao, Shupeng Gui, Ji Liu, Yu Cheng, Xiaoning Qian & Shuai Huang (2019) Longitudinal planning for personalized health management using daily behavioral data, IISE Transactions on Healthcare Systems Engineering, 9:3, 243-249, DOI: 10.1080/24725579.2019.1640814

To link to this article: https://doi.org/10.1080/24725579.2019.1640814

Published online: 25 Sep 2019.

IISE TRANSACTIONS ON HEALTHCARE SYSTEMS ENGINEERING
2019, VOL. 9, NO. 3, 243–249
https://doi.org/10.1080/24725579.2019.1640814

Taylor & Francis
Taylor & Francis Group

Longitudinal planning for personalized health management using daily behavioral data

Cao Xiao[a†], Shupeng Gui[b†], Ji Liu[b], Yu Cheng[c], Xiaoning Qian[d], and Shuai Huang[e]

[a]Analytics Center of Excellence, IQVIA, Cambridge, USA; [b]University of Rochester, Rochester, New York, USA; [c]Microsoft AI & Research, Redmond, USA; [d]Texas A&M University, College Station, Texas, USA; [e]University of Washington, Seattle, Washington, USA

ABSTRACT
Mitigating globally emerging health problems such as obesity needs scalable solutions that can facilitate health management and promote healthier lifestyles outside of clinical settings. Such scalable solutions, while targeting general population, need to provide personalized behavior change plans that not only fit users' own underlying physiologic dynamics but also suit their preferences and needs. There has been fast-growing development of mobile health devices and applications for monitoring of human behavior (such as physical activity and food intake) and health status such as BMI. However, there are challenges to translate these noisy and dynamic behavioral data into personalized longitudinal planning. To address such challenges, we develop an integrated framework that unifies dynamic modeling, sparse learning, dictionary learning and matrix completion to translate users' behavioral data into personalized dynamic system models and use them as constraints for deriving deeply personalized longitudinal health plans. We evaluate the proposed framework on a real-world user behavioral dataset and demonstrate its promising utility and efficacy.

ARTICLE HISTORY
Received 17 September 2018
Accepted 3 July 2019

KEYWORDS
Wearable data; health monitoring; dynamical modeling

1. Introduction

Emerging mobile health technologies, such as mobile applications and wearable devices, have provided health professionals unprecedented monitoring and management capacity to materialize the envisioned personalized and preventative healthcare, particularly for those health management problems that need scalable solutions to promote healthier lifestyles outside of clinical settings, e.g., weight management. In addition, they collect large amounts of longitudinal behavioral data from users. These data are hypothesized to encode underlying user-specific physiological dynamics that govern the latent relations between the behavioral variables (e.g., activities and food intake) and target variables (Liu et al., 2015). Such physiological dynamics, if fully explored, could be useful for designing personalized behavior change plans that are optimized to maximize health improving effect for the targeting users (Spring, Gotsis, Paiva, & Spruijt-Metz, 2013).

However, it is challenging to learn physiological dynamics from longitudinal user behavior data due to their noisy and irregular-sampled nature and lots of missing values (Qiao et al., 2015). Moreover, to leverage the learned dynamics, e.g. provide health recommendation, the objectives need to be integrated with the physiological dynamics learning, so that the recommendation could be truly personalized to each user, and the framework could mimic what doctors have been doing on a larger scale, and in an intelligent and automatic fashion.

As summarized in Section 2, various approaches were proposed to either learn physiological patterns from behavioral data or provide personalized health recommendation or both. However, the existing strategies only exploit limited value of these data so that feedback to individuals is often limited to either overall statistics (Consolvo, Klasnja, McDonald, & Landay, 2014), visualization of self-tracked data (Fitbit Inc.) or generic suggestions not being personalized to users' lifestyles (Kukafka, 2005). Note that although there are many studies that leverage deep models for pattern mining and use reinforcement learning for planning, they do not fit our case. Deep learning requires large training data and is a bit "black-box," while behavioral models can only be learned from limited data and need to be interpretable. For reinforcement learning models, they learn optimal policy according to predefined reward functions, but in our case there is no clear reward function. Therefore, to the best of our knowledge, there still lacks an integrated method for simultaneous accurate dynamic modeling and personalized planning. To make the best use of the data, we recognize the following learning and planning challenges:

- *Challenges in learning from noisy behavior data with considering delayed effect*: Besides underlying complex

CONTACT Cao Xiao cao.xiao@iqvia.com Analytics Center of Excellence, IQVIA.
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/uhse.
†These authors contributed equally to this work.

dynamics, the missing values and outliers in users' behavioral data present difficulty in learning (Fung & Sheung, 2006; Zhang, Meratnia, & Havinga, 2010). In addition, there exists delayed effect in users' health outcome.

- *Challenges in personalized longitudinal planning*: It is challenging to formulate the optimal planning on the foundation of the dynamic model with simultaneous (1) complying with users' personal physiological dynamics that health planning needs; (2) incorporate users' preferences and (3) referencing peers' behavioral routines.

