

Activity Recommendation: Optimizing life in the long-term

Julian A. Ramos Rojas
Human-computer
Interaction Institute
Carnegie Mellon
University
Pittsburgh, PA
ingenia@andrew.cmu.edu

Johana Rosas
Psychology
Department
University of
Pittsburgh
Pittsburgh, PA
Jmr259@pitt.edu

Yilin Shen
Samsung Research
America
Mountain View, CA
Yilin.shen@samsung.com

Hongxia Jin
Samsung Research
America
Mountain View, CA
hongxia.jin@samsung.com

Anind K. Dey
Information School
University of
Washington
Seattle, WA
anind@uw.edu

Abstract— College students every day decide and plan how to best spend their time to balance academic, physical, and social goals under uncertainty. This process is likely suboptimal where long-term life satisfaction and success is not guaranteed, and poor decision-making may lead to longer-term problems like depression. To support everyday planning, we introduce activity recommendation, a novel method that combines artificial intelligence, machine learning, and a psychology-informed approach to automatically generate activity-recommendations that optimize long-term life satisfaction. We tested our method with an existing dataset and derived activity recommendations for depressed and non-depressed students. We evaluated the recommendations through interviews with college students who rated the suggestions positively. Our model can be optimized for different goals and domains and is easy to interpret. Our results demonstrate the feasibility of our approach and lay the groundwork towards implementing a live system.

Keywords— *Q-Learning; Reinforcement Learning.*

I. INTRODUCTION

College life is full of challenges. For many, college is the first experience of living as an independent adult and a time when new responsibilities are acquired. In addition to academic requirements, students have to adjust to a new environment and make decisions, such as what time to wake up on a given day, which major they want to study, *etc.* Having these many different sources of pressure disrupts student life and may lead them to take decisions that may not be the best for their physical or mental health in the long-term. Results from the National college health assessment (ACHA) IIc [1] showed that 53.7% of students reported experiencing levels of stress ranging from “more than average stress” to “tremendous stress”. Results also showed that 85.7% of students felt overwhelmed by all they have to do, 82% felt exhausted but not from physical activity, and 59.5% felt very lonely within the past 12 months [1]. In 2014, the National Institute of Mental Health estimated that 9.3% of adults between 18 and 25 years of age had at least one major depressive disorder in the past year [2]. Depression is a mood disorder characterized by low mood, lack of interest in pleasurable activities, sleep problems, and difficulties concentrating and making decisions [3]. Balancing all responsibilities, social life, family relationships, and hobbies requires a lot of planning and having difficulties managing all of these activities could lead to feeling stress and have a negative impact on long-term mental health.

In this article, we present a novel method to automatically build an activity recommendation model that can help college students to better plan their everyday life in a way that will lead to long-term life satisfaction and health. Our method leverages state of the art machine learning (ML) and artificial intelligence (AI) techniques to suggest activities, based on the current mental and physical state of the student, and maximizes positive affect in the long term. Our method was created with interpretability in mind, which is crucial for understanding the recommendations and for offline evaluations (*e.g.*, simulations of the system before deployment to test for expected effects of the recommendations) that can better support field deployment of this technology. Our contributions can be summarized as follows:

- An artificial intelligence, machine learning, and psychology-informed approach to recommend activities to students that optimize their long-term positive affect.
- An analysis of our approach and recommendation results from both an algorithmic and psychological perspective.
- A user study validating our recommendation results with college students that found that students positively rate the recommendations and would follow them in most situations.

II. COLLEGE LIFE AND LONG-TERM REWARD ESTIMATION

Estimating the consequences (long-term rewards) of daily life decisions is not an easy task as they are greatly influenced by basic emotions. Studies by Tanaka [4] and McClure [5] demonstrate that decisions over delayed rewards are processed by two different systems in the brain. Short-term rewards are processed by the limbic system, which is concerned with instinct, mood and basic emotions. Long-term rewards are processed by the lateral prefrontal cortex and posterior parietal cortex. This may explain why people with major depressive disorder [6] and people undergoing drug cravings [7] make worse long-term financial decisions than control subjects. Moreover, the famous marshmallow experiments [8] and follow up studies [9]–[12] have demonstrated the importance of delayed gratification. In those studies, preschool children (4 years old) were given the option of getting an immediate reward vs. a delayed reward. In general, follow-up studies 10 years later showed that adolescents who chose the delayed reward as a preschooler, were rated by their parents as having higher ability

to concentrate and pay attention, and more likely to exhibit self-control in frustrating situations [9]. The ability to recognize and choose between short- and long-term rewards is fundamental when making life decisions and could influence an individual’s life satisfaction. Thus, it is critical and desirable to support college students’ daily planning with activity recommendations that can enable their long-term success, and physical and mental health.

III. RELATED WORK

Activity recommendation is a nascent field in the computer science, pervasive computing and human computer-interaction areas. We now review the most representative work in this domain. The first work that looked at behavior change and activity recommendation is MyBehavior [13], an application for generating food and physical activity recommendations. MyBehavior uses a Multi-Armed Bandit (MAB) algorithm as the core activity recommendation engine. It also uses behavior change models like the BJ-Fogg Behavior Model [14], and the Transtheoretical Model of Behavior Change [15] to guide design decisions of their app and research study, and to tune the recommendations generated through their MAB. MyBehavior supports achieving a health goal by suggesting activities to a user that she is likely to perform and will move her towards the goal. Making recommendations in this way reduces the effort a user has to make towards a goal and increases the success of a recommendation, as predicted by the BJ-Fogg model. MyBehavior is strongly focused on personalized models, which is a very interesting and well supported approach; *however*, it does not take advantage of data from all of its users to generate recommendations with even better long-term outcomes for each user. Also, the reliance on MABs for generating recommendations limits their outcomes to short-term optimization. Formally, MABs are one-state Markov Decision Processes and this limits long-term reward estimation. Potentially, an approach that could be used for activity recommendation is routine behavior modeling [16], [17]; however, this will only work in well-known or relatively simple environments due to the reliance of this method on a multi-state Markov Decision Process. Another difficulty with this approach is that it does not take into account long-term reward maximization, which is key for our goal of supporting college students in planning their everyday activities.

IV. NOTATION AND PROBLEM FORMULATION

As described above, there is no work focused on optimizing for long-term rewards. To estimate the effects of activities in the long-term, we formulate activity recommendation as a reinforcement learning (RL) problem: a set of problems and methods designed to deal with delayed or long-term effects of actions. Other methods exist for the forecast of delayed effects like regressions with lag variables or the Autoregressive Integrated Moving Average (ARIMA), but unlike reinforcement learning methods, they cannot model the interaction between actions and states, where state is a set of variables representative of the evolution of the problem and actions are the decisions taken to optimize a desired outcome. Modeling both the state

and actions are a necessary condition for modeling an activity recommendation problem. In this section, we introduce the notation and definitions fundamental for reinforcement learning problems in the context of a college student’s life. We start with a core definition of RL and then describe Q-Learning, the specific RL algorithm we use in this work.

A. The problem

We define our problem as one of balancing college-life, by recommending activities that will maximize students’ long-term wellbeing. We assume that these activity recommendations are provided each night when the student is planning how many hours to sleep that night, and how to spend her time the next day. A complete model of student behavior would consider the likelihood of a student following each potential recommendation, however this would require a field evaluation of at least an entire semester. Instead, we use an offline evaluation (*i.e.*, simulation) of our approach and assume that students follow suggestions exactly as proposed by the system.

B. Reinforcement Learning

Reinforcement Learning (RL) [18] is defined both as a framework and a set of decision-making problems where an agent in a given environment is trying to maximize a reward through a specific set of actions that have an effect on the environment. More formally, RL models consists of:

- A set of environment states $s \in \mathcal{S}$;
- A set of agent actions $a \in \mathcal{A}$;
- A reward r obtained at some state s after taking action a .

In this paper, the agent is a college student; the actions are the different activities the student is planning to do next such as, how many hours to sleep tonight, how long to exercise the next day, and how much to socialize the next day. The environment is a simplified representation of a student’s daily life, including hours slept, socialization level, exercise level and deadlines, any combination of which produces a state. A state represents a summary of a student’s day. As an example, socialize with less than 5 people, slept 6 hours or more, and 0 deadlines could be labeled as a lazy day state. The reward is a general measure of wellness, a numerical value that represents how good or bad the student is feeling on a specific day. There are many different goals that can be achieved with RL, but in this work we focus on calculating the value function V and a policy π . A value function estimates the long-term reward¹ for a given state $V(s)$ or state-action pair $Q(s, a)$, with the latter known as the Q-values function. The policy is a function that defines the actions that are going to be suggested to the student each day to maximize long-term reward. The policy then is the suggested student’s behavior in his environment. For a more thorough review of reinforcement learning please see [18]–[20]. There are many different variations and assumptions that change the methods in a RL problem; here we focus on Q-Learning.

¹ In the RL literature this is referred to as the expected sum of discounted rewards.

C. Q-Learning and delayed reward

Q-Learning [21], [22] is a temporal difference learning method used for estimating an optimal policy that maximizes long-term reward. It also belongs to a family of methods called model free, which means that the environment is not modeled explicitly. Reasons for not explicitly modeling the environment vary, but for the purpose of this work, the environment is too complicated to be captured analytically [23]. Also, the quantity of data available for building our model was not large enough to estimate a reliable model of the environment. In comparison with MABs, which were used in MyBehavior system, MABs can capture local minima more easily than Q-Learning. However, Q-learning does consider temporal correlation between states that is not covered by MABs and this is a desirable property for our problem given that naturally occurring everyday life activities are correlated with previous activities and previous days (e.g., sleepiness depends among other factors on the amount of sleep from the previous night). Also, Q-Learning can produce the optimal policy even without converging to the right values; this is also desirable since the convergence of Q-Learning to the optimal values cannot be measured. Another desirable property in the context of the college life-balance problem is that Q-Learning can find an optimal policy independent of the policy used to obtain the data. In the context of this problem, the policy used to obtain the data is learned from each student decision on which activities to do and for how long. Q-Learning is defined by:

- A learning factor α that adjusts how much a previous estimate of the Q-value function will change when combined with more recently observed data
- A discount factor γ which weights the influence of future rewards based on previously observed Q-values
- A reward r
- A future state $s' \in \mathcal{S}$
- A future action $a' \in \mathcal{A}$

The main components of Q-Learning are the Q-values estimate defined as:

$$\hat{Q}(s, a) \leftarrow (1 - \alpha) \cdot \hat{Q}(s, a) + \alpha \left(r + \gamma \cdot \max_{a'} \left(\hat{Q}(s', a') \right) \right)$$