To address the arising challenges, we develop a longitudinal planning framework that firstly learn personal physiological dynamics with simultaneous missing value imputation and outlier detection, and then the learned dynamic mechanism is used to guide personalized longitudinal planning with simultaneous consideration of users' preferences and peers' examples, all in the forms of constraints. To summarize, we propose the following contributions:

- *Dynamic modeling*: We propose to learn personal health dynamic with dynamic System identification with Simultaneous Missing value estimation and Outlier detection (SSMO). It automatically removes the effects of outliers in the dataset, imputes missing values and conducts model identification, with considering delayed effect of health outcomes.
- *Longitudinal planning*: We propose a longitudinal planning method that learns optimal behavior change plans guided by personal physiologic dynamics. To improve user adoption, we further formulate their preferences and needs as constraints, and constructing an action polyhedron construction (APC) engine constructed using dictionary learning for each user to learn from peers. It uncovers the regularity underlying heterogeneous human behavior, as well as provides users with more flexibility to learn from feasible actions of peers.
- *Efficient solution*: We developed efficient algorithms to solve the learning and planning problems with specific optimization strategies to ensure the feasibility and robustness of the algorithms. In particular, block coordinate descent for SSMO and ADMM for APC engine.

2. Related work

To model behavior change, Rabbi, Aung, Zhang, and Choudhury (2015) proposed to use the multi-armed Bandit algorithm (Gittins, Glazebrook, & Weber, 2011) to automate the behavioral change plan; however, it needs pre-specified behavioral change options. Our method is also remotely related to time series analysis. Time series analysis with handling missing values or outliers include mean value filling, cubic fitting, polynomial fitting for missing value estimation (Fung & Sheung, 2006) and autoregressive moving average and vector autoregression for outlier detection (Zhang et al., 2010). However, these methods are based on the assumptions very different from our problem setting, and often adopt the two-stage strategy (pre-process the data first and then fit the model, or in reverse) that is suboptimal. Also, they rarely address missing values and outliers simultaneously under a consistent model assumption.

In recommendation system research, Koren (2009) proposed to model the temporal dynamics of the user's preference during the whole time period and applied this in a factorization model. Xiong, Chen, Huang, Schneider, and Carbonell (2010) formalized Bayesian probabilistic tensor factorization model with a special constraint on the time dimension for the temporal recommending setting. However, these methods are motivated by temporal changes in customer preferences, which is fundamentally different from our dynamic system perspective. Also, most of the existing methods on recommendation systems focus on short-term recommendation, our method aims to provide long-term planning that depicts the trajectory from the user's current health status to the target status.

3. Method

Our proposed learning pipeline comprises three major components: (1) a dynamic system identification engine that automatically handles missing values and outliers; (2) a dictionary learning approach for action polyhedron construction and (3) a personalized longitudinal planning algorithm.

3.1. Dynamic SSMO detection engine

The learning component, SSMO, is the first step of the proposed system. Given a user record with a series of observations, $\{\mathbf{x}_t, \mathbf{u}_t\}_{t=1}^{T}$, where $\mathbf{x}_t$ denotes the outcome at time t which can be extended to any dimension and $\mathbf{u}_t$ denotes the behavior profile as a column vector at time t, we propose to adopt a linear dynamic system as the underlying model to characterize the relationships between the behavioral and outcome variables. The linear dynamic system is a very flexible model that can characterize a wide range of dynamics. For instance, in what follows we illustrate the specific use of a linear dynamic system to model the 3rd-order dynamics:

$$\begin{bmatrix} \mathbf{x}_{t+1} \\ \mathbf{x}_t \\ \mathbf{x}_{t-1} \end{bmatrix} = A \begin{bmatrix} \mathbf{x}_t \\ \mathbf{x}_{t-1} \\ \mathbf{x}_{t-2} \end{bmatrix} + B \begin{bmatrix} \mathbf{u}_t \\ \mathbf{u}_{t-1} \\ \mathbf{u}_{t-2} \end{bmatrix} + C + \mathbf{w}_t, \qquad (1)$$

where $\mathbf{w}_t$ is white noise and C is a bias term. Such a formulation can capture both spontaneous effect (i.e., from $\mathbf{u}_t$ to $\mathbf{x}_t$) as well as delayed effect (i.e., from $\mathbf{u}_{t-2}$ to $\mathbf{x}_{t+1}$). Apparently, this formulation is generic and can be further extended to capture higher-order dynamics. It can also be recognized as an equivalent form with the common continuous linear dynamic system that models $[\mathbf{x}(t); \mathbf{x}.(t); \mathbf{x}..(t)]$ (Antsaklis & Michel, 2007) if we rewrite $[\mathbf{x}_t; \mathbf{x}_{t-1}; \mathbf{x}_{t-2}]$ as $[\mathbf{x}_t; (\mathbf{x}_t - \mathbf{x}_{t-1}); (\mathbf{x}_t + \mathbf{x}_{t-2} - 2\mathbf{x}_{t-1})]$ (note that the parameters A, B and C would be different then).

The basic idea of SSMO for better dynamic system identification is to simultaneously impute missing values, handle outliers, take care of delayed effects and train the dynamic model, so that the imputation of missing values and detection of outliers are in a proper context defined by the learned dynamic model.

Let $\mathbf{X} = [\mathbf{x}_0, \mathbf{x}_1, ..., \mathbf{x}_T]$ be the state matrix and $\mathbf{U} = [\mathbf{u}_0, \mathbf{u}_1, ..., \mathbf{u}_{T-1}]$ be the action matrix. Ω_x and Ω_u represent the observed elements in X and U, respectively, with their complement sets denoting missing values. Define $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_0, \hat{\mathbf{x}}_1, ..., \hat{\mathbf{x}}_T]$ and $\hat{\mathbf{U}} = [\hat{\mathbf{u}}_0, \hat{\mathbf{u}}_1, ..., \hat{\mathbf{u}}_{T-1}]$ as the estimates of $\mathbf{X}$ and $\mathbf{U}$. We propose to learn $\hat{\mathbf{X}}$ and $\hat{\mathbf{U}}$ in SSMO to be consistent with $\mathbf{X}$ and $\mathbf{U}$ on the observed sets Ω_x, Ω_u, respectively. The identified outlier and missing elements in $\mathbf{X}$ are essentially free variables in $\hat{\mathbf{X}}$.

$$\min_{\substack{A,B,C,\\ \hat{\mathbf{X}},\hat{\mathbf{U}}}} \frac{1}{2}\sum_{t=0}^{T-1}\left\| \begin{bmatrix} \hat{\mathbf{x}}_{t+1} \\ \hat{\mathbf{x}}_t \\ \hat{\mathbf{x}}_{t-1} \end{bmatrix} - \left(A \begin{bmatrix} \hat{\mathbf{x}}_t \\ \hat{\mathbf{x}}_{t-1} \\ \hat{\mathbf{x}}_{t-2} \end{bmatrix} + B \begin{bmatrix} \hat{\mathbf{u}}_t \\ \hat{\mathbf{u}}_{t-1} \\ \hat{\mathbf{u}}_{t-2} \end{bmatrix} + C \right) \right\|^2 \tag{2a}$$

$$\text{s.t.} \quad ||(\hat{\mathbf{X}}-\mathbf{X})_{\Omega_x}||_0 \le a; \quad ||(\hat{\mathbf{U}}-\mathbf{U})_{\Omega_u}||_0 \le b. \tag{2b}$$

The objective (2a) is a squared loss function to evaluate the goodness-of-fit of the system parameters A, B, C and the estimates $\hat{\mathbf{U}}, \hat{\mathbf{X}}$. The constraints (2b) are to control the number of different estimated elements from the observed elements in Ω_x and Ω_u, which essentially controls the number of outliers among the observed elements. The parameters a and b restrict the maximal number of outliers. When they are set to 0, Eq. (2) only handles missing values. The values for a and b actually are not hard to decide. For example, we could estimate the upper bound of the percentage of outliers that is easily accessible in many applications. This algorithm is indeed robust as long as a and b are greater than the actual number of outliers but not far away from it.

3.1.1. Solving SSMO by block coordinate descent

We apply block coordinate descent (BCD) (Tseng, 2001) to solve Eq. (2) by alternatively optimizing two groups of variables $\{A, B, C\}$ and $\{\hat{\mathbf{X}}, \hat{\mathbf{U}}\}$:

- To optimize $\{A, B, C\}$, it is a least squares optimization with a closed-form solution.
- To optimize $\{\hat{\mathbf{X}}, \hat{\mathbf{U}}\}$, we adopt the projected gradient descent method to iteratively update:

$$\hat{\mathbf{X}}_{k+1} = \arg\min_{\hat{X}}\{||\hat{\mathbf{X}}-(\hat{\mathbf{X}}_k-\gamma g_{\hat{X}_k})||_F^2; \\ \text{s.t.}||(\hat{\mathbf{X}}-\mathbf{X})_{\Omega_x}||_0 \le a\},$$

where $g_{\hat{X}_k}$ is the partial derivative of the objective function (2a) w.r.t. $\hat{\mathbf{X}}_k$; γ is the step size that could be chosen to be a sufficiently small constant and $||\cdot||_F$ denotes Frobenius norm. It actually also admits a closed-form solution that can be found by: first, selecting a elements in $(\hat{\mathbf{X}}_k-\gamma g_{\hat{\mathbf{X}}_k}-\mathbf{X})_{\Omega_x}$ with the largest magnitudes as the outliers at the current iteration and forming a set S; second, setting the elements outside of Ω_x and in set S: $(\hat{\mathbf{X}}_{k+1})_{\bar{\Omega}_x\cup S} = (\hat{\mathbf{X}}_k-\gamma g_{\hat{\mathbf{X}}_k})_{\bar{\Omega}_x\cup S}$; third, setting the remaining elements in $\hat{\mathbf{X}}_{k+1}$ to take the same values in $\hat{\mathbf{X}}_k$. To update $\hat{\mathbf{U}}_{k+1}$ from $\hat{\mathbf{U}}_k$, one can follow a similar procedure. Due to the space limit, we omit the detailed derivation. We summarize all the steps in Algorithm 1.

Algorithm 1. BCD for SSMO

Require: $\mathbf{X}_{\Omega_x}, \mathbf{U}_{\Omega_u}, a, b$
Ensure: $A, B, C, \hat{\mathbf{X}}, \hat{\mathbf{U}}$
1: **repeat**
2: Optimize A, B, C by minimizing the least squares problem (2a) without any constraint.
3: Optimize $\hat{\mathbf{X}}$: Select top a largest elements in $(\hat{\mathbf{X}}-\gamma g_{\hat{X}}-\mathbf{X})_{\Omega_x}$, which forms the index set S; Update elements of $\hat{\mathbf{X}}$ in $\bar{\Omega}_x \cup S$ by $(\hat{\mathbf{X}})_{\bar{\Omega}_x\cup S} \leftarrow (\hat{\mathbf{X}}-\gamma g_{\hat{X}})_{\bar{\Omega}_x\cup S}$.
4: Optimize $\hat{\mathbf{U}}$: similar to the updates of $\hat{\mathbf{X}}$;
5: **until** convergence.
6: **return** $A, B, C, \hat{\mathbf{X}}, \hat{\mathbf{U}}$;