Equation 1. Q-Values estimation

Intuitively, the above equation adjusts the long-term or delayed rewards for a given state, action and future state s' by weighting the previous Q-value estimate, the reward received and the best possible long-term reward obtained in the future state. Equation (1) can be tuned for seeking any delayed reward desired. For instance, to seek short-term rewards exclusively, we can set $\gamma = 0$, while for estimating the history of all rewards it would be $\gamma = 1$. Then the behavior of a student who is not very good at long-term planning will be very similar to that calculated using the Q-values for a low gamma. The gamma value can then be tuned to generate activity recommendations that can change smoothly over time from short- to long-term reward. This could prove beneficial for behavior change since by suggesting activities that are similar to what the student is already doing and slowly changing them to more optimal ones, this decreases the effort the student has to make to adopt those changes. This, according to the BJ-Fogg model of behavior [14], would

generate a higher likelihood for following the activity suggestion even when the student has very low motivation to do so.

V. ACTIVITY RECOMMENDATION MODEL

To support student planning and life satisfaction in college we now introduce our activity recommendation model. This model suggests activities to perform taking into account the current state of the student. Our activity recommendation model seeks to maximize long-term satisfaction of a student although it could be used to improve any other aspect of a student's life and also could be used with a different population. To evaluate the feasibility of our activity recommendation model we used the StudentLife [24] dataset. Details on this dataset, measures used, data imputation and preprocessing are provided in the StudentLife dataset section. Our method can be summarized in the next steps:

- Choose a reward variable
- Separate data into groups that may experience short-term reward in a different way.
- Compute features that convey the most information related to the long-term reward.
- Cluster the features into a reasonable number of clusters that allows for good compression of the data and can still be interpreted by a human in a reasonable amount of time.
- Discretize the actions.
- Compute the value function and policy using Q-Learning.

Next the reasoning behind each step is explained.

A. Choosing a reward variable

In our activity recommendation model, we are interested in maximizing the long-term positive affect of a student. Ideally, we would have liked to have a multivariate reward function that could account for the student's physical and mental health, and academic performance; however, academic performance was not available at a high enough granularity to be used for this purpose. We used the Photographic Affect Meter (PAM) administered on a mobile EMA which was the only measure available in the StudentLife dataset that indirectly estimates the physical and mental health of a student. The PAM measures affect, both positive and negative. There are consistent patterns of the association of positive affect with physical health in the literature. Positive affect has been linked to: lower morbidity, decreased symptoms and pain, and increased longevity among older community-dwelling individuals [25]. The numerical values of the PAM vary from 1 to 16, where 16 represents the highest arousal and positive valence while 1 represents the lowest arousal and negative valence. Since the numerical values of the PAM are directly proportional to affect, they were used directly as our reward, and then our method maximized the long-term positive affect. For all the data used in this work, students had at least one PAM measure per day, in cases were more than one PAM was logged the average was calculated and used instead.

B. Separating into groups

To account for the differences in how people experience positive and negative affect, we divided the StudentLife dataset into depressed and non-depressed students using the PHQ9 [26]. The PHQ9 is a self-administered questionnaire that includes the

nine affective, behavioral, cognitive, and somatic symptoms of depression as described by the DSM-IV [27]. Depression affects mood and cognition; thus, what generates positive emotions in a non-depressed student may not have the same effect or could even have an opposite effect in a depressed student. Personalized activity recognition models were not considered due to the small size of the dataset per student (46 days on average). Instead, by grouping data from different students, we assume that their PAM reports are similar under similar conditions (similar state and action). By making this assumption, we switch from having unique days for each participant to having unique states (*i.e.*, types of day) and by doing so we can now observe transitions across different states enough times to make Q-Learning applicable to this kind of problem. Note that other factors like personality traits could also have been used to divide the population into groups.

C. Computing features

In our activity recognition model, we assume that each night, a student is planning what to do tomorrow while providing our model with his current state. Our activity recommendation model then can suggest activities to the student each night. These suggestions are comprised of how many hours to sleep tonight, and how much to socialize and exercise tomorrow. With this in mind, to create an activity recommendation, a good state representation is needed. This state representation can only be accomplished by collecting enough information such that the algorithm can differentiate between similar states that could have potentially different long-term rewards. We start by identifying the factors that have an influence on a student’s positive affect and the information that we have access to through the StudentLife dataset. We were restricted by the dataset we used, however this does not mean that more information cannot be included in future datasets. Another dimension of importance when searching for a correct state representation is the time span of a state. Short time spans could produce frequent yet unnecessary suggestions while long time spans may not be frequent enough to be useful or have an effect on the desired variable that we are seeking to maximize. In this work we used the time span of one day. The variables used in our state representation include:

Physical activity: was measured using participants’ self-report of the time they spent exercising. This variable summarizes the amount of time the student walked and/or exercised during the current day. Although this variable does not capture health itself, it does capture promotion of physical health. Note that excessive exercise or exercise at times when there is a high number of deadlines could also adversely affect the student, so this variable can capture both positive and negative effects.

Affect: was measured with the Photographic Affect Meter (PAM) [28]. This variable captures both mental and physical health of the student. Through the PAM, our model can learn the level of happiness (valence) and energy (arousal) the student is experiencing.