3.2. Personalized longitudinal planning

The outputs of SSMO are used to constrain the personalized longitudinal planning. Specifically, it is to identify an optimal sequence of actions, denoted by $\mathbf{u}_0, \mathbf{u}_1, ..., \mathbf{u}_{T-1}$, to drive the user's initial health status $\mathbf{x}_0$ to the target status $\mathbf{x}_T$ in T days. For example, with BMI as the health status, the user may want to reduce BMI from the current level $\mathbf{x}_0 = 30$ to the target 28 in 90 days. With a dynamic model, any proposed planning can be evaluated with the predicted future health status. The challenge is how to utilize this capacity to derive the optimal planning. On the other hand, we should formalize the user's preferences as optimization constraints to enhance the quality of the generated optimal planning. This leads to the following formulation:

$$\min_{\mathbf{u}_0, \mathbf{u}_1, ..., \mathbf{u}_{T-1}} \quad \sum_{t=1}^{T} \mathbf{c}^\top \mathbf{u}_t + \lambda \sum_{t=1}^{T-1} ||\mathbf{u}_t-\mathbf{u}_{t-1}||_1 \tag{3a}$$

$$\text{s.t.} \quad \hat{\mathbf{x}}_T \le \text{target} \tag{3b}$$

$$\mathbf{u}_- \le \mathbf{u}_t \le \mathbf{u}_+ \tag{3c}$$

$$\mathbf{u}_0 = \mathbf{h} \tag{3d}$$

$$\mathbf{u}_t \in \text{conv}(D). \tag{3e}$$

The objective function (3a) consists of two terms: The first term is to measure the cost of the adopted action, as different users might have different preferences or difficulties in conducting the actions; and the second term is to measure the smoothness of actions across all time points, assuming that users do not like sudden changes between consecutive actions as what the low-effort theory implies (Fogg, 2009). The first constraint (3b) is to ensure that, by following the planning, the user will achieve the pre-specified goal, where $\hat{\mathbf{x}}_T$ is the estimated final health status using the dynamic model (1). Specifically, based on the initial status $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ and the action sequence $\mathbf{u}_0, \mathbf{u}_1, ..., \mathbf{u}_{T-1}, \hat{\mathbf{x}}_T$ can be estimated from $[\hat{\mathbf{x}}_T; \hat{\mathbf{x}}_{T-1}; \hat{\mathbf{x}}_{T-2}] = A^{T-2}[\mathbf{x}_2; \mathbf{x}_1; \mathbf{x}_0] + \sum_{t=2}^{T-1} A^{T-t-1}B[\mathbf{u}_t; \mathbf{u}_{t-1}; \mathbf{u}_{t-2}]$. The second box constraint (3c) is to avoid unwanted and

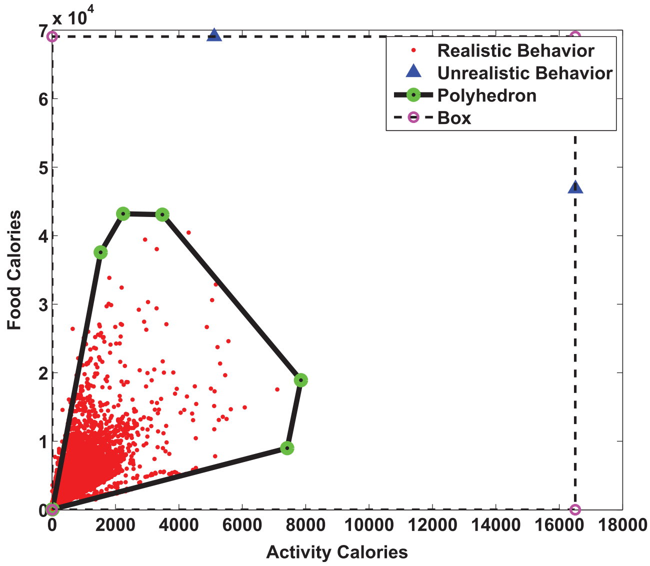

Figure 1. The constructed action polyhedron D learned from over 10,000 behavioral actions collected from over 30 users.

unrealistic actions for a specific user. The third constraint (3d) ensures the recommendation to start from the habits $\mathbf{h}$ of the user. We further impose another constraint (3e) to enforce the recommended actions to be within a set $\text{conv}(D)$. This is named as the *action polyhedron* that specifies the action space, in which a realistic planning can be constructed. It provides us great flexibility to incorporate domain knowledge or any kind of prior knowledge to guide the personalized longitudinal planning. In Section 3.3, we will introduce a learning-from-peers approach to construct this action space by a novel dictionary learning approach.

3.3. APC engine

We now introduce the dictionary learning approach to construct the action polyhedron $\text{conv}(D)$ as the feasible region of actions, ensuring the recommendations are realistic and reasonable in practice. In particular, $\text{conv}(D)$ can be viewed as a summary of the typical action patterns of other users that form a potential mentor group for the target user.

Let $\mathbf{U} \in \mathbf{R}^{p\times n}$ be the action matrix that we could use to learn $\text{conv}(D)$, i.e., each column represents an action vector that has been undertaken in real life by a certain user. $\mathbf{U}$ can be constructed by collecting n actions from different users whose behavioral patterns could inspire the planning for the target user. Then, the formulation of the dictionary learning can be written as follows (Eq. (4)):

$$\begin{aligned} \min_{w} \quad & \frac{1}{2}||\mathbf{U}-\mathbf{U}\mathbf{W}||_F^2+\gamma||\mathbf{W}||_{2,1} \\ \text{s.t.} \quad & \mathbf{W}\geq 0, 1^\top\mathbf{W}=1. \end{aligned} \tag{4}$$

The first term is to measure the approximation adequacy to represent all actions (vectors) in $\mathbf{U}$ using a set of basis actions (vectors). The second term is to enforce sparsity in the basis matrix $\mathbf{W}$, i.e., the nonzero columns of the learned $\mathbf{W}$ indicate the selected actions. Figure 1 illustrates the goal of finding a polyhedron to represent the whole action set $\mathbf{U}$ using a convex hull as the action polyhedron D. Figure 1 is

Table 1. Comparison of estimation error (in RMSE)

SSMO	Mean value	Last value	Last + Med
0.67 ± 0.49	1.31 ± 2.13	1.06 ± 1.06	0.86 ± 0.57
PACE	MICE	Amelia	MI
0.71 ± 0.55	1.34 ± 0.46	1.83 ± 0.95	1.30 ± 0.44

based on over 10,000 behavioral actions collected from over 30 users. It shows that the behavioral actions undertaken by this cohort exhibit a clear regularity, indicating that human behavioral actions follow certain principles and are not totally random. Thus, to generate personalized longitudinal planning, it is required that the planning should consist of reasonable actions that fit the "human patterns." Further, Figure 1 shows that the dictionary learning formulation provides a very effective approach to extract patterns and summarize the massive data matrix $\mathbf{U}$.

Challenges in Optimization: Note that the proposed dictionary learning method is different from the existing methods that have been used in pattern recognition (Mairal, Bach, Ponce, & Sapiro, 2009; Yeh & Yang, 2012), and event detection (Cong, Yuan, & Liu, 2011). To solve Eq. (4), we face two challenges: The first one is that Eq. (4) involves high dimensional $\mathbf{W} \in R^{n\times n}$; the other challenge lies on that Eq. (4) includes nonsmooth regularization term and constraints. It takes hours to solve it when $n = 1000$ if using general solvers, for example, CVX (CVX Research, 2012). In the following, we propose an efficient algorithm to solve it with ADMM. First, we duplicate the variable $\mathbf{W}$ with another variable $\mathbf{V}$, and rewrite Eq. (4) as follows:

$$\begin{aligned} \min_{w,v} \quad & \frac{1}{2}||\mathbf{U}-\mathbf{U}\mathbf{W}||_F^2+\gamma||\mathbf{V}||_{2,1} \\ \text{s.t.} \quad & \mathbf{V}\geq 0, 1^\top\mathbf{W}=1^\top, \mathbf{W}=\mathbf{V}. \end{aligned} \tag{5}$$

Then, we define the augmented Lagrangian of Eq. (5) and further convert it to be

$$\begin{aligned} L_\rho(\mathbf{W},\mathbf{V},\Lambda) = & \frac{1}{2}||\mathbf{U}-\mathbf{U}\mathbf{W}||_F^2+\frac{\rho}{2}||\mathbf{W}-\mathbf{V}||_F^2 \\ & +\gamma||\mathbf{V}||_{2,1}+\langle\Lambda,\mathbf{W}-\mathbf{V}\rangle \\ & +\mathbf{I}_{\mathbf{V}\geq 0}(V)+\mathbf{I}_{1^\top\mathbf{W}=1^\top}(\mathbf{W}), \end{aligned} \tag{6}$$

where $\mathbf{I}_{\text{condition}(\cdot)}$ is the Delta function: It gives zero if the condition is satisfied by the augment; Otherwise, it returns $+\infty$. ρ could be an arbitrary positive number.

Following the ADMM procedure, we then iterate over three steps:

Minimize $L_\rho(W, V, \lambda)$ **w.r.t. W:** It essentially solves the following optimization:

$$\min_{W} \frac{1}{2}||\mathbf{U}-\mathbf{U}\mathbf{W}||_F^2+\frac{\rho}{2}||\mathbf{W}-\mathbf{V}||_F^2+\langle\Lambda,\mathbf{W}\rangle+\mathbf{I}_{1^\top\mathbf{W}=1^\top}(\mathbf{W}).$$

By deriving the KKT conditions (Boyd & Vandenberghe, 2004), it is equivalent to solve

$$\begin{bmatrix} \mathbf{U}^\top\mathbf{U}+\rho\mathbf{I} & 1 \\ 1^\top & 0 \end{bmatrix}\begin{bmatrix} \mathbf{W} \\ \Phi^\top \end{bmatrix}=\begin{bmatrix} \mathbf{U}^\top\mathbf{U}+\rho\mathbf{V}-\Lambda \\ 1^\top \end{bmatrix}, \tag{7}$$

where Φ is the dual variable. By solving the Eq. (7), we obtain a closed-form solution as:

$$\mathbf{W} = (\mathbf{I} - N \cdot 1^\top)A^{-1}B + N \cdot 1^\top,$$
$$N = \frac{A^{-1}1}{1^\top A^{-1}1}, \quad A = U^\top U + \rho\mathbf{I}, B = U^\top U + \rho V - \Lambda, \tag{8}$$

and $\mathbf{I}$ is an identity matrix.

Minimize $L_\rho(\mathbf{W}, \mathbf{V}, \Lambda)$ w.r.t. **V:** It essentially solves the following optimization:

$$\min_{V} \frac{\rho}{2}||\mathbf{W} - \mathbf{V}||_F^2 - \langle \Lambda, \mathbf{V} \rangle + \gamma||\mathbf{V}||_{2,1} + \mathbf{I}_{\mathbf{V} \geq 0}(\mathbf{V}). \tag{9}$$

Update Λ**:** This step mimics the dual gradient ascent: $\Lambda = \Lambda + \rho(\mathbf{W} - \mathbf{V})$. We finally summarize the algorithm in Algorithm 2.