Deadlines: was measured counting the number of self-reported daily number of homework assignments, projects and other deadlines the student had each day of the study. This variable captures some aspects of academic pressure.

Social: refers to the level of socialization the student experienced the current day, as reported through a daily survey.

Sleep hours: number of hours slept as reported through a daily survey. Sleep is fundamental for both mental and physical health. For example, many biological processes in the human body depend on sleep.

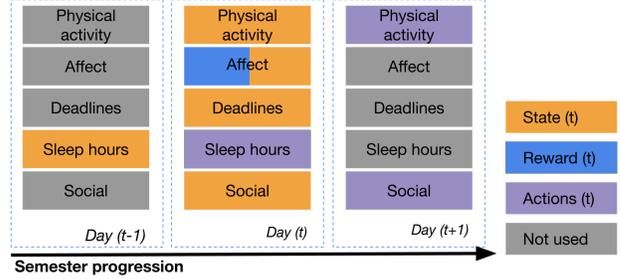


Figure 1. State representation. The color for each variable represents whether it is part of the state, reward or action. Affect is part of both the state and the reward.

The state representation for any day t is shown in Figure 1. The state for a day t is represented by the physical activity, social interaction, deadlines and affect for that day (t) and the number of hours slept the day before ($t - 1$). The state representation used captures the idea that the number of hours the student slept the previous night has a strong effect on the current state and more importantly in the reward function (*i.e.*, student’s mood), while other activities like physical activity and affect have the strongest effect the same day. The Actions, which are the suggestions our model is going to generate, are the recommended sleep duration that night (t) and the recommended amount of physical exercise and socialization for the next day ($t + 1$).

D. Clustering

Although the features stated above capture many aspects of importance in the daily life of a student, its interpretation is difficult due to a large state space that grows exponentially with the number of variables. Also, this state representation cannot be used to train a Q-Learning model, due to a Q-Learning property that guarantees convergence when actions are repeatedly observed in all states and the action-values are represented as discrete values [21]. The first condition is not met as there are almost as many states as days in our dataset. To overcome this challenge, instead of using continuous values for the state representation, the dataset was clustered, with each state being represented by a centroid of one of the clusters discovered. Clustering refers to a set of unsupervised machine learning methods used to find the natural groupings of a dataset. However, clustering methods can also be used to compress a dataset into a low dimensional space. To compress our dataset and meet the requirements for Q-Learning to converge, we clustered our data using k-means with 22 clusters. The number of clusters was determined using a variation of the Bayesian information criterion modified for k-means [29]. The number of clusters for both the depressed and non-depressed dataset was picked such that further improvement from increasing the number of clusters was negligible and the number of clusters was low enough to support interpretation of the results. Due to the random initialization of k-means, clusters found with

after finding problems with the way the EMA was asked, which may have resulted in unreliable responses. A student’s data was excluded from analysis if any of the features from their data collected was completely missing. Our final dataset is composed of 44 students and 2024 days of data for an average of 46 days logged of data per person.

A. Data imputation

Although the StudentLife dataset is a great source of information, there was a high percentage of missing data. Considering only the variables of interest, we have a total of 70% missing data. Among all variables, PAM has a very small number of missing data points.

To calculate the values of the missing data points we used an imputation technique, testing it first with our dataset. In this data imputation test, we first took rows of our dataset without missing data points, and then randomly deleted data points (generating fake missing data points). We imputed this fake dataset using several data imputation techniques from a data imputation library [31]. We then measured the root mean squared error of the recovered dataset against the original dataset. We found that the best technique for imputing with our particular dataset was Soft-Impute [32].

B. Preprocessing

After imputing missing values, we computed the average value for each variable of interest. For example, PAM in many cases was reported more than once per day. In most cases, all values (or a subset of them) were directly averaged with the exception of exercise related values. For exercise, there were multiple questions, but we only took into account the answers to: *How long did you walk for today? and If you exercised how long did you exercise for?* The possible answers to these questions were time ranges that varied from none to greater than 90 minutes. We multiplied the time spent exercising by two and added this value to the time spent walking and we called this value Physical Activity. The multiplication factor is based on the Physical Activity Guidelines for Americans [33], where it is stated that the time spent in vigorous activities counts for double the time spent in moderate activities. For sleep, we only used the number of hours slept the previous night. In order to discretize both activities and state features, we used percentiles. Also, we divided our dataset into depressed and non-depressed students using the average of the PHQ-9 [26] questionnaires they filled at the beginning and the end of the study.

VII. POLICY OBTAINED

As described earlier, subjects in our dataset were split into non-depressed and depressed student groups. For each group of students, a policy was generated from the Q-values calculated. The Q-values for the non-depressed students can be observed in Figure 2 (left). The labels generated for the non-depressed students are shown in Table 3. The Q-values calculated for the depressed students are shown in Figure 2 (right) and the labels for the states are shown in Table 4.

S	Label	Recommendation
0	Distressed, slept less than 4.5 hours, socialized with less than 9 people, exercised less than 20 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more

1	Contented, slept less than 4.5 hours, socialized with 10 to 19 people, exercised 20-30 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 19 people or more, Exercise 5 minutes or less
2	Distressed, slept 4.5 to 6 hours, socialized with 19 people or more, exercised 20-30 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
3	Excited, slept 6 hours or more, socialized with 19 people or more, exercised 20-30 minutes and 1 deadline	Sleep more than 6 hours, Socialize with 19 people or more, Exercise 20-30 minutes
4	Contented, slept 6 hours or more, socialized with 10 to 19 people, exercised 20-30 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
5	Contented, slept 6 hours or more, socialized with 10 to 19 people, exercised 30 minutes or more and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
6	Contented, slept 4.5 to 6 hours, socialized with 10 to 19 people, exercised 20-30 minutes and 2 deadlines	Sleep more than 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
7	Distressed, slept 4.5 to 6 hours, socialized with 10 to 19 people, exercised 20-30 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
8	Excited, slept 6 hours or more, socialized with 10 to 19 people, exercised 20-30 minutes and 1 deadline	Sleep more than 6 hours, Socialize with less than 9 people, Exercise 20-30 minutes
9	Depressed, slept 4.5 hours or less, socialized with less than 9 people, exercised less than 20 minutes and 0 deadline	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
10	Distressed, slept 4.5 to 6 hours, socialized with 10 to 19 people, exercised 20-30 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
11	Contented, slept 6 hours or more, socialized with 10 to 19 people, exercised 20-30 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
12	Distressed, slept 6 hours or more, socialized with 10 to 19 people, exercised 20-30 minutes and 0 deadlines	Sleep less than 4.5 hours, Socialize with 19 people or more, Exercise 20-30 minutes
13	Contented, slept 4.5 to 6 hours, socialized with 19 people or more, exercised 20-30 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
14	Contented, slept 6 hours or more, socialized with 10 to 19 people, exercised 20-30 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
15	Contented, slept 4.5 to 6 hours, socialized with 10 to 19 people, exercised 20-30 minutes and 0 deadlines	Sleep more than 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
16	Depressed, slept 6 hours or more, socialized with less than 9 people, exercised less than 20 minutes and 0 deadlines	Sleep more than 6 hours, Socialize with less than 9 people, Exercise 20-30 minutes
17	Distressed, slept 4.5 to 6 hours, socialized with 10 to 19 people, exercised 20-30 minutes and 2 deadlines	Sleep more than 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more

18	Distressed, slept 4.5 to 6 hours, socialized with 5 to 10 people, exercised 20-30 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 30 minutes or more
19	Distressed, slept 4.5 hours or less, socialized with less than 9 people, exercised less than 20 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 10 to 19 people, Exercise 20-30 minutes
20	Distressed, slept 4.5 to 6 hours, socialized with 10 to 19 people, exercised 20-30 minutes and 1 deadline	Sleep more than 6 hours, Socialize with 10 to 19 people, Exercise 20-30 minutes
21	Distressed, slept 4.5 to 6 hours, socialized with 10 to 19 people, exercised 20-30 minutes and 1 deadline	Sleep less than 4.5 hours, Socialize with 19 people or more, Exercise 20-30 minutes

Table 3. Non-depressed students' state labels. S denotes the state

S	Label	Recommendation
0	Distressed, slept 6 hours or more, socialized with 5 to 19 people, exercised 5-12 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
1	Contented, slept 4.5 to 6 hours, socialized with 5 to 19 people, exercised 5-12 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
2	Distressed, slept 4.5 hours or less, socialized with 5 or less people, exercised 5 minutes or less and 0 deadlines	Sleep more than 6 hours, Socialize with less than 5 people, Exercise 5-12 minutes
3	Excited, slept 4.5 to 6 hours, socialized with 19 people or more, exercised 5-12 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
4	Excited, slept 6 hours or more, socialized with 5 to 19 people, exercised 5-12 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
5	Excited, slept 4.5 to 6 hours, socialized with 5 to 19 people, exercised 5-12 minutes and 0 deadlines	Sleep more than 6 hours, Socialize with less than 5 people, Exercise 12 minutes or more
6	Distressed, slept 4.5 hours or less, socialized with less than 5 people, exercised less than 5 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
7	Contented, slept 4.5 to 6 hours, socialized with 5 to 19 people, exercised 5-12 minutes and 0 deadlines	Sleep more than 6 hours, Socialize with 19 people or more, Exercise 5-12 minutes
8	Contented, slept 6 hours or more, socialized 19 people or more, exercised 5-12 minutes and 0 deadlines	Sleep less than 4.5 hours, Socialize with less than 5 people, Exercise 5-12 minutes
9	Distressed, slept 6 hours or more, socialized with less than 5 people, exercised less than 5 minutes and 0 deadlines	Sleep less than 4.5 hours, Socialize with 19 people or more, Exercise 5-12 minutes
10	Contented, slept 4.5 to 6 hours, socialized with 5 to 19 people, exercised 5-12 minutes and 1 deadline	Sleep more than 6 hours, Socialize with 19 people or more, Exercise 5-12 minutes
11	Depressed, slept 4.5 hours or less, socialized with 5 people or less, exercised less than 5 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 19 people or more, Exercise 12 minutes or more