Algorithm 2. ADMM for APC engine

Require: $\mathbf{U}$, $\rho > 0$, and γ
Ensure: $\mathbf{W}$
1: **repeat**
2: Minimize $L_\rho(\mathbf{W}, \mathbf{V}, \Lambda)$ in terms of $\mathbf{W}$ and update $\mathbf{W} \leftarrow (I - N1^\top)A^{-1}B + N1^\top$
3: Minimize $L_\rho(\mathbf{W}, \mathbf{V}, \Lambda)$ in terms of $\mathbf{V}$ and update $V_{i\cdot} \leftarrow \max\left(0, 1 - \frac{\beta}{||\mathbf{Y}_{i\cdot}||}\right) \cdot \mathbf{Y}_{i\cdot}$;
4: $\Lambda \leftarrow \Lambda + \rho(\mathbf{W} - \mathbf{V})$;
5: **until** convergence
6: **return W;**

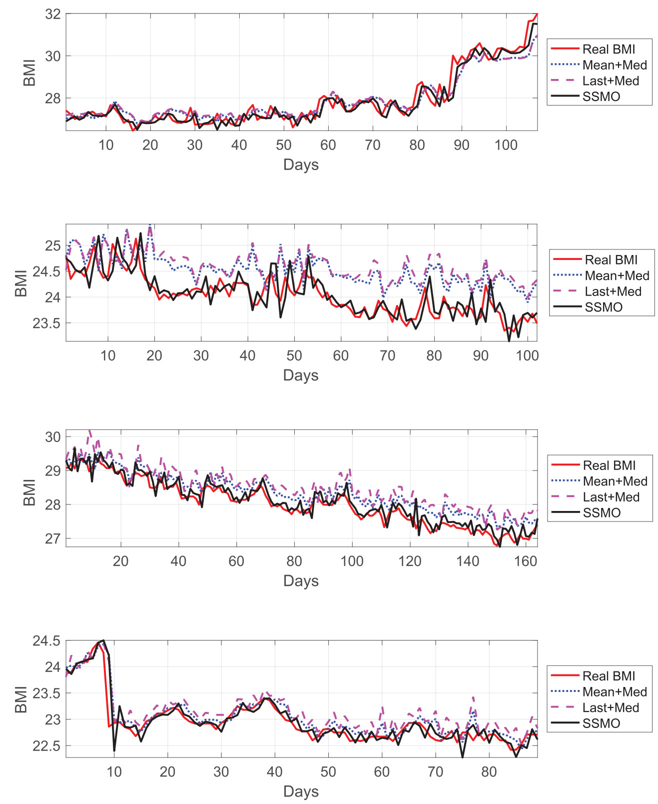

Figure 2. BMI estimation with three different methods.

4. Experiment

In experiment, our task is to learn personalized health planning for users to manage their BMI. We used proprietary data collected in a longitudinal study of obesity that involves more than 1000 real-world users and each user has several years' daily measurements (collected from wearable devices, including diet, sleep, exercise information and BMI). We evaluated the framework based on this dataset and also evaluated the effect of personalized recommendation using 25 users whose data are preprocessed and ready for analysis. For the 25 users, in total we have more than 10,000 days' measurements. The code will be released upon the acceptance of this paper with a full dictionary of the variables.

4.1. Evaluation of SSMO

We first evaluate the performance of SSMO. Baselines are benchmark methods for imputing missing values and removing outliers, including the "mean imputation," the "last value carried forward" method and "last value carried forward" with a median filter to remove outliers, as mentioned by Gelman and Hill (2007). In addition, we also compared it with the non-parametric Principal Analysis by Conditional Expectation (PACE) (Chen, 2015), MICE (van Buuren & Groothuis-Oudshoorn, 2011), Amelia (Honaker, King, & Blackwell, 2011) and MI (Gelman & Hill, 2015). Results in Table 1 demonstrate that SSMO consistently outperforms all benchmark methods. Note that, SSMO has additional benefits of simultaneous anomaly detection while these baselines do not.

We further evaluate the performance of SSMO when both missing values and outliers present in the data. Here, we analyze a real-life fitness data with users' daily fitness behaviors variables including diet, sleep, exercise information and BMI as health outcomes. There are 25 users' data in this dataset, while almost every user's data show significant missing values and outliers with similar patterns as shown in Figure 2. Again SSMO achieves more accurate model estimation, as reflected by the prediction errors in those users. Figure 2 shows the details of the prediction results by SSMO, and the other two imputation methods. For each user, the dynamic model has been built based on the data generated during the first half of days, and evaluated on the other half for prediction errors.