12	Distressed, slept 4.5 to 6 hours, socialized with 19 people or more, exercised 5-12 minutes and 0 deadlines.	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 5-12 minutes
13	Distressed, slept 4.5 to 6 hours, socialized with 5 to 19 people, exercised 5-12 minutes and 0 deadlines	Sleep more than 6 hours, Socialize with 5 to 19 people, Exercise 5 minutes or less
14	Contented, slept 6 hours or more, socialized with 5 to 19 people, exercised 5-12 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
15	Contented, slept 4.5 to 6 hours, socialized with 5 to 19 people, exercised 5-12 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 19 people or more, Exercise 12 minutes or more
16	Excited, slept less than 4.5 hours, socialized with 5 to 19 people, exercised 5-12 minutes and 0 deadlines.	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
17	Contented, slept 4.5 to 6 hours, socialized with 19 people or more, exercised 5-12 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 19 people or more, Exercise 12 minutes or more
18	Depressed, slept 6 hours or more, socialized with less than 5 people, exercised less than 5 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with less than 5 people, Exercise 5 minutes or less
19	Excited, slept 4.5 to 6 hours, socialized with less than 5 people, exercised 5-12 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
20	Distressed, slept 6 hours or more, socialized with 5 to 19 people, exercised 5-12 minutes and 0 deadlines	Sleep 4.5 to 6 hours, Socialize with 5 to 19 people, Exercise 12 minutes or more
21	Distressed, slept 6 hours or more, socialized with 5 to 19 people, exercised 5-12 minutes and 1 deadline	Sleep 4.5 to 6 hours, Socialize with 19 people or more, Exercise 5 minutes or less

Table 4. Depressed students' state labels. S denotes the state

VIII. POLICY EVALUATION

In order to evaluate the policy obtained we performed an interview in which we asked college students to imagine themselves feeling a certain way after experiencing a described day, and then asked them to rate activity recommendations suggested for them by our model according to their depression level.

Participants: A total of 14 undergraduate students participated in the study (36% women). Participants' age ranged between 17 and 23 years of age ($M = 18.86$ years; $SD = 1.79$). The sample was composed of 3 students that identified themselves as Caucasian (22%), 2 (14%) Hispanic; 7 (50%) Asian; 1 (7%) Black/African American; and 1 (7%) Native Hawaiian or Other Pacific Islander. Ten students were in their Freshman year (71%), 1 Sophomore (7%), 2 Junior (14%), and 1 Senior (7%) at a private University. The students were from many different majors: Economics, Computer Science, Information Security, Electrical engineering, Physics, Design, Civil engineering, Mechanical Engineering and Business administration. The study took about 25 minutes to complete. Participants were compensated \$5 for their time.

Measures and Procedure: Participants first filled out a demographics questionnaire including age, gender, race, and grade, and the PHQ-8 depression questionnaire [34]. Next, researchers scored the PHQ-8 to determine if the student was depressed or not in order to choose the appropriate set of recommendations. After that, students were asked by the researchers to imagine being in a certain state (see below) and were given a recommendation to improve long-term reward as estimated with our method. Students were then asked how much they agree with the activity suggestion using a Likert scale from 0 to 5 with 0 being “Strongly Disagree”, 3 “Neutral” and 5 “Strongly Agree”. Next, the students were asked whether they would follow this recommendation. The possible answers were “Yes”, “Maybe” or “No”.

Here is an example: “Imagine that you are feeling distressed, you slept less than 4.5 hours last night, and you socialized with less than 9 people, exercised for less than 20 minutes, or walked for less than 40 minutes and have 0 deadlines or homework assignments for today. **The Recommendation for you is:** To sleep 4.5 to 5.9 hours tonight, and to socialize with 10 to 19 people, and exercise for 20 to 30 minutes or walk for 40 to 60 minutes tomorrow. How much do you agree with this recommendation? Would you follow this recommendation?” Terms like socialization and exercise were defined to the student before starting the evaluation of the recommendation. Both socialization and exercise were defined as in the StudentLife dataset: “Socialization refers to people you had contact with, including anyone you said hello to, chatted, talked or discussed matters with, whether you did it face-to-face, by telephone, by mail or on the internet, and whether you personally knew the person or not”; “Exercise refers to higher than moderate physical activity and excludes walking”. The assessment of the recommendations was repeated for each of the 22 different states to get feedback on how the participants would feel about the different recommendations for each state (type of day). Lastly, we asked some follow up questions to understand students reasoning for agreeing or disagreeing with the recommendations. No identifiable information was collected in order to maintain anonymity.

IX. RESULTS

Half of the students in our study reported that they “Sometimes” plan their day the day before, 35% “Often” and 14% “Rarely”. Scores for the PHQ-8 ranged from 0 to 13 ($M = 4.43$, $SD = 3.88$). Only 1 student reported “Moderate Depression”, so we excluded this student’s ratings of the recommendations since a single student would not give us a representative opinion on how moderately depressed students feel about recommendations generated specifically for them.

To the question “How much do you agree with the recommendation?” the median rate was 4 (1 Strongly disagree, 3 Neutral, 5 Strongly Agree) the distribution is shown in Figure 4. To further analyze the responses, we divided them into positive (Strongly Agree and Agree), neutral, and negative (Strongly Disagree and Disagree). For each of the 22 situations, a median of 7 (53%) people gave positive responses, a median of 2 (15%) people gave neutral responses and a median of 4 (23%) people gave negative responses across all recommendations, as shown in Figure 4. To the question

“Would you follow this recommendation?” The students answered, “yes” 50% of the time, “maybe” 17% and “no” 31%. Of the 22 recommendations designed by the system for non-depressed students, two suggested to sleep for less than 4.5 hours. The following recommendation “To sleep less than 4.5 hours tonight, and to socialize with more than 19 people, and exercise for 20 to 30 minutes or walk for 40 to 60 minutes tomorrow” was given for states 12 and 21. For state 12, only 3 people “Agree” or “Strongly Agree” with the suggestion and 6 people noted that they would not follow it. For state 21, only 2 people “Agree” or “Strongly Agree” with the suggestion and 9 people noted that they would not follow it. There was only one suggestion to which nobody disagreed. For state 15 the suggestion was “To sleep 5.9 or more hours tonight, and to socialize with 10 to 19 people, and exercise for more than 30 minutes or walk for more than 60 minutes.” Eleven people said they would follow the suggestion and 2 noted that they would maybe follow it.

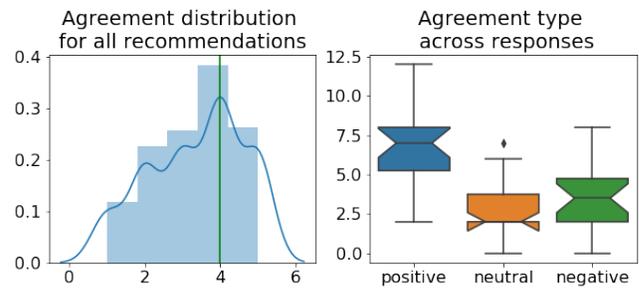


Figure 4. Left) Agreement distribution of all recommendations for all participants, the median is the vertical line. Right) Agreement by type across all responses, i.e., for the positive agreement, a median of 7 people agreed or strongly agreed to the recommendation provided.

X. DISCUSSION

Students had a mostly positive response to the recommendations suggested by our method as shown by the average agreement rate. Also, students reported that they would follow the recommendations at a higher rate than they would not. After further questioning about the reasons why they would not follow a recommendation, some of the students reported that they did not like suggestions to sleep for a short amount of time (i.e., 4.5 or less) or those with a limited amount of socialization (i.e., socialize with fewer than 9 people). Our participants said that they would follow recommendations that had what they perceived as a good balance for the state described. For “maybe” responses our participants said they would do more of some of the suggested activities, for example they would sleep more or socialize more given the opportunity. Like in other systems where a black box approach is taken, our study reveals that for our activity recommendation system to work in a live deployment, suggesting activities will not be enough and providing some explanation to support the recommendation will be necessary. In cases where the student does not agree with the optimal recommendation of the system, it may only mean that the student is not ready for a drastic change in her everyday life. In such cases, suboptimal (but better than a student’s own choices) recommendations have a higher likelihood of being

followed by the student as predicted by the BJ-Fogg behavior model [14] and self-efficacy literature [35].

A. *Sacrificing in the short-term for greater long-term satisfaction*

One of the advantages of our activity recommendation model is that it estimates long-term outcomes, however students may not always agree with such recommendations, as students are often unable to correctly anticipate future outcomes. Our model provides an automatic way to consider past and present states to predict and improve future states. In our interviews for example, students mostly disagreed with two recommendations for which the system is recommending to sleep less than 4.5 hours. At first sight, these recommendations seem unreasonable especially in the short-term. However, given the nature of our method and how recommendations are estimated, those recommendations are the best option for optimizing positive affect over the long term, as estimated from the data. This recommendation to sleep less make some sense, for example, when in an intense academic program that demands students to have long days filled with homework and projects and little sleep. The algorithm in this case is confounding causality with correlation since students may be sleeping less because they are busy and hence the recommendation should instead focus on making sure that homework and projects are up to date even if it results in a less sleep. This will ultimately result in higher long-term positive affect, as estimated from data.

This brings up the issue of having an automated method that is generating recommendations that do not line up with suggestions from medical institutions or even the goals of the user. As an example, our model suggests for two different states to sleep less than 4.5 hours as the best action. However, the American College Health Association (ACHA) suggests young adults sleep more than 6 hours every day. What this suggestion highlights is the need for having a supervisory system that checks the recommendations and compares them to guidelines and picks suggestions that meet official health standards. Looking at state 21 in Figure 2, the second best recommendation for that state suggests Sleep 4.5 to 6 hours, Socialize with 19 people or moreization and Low Exercise; this recommendation is in line with the ACHA recommendations and should be suggested instead. It remains as an open question as to why the system is suggesting what appears to be an unhealthy recommendation, despite the fact that it is generated from the data. We can observe that in both state 12 and 21, where low sleep suggestions are made, the student is distressed. This could mean that despite not having a current deadline there might be upcoming deadlines of importance that were not captured directly by our state representation but are having a significant effect in the student's affect and are captured through their distress.

B. *Shortcomings of our method*

As highlighted by the controversial recommendations to sleep less than 4.5 hours, the proposed method suffers from two main problems: 1) In a few cases, it confuses correlation with causation and 2) Optimizing over a single measure like affect could result in negative measures for other important outcomes: for example, students with a chronic condition like diabetes or asthma may need more sleep and physical activity and hence

they both need to be prioritized ahead of positive affect. To address the first issue, more fine-grained data should be collected to differentiate more clearly between actions and resulting outcomes from those actions. However, even having this fine grained data may not be a complete solution. For example, sleep duration can be an exogenous process affected by external constraints like school schedule, deadlines, *etc.*, or an endogenous process when for example a student decides to sleep late during the weekend. Detecting and modeling this duality of sleep as an exogenous or endogenous process requires further advances in this domain. To address the second issue, the goal of the user should be aligned with the goal in our method. In our specific study, the student feedback may have been different if each student's goal was indeed to maximize their positive affect in the long-term. In our study the students' actual goals are unknown.