4.2. Evaluation of dictionary construction with APC

Dictionary construction restricts the recommended actions within a space spanned by some existing users' action data. There is an implicit assumption of this method that hypothesizes that, although people are different and have heterogeneous behavioral patterns, there are some regularities or canonical structures governing the human behavior. Therefore, the effectiveness of the dictionary construction method APC depends on how valid this assumption holds true in reality. Here, we apply APC on the 25 users' behavioral data. Figure 3 provides the results regarding how many basis vectors we need to represent all the behavioral data of all the users. Apparently, the larger the dictionary size is, the better (lower) the error of representation by the squared Frobenius norm that it can provide. On the other hand, we can also observe that the error of representation drops quickly with the increasing number of basis vectors in the

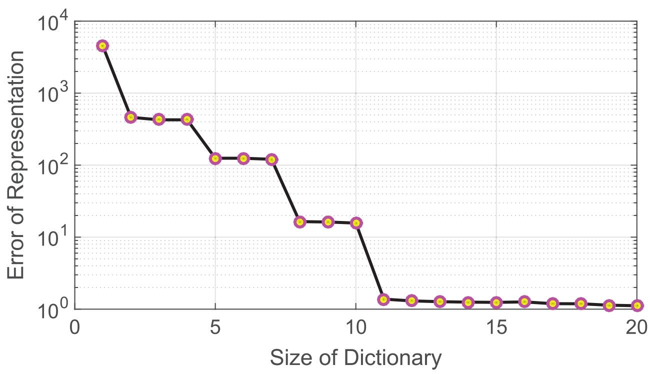

Figure 3. Errors of representation of **U** based on the size of dictionary.

Figure 4. Recommended Routines (in Eq. (4.2)) and Plans (in Eq. (4.4)).

dictionary. With eleven basis vectors, the error of representation approaches 1.0.

A visualization of the five basis vectors are shown in Figure 4, representing five typical health management routines used in the cohort and three levels of combinations of the routines. For example, Pattern 1 highlights the decent amount of calories consumed from food with less activities, while Pattern 3 is a more balanced diet and exercise routine. Interestingly, the learned patterns can be mapped to the official guideline for obesity prevention (Fitch et al., 2013).

4.3. Evaluation of the planning

While the ultimate evaluation for any healthcare planning strategy is to conduct clinical trials, it is expensive and often can only be realistic at the later stage of the health improvement plan development. On the other hand, the literature shows that compliance to health recommendations is such a complex behavior that the compliance levels vary from user to user and even for the same user. Therefore, we take a pragmatic approach to evaluate the efficacy of our strategy via data-based simulations with scenarios that reflect different user compliance levels. We first simulate the dynamic change of the health outcome. Specifically, we randomly pick up three users with their behavioral data (with last M measurements held out for evaluation) to train the dynamic model using SSMO and further derive the optimal planning as a temporal action set **U** using our planning formulation. We investigate a range of

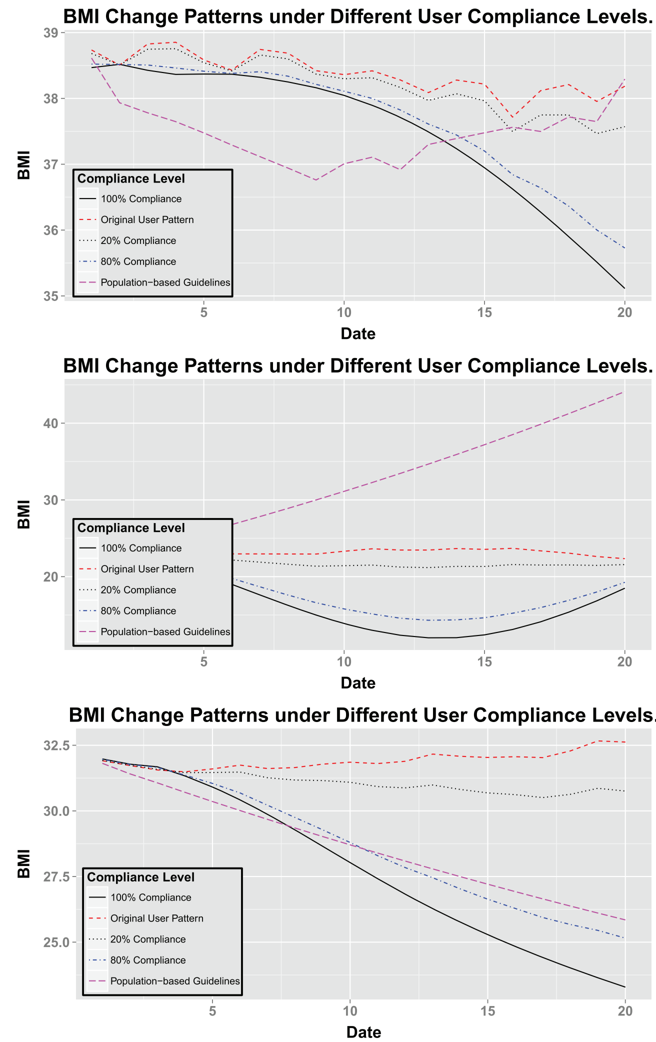

Figure 5. BMI change patterns under different user compliance levels.

compliance levels. For example, a 80% compliance level means 80% actions are randomly picked from the optimal plan, and the rest are from the originally observed behavior. We then predict the BMI change based on users' learned dynamic model.

The results in Figure 5 indicate that outcomes from users' original routines will either fluctuate and then become worse, or stay in a plateau or continue to become worse. While for all users, complete or partial compliance to the optimal plans always leads to better and steady results. Note that, we also compare our method with the population-based planning, i.e., recommending the user with the mean level of activities from the user's peers. The population-based planning might be reasonable for users with average conditions, but may not perform as well or even result in condition deterioration for the users with very high or low BMI profiles. Again, as shown in Figure 5, our method outperforms population-based planning by providing more effective personalized plans and recommendations.

4.4. Effect of the smoothness and dictionary constraints on recommendation

We also investigate the effectiveness of the smoothness constraint and the action polyhedron on the quality of the generated plannings. While the planning quality is a multi-facet concept, a simple criterion is that the derived plan should fit the existing behavioral patterns. Thus, our strategy here is to randomly pick up an user and generate three optimal plans: Plan A from the full model as depicted in Section 3.2; Plan B from a reduced model without the constraint (3e) and Plan C from a further reduced model by removing the constraint (3e) and setting λ in (3a) to 0. The three plans are drawn on top of the five typical patterns in Figure 4. It is obvious that the optimal plan derived from the full model fits the existing patterns better, which seems to be more realistic and has a higher likelihood to be adopted by the users than the other two plans that are quite different from the existing patterns. In addition, we quantitatively evaluated this conclusion by calculating the distance between the plans to the space defined by the basis vectors (as a score representing how similar the plans with real-world actions). We used cosine similarity that showed the scores for the Plan A, B, C as 0.901, 0.856 and 0.855, respectively.

5. Conclusions

In this work, we have developed a systematic framework that unifies dynamic modeling, sparse learning, dictionary learning and matrix completion, to automate the personalized health planning. Our method is generic and can be applied to a wide range of health management problems such as obesity, fitness or any chronic conditions, where the disease process is a complex dynamic process that can be modified by exogenous variables such as environmental, behavioral and clinical variables. Our framework holds great potential to provide scalable solutions for mitigating these health problems, which can promote healthier lifestyles outside of clinical settings.

References

Antsaklis, P. J., & Michel, A. N. (2007). A linear systems primer. Springer.

Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge, UK: Cambridge University Press.

Chen, K. H. (2015). *PACE package for functional data analysis and empirical dynamics*.

Cong, Y., Yuan, J., & Liu, J. (2011). *Sparse reconstruction cost for abnormal event detection*. CVPR '11 (pp. 3449–3456).

Consolvo, S., Klasnja, P., McDonald, D. W., & Landay, J. A. (2014). Designing for healthy lifestyles: Design considerations for mobile technologies to encourage consumer health and wellness. *Foundations and Trends in Human–Computer Interactions*, 6 (pp. 169–307).

CVX Research. (2012). CVX: Matlab software for disciplined convex programming, version 2.0. Retrieved from http://cvxr.com/cvx

Fitch, A., Fox, C., & Bauerly, K., Gross, A., Heim, C., Judge-Dietz, J., … Webb, B. (2013). *Guideline of prevention and management of obesity for children and adolescents*.

Fogg, B. J. (2009). *A behavior model for persuasive design*. Proceedings of the 4th International Conference on Persuasive Technology, Persuasive '09. Claremont, California, USA, 2009.

Fung, D., & Sheung, C. (2006). *Methods for the estimation of missing values in time series* (PhD thesis). Edith Cowan University Perth, Perth, Western Australia.

Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel and hierarchical models*. Cambridge, UK: Cambridge University Press.

Gelman, A., & Hill, J. (2015). *mi: Missing data imputation and model checking*.

Gittins, J., Glazebrook, K., & Weber, R. (2011). *Multi-armed bandit allocation indices* (2nd ed.). New Jersy: John Wiley & Sons.

Honaker, J., King, G., & Blackwell, M. (2011). Amelia II: A program for missing data. *Journal of Statistical Software*, *45*(7), 1–47.

Koren, Y. (2009). *Collaborative filtering with temporal dynamics*. KDD '09. Paris, France.

Kukafka, R. (2005). Tailored health communication. In D. Lewis, G. Eysenbach, R. Kukafka , P. Z. Stavri, H. Jimison (Eds.), *Consumer health informatics*: Informing *consumers and improving health care,* pp. 22–33. Springer, USA.

Liu, B., Li, J., Chen, C., Tan, W., Chen, Q., & Zhou, M. (2015). Efficient motif discovery for large-scale time series in healthcare. *IEEE Transactions on Industrial Informatics*, *11*(3), P583–590.

Mairal, J., Bach, F., Ponce, J., & Sapiro, G. (2009). *Online dictionary learning for sparse coding*. Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09. Montreal, Canada.

Qiao, X., Yu, W., Zhang, J., Tan, W., Su, J., Xu, W., & Chen, J. (2015). Recommending nearby strangers instantly based on similar check-in behaviors. *IEEE Transactions on Automation Science and Engineering*, *12*(3), P1114–1124.

Rabbi, M., Aung, M. H., Zhang, M., & Choudhury, T. (2015). *Mybehavior: Automatic personalized health feedback from user behaviors and preferences using smartphones*. Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '15. Osaka, Japan.

Spring, B., Gotsis, M., Paiva, A., & Spruijt-Metz, D. (2013). Healthy apps: Mobile devices for continuous monitoring and intervention. *IEEE Pulse*, *4*(6), 34–40.

Tseng, P. (2001). Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, *109*(3), 475–494.

van Buuren, S., & Groothuis-Oudshoorn, K. (2011). mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software*, *45*(3), 1–67.

Xiong, L., Chen, X., Huang, T., Schneider, J., & Carbonell, J. (2010). *Temporal collaborative filtering with bayesian probabilistic tensor factorization*. Proceedings of the SIAM International Conference on Data Mining, SDM 2010 (pp. 211–222), April 29–May 1, 2010, Columbus, OH.

Yeh, C.-C. M., & Yang, Y.-H. (2012). *Supervised dictionary learning for music genre classification*. ICMR '12. Hong Kong, China.

Zhang, Y., Meratnia, N., & Havinga, P. (2010). Outlier detection techniques for wireless sensor networks: A survey. *IEEE Communications Surveys & Tutorials*, *12*(2), 159–170.