XI. CONCLUSION AND FUTURE WORK

In this work, we have presented a method for automatically computing activity recommendations that maximize long-term positive affect for non-depressed and depressed college students. To our knowledge, this is the first attempt to overcome the gap between recommendations of current short-term suboptimal and long-term optimal lifestyles for college students. Our method takes into account a variety of factors, such as how the student's current day went (physical, social and academic aspects), how many hours she slept the previous night and her depression level. We evaluated our activity recommendation by interviewing and asking college students to evaluate the recommendations generated by our model. The results showed that overall, students positively rate the activities suggested and they would follow most of them.

Our recommendations, however, are as good as the dataset we used to generate them from. The StudentLife dataset is very comprehensive and interesting, however it lacks some of the contextual information we would need to better understand the activity recommendations. Another limiting factor was the quality of the data itself and more specifically the high amount of missing data. Although, we were able to impute missing data in the dataset, this is a factor that limits how much can be learned from the dataset.

In future work, we plan to collect our own data and incorporate more contextual information into our approach, in order to better understand the activity recommendations generated. For example, interesting sources of information we would have liked to have access to, include entertainment related information, complexity of homework and projects, quality of social interactions (rather than just quantity) and health status. This will be very useful not only for interpretability, but also for generating richer activity recommendations. Although, we addressed personalization by separating students into depressed and non-depressed students, other ways to personalize further could generate even better recommendations that take into account personality traits, health issues, college majors, and other factors. Another interesting research topic that we would like to explore is on supporting students in staying away from states that could lead to depression or helping them move out of depressed states if already experiencing symptoms of depression.

REFERENCES

- [1] "American College Health Association-National College Health Assessment IIc," Hanover, MD, 2015.
- [2] S. L. Hedden, "Behavioral health trends in the United States: Results from the 2014 National Survey on Drug Use and Health," 2015.
- [3] American Psychiatric Association, *Diagnostic and statistical manual of mental disorders: DSM-IV*, 4th ed. Washington, DC, 1994.
- [4] S. C. Tanaka, K. Doya, G. Okada, K. Ueda, Y. Okamoto, and S. Yamawaki, "Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops," *Behav. Econ. Prefer. Choices, Happiness*, no. August, pp. 593–616, 2016.
- [5] S. M. McClure, D. I. Laibson, G. Loewenstein, and J. D. Cohen, "Separate neural systems value immediate and delayed monetary rewards.," *Science*, vol. 306, no. 5695, pp. 503–507, 2004.
- [6] E. Pulcu *et al.*, "Temporal discounting in major depressive disorder.," *Psychol. Med.*, vol. 44, no. 9, pp. 1825–34, 2014.
- [7] L. A. Giordano, W. K. Bickel, G. Loewenstein, E. A. Jacobs, L. Marsch, and G. J. Badger, "Mild opioid deprivation increases the degree that opioid-dependent outpatients discount delayed heroin and money," *Psychopharmacology (Berl.)*, vol. 163, no. 2, pp. 174–182, 2002.
- [8] W. Mischel, E. B. Ebbesen, and a R. Zeiss, "Cognitive and attentional mechanisms in delay of gratification.," *J. Pers. Soc. Psychol.*, vol. 21, no. 2, pp. 204–218, 1972.
- [9] W. Mischel, Y. Shoda, and M. I. Rodriguez, "Delay of gratification in children.," *Science*, vol. 244, no. 4907, pp. 933–8, 1989.
- [10] O. Ayduk, R. Mendoza-Denton, W. Mischel, G. Downey, P. K. Peake, and M. Rodriguez, "Regulating the interpersonal self: strategic self-regulation for coping with rejection sensitivity.," *J. Pers. Soc. Psychol.*, vol. 79, no. 5, pp. 776–792, 2000.
- [11] T. R. Schlam, N. L. Wilson, Y. Shoda, W. Mischel, and O. Ayduk, "Preschoolers' delay of gratification predicts their body mass 30 years later.," *J. Pediatr.*, vol. 162, no. 1, pp. 90–93, 2013.
- [12] Y. Shoda, W. Mischel, and P. K. Peake, "Predicting adolescent cognitive and self-regulatory competencies from preschool delay of gratification: Identifying diagnostic conditions.," *Dev. Psychol.*, vol. 26, no. 6, pp. 978–986, 1990.
- [13] M. Rabbi, M. H. Aung, M. Zhang, and T. Choudhury, "MyBehavior: automatic personalized health feedback from user behaviors and preferences using smartphones," *Proc. 2015 ACM Int. Jt. Conf. Pervasive Ubiquitous Comput.*, pp. 707–718, 2015.
- [14] B. Fogg, "A behavior model for persuasive design," *Proc. 4th Int. Conf. Persuas. Technol. - Persuas. '09*, p. 1, 2009.
- [15] P. Kraft, S. R. Sutton, and H. M. Reynolds, "The transtheoretical model of behaviour change: Are the stages qualitatively different?," *Psychol. Health*, vol. 14, no. 3, pp. 433–450, Dec. 2007.
- [16] N. Banovic *et al.*, "Leveraging human routine models to detect and generate human behaviors," *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. 2017.
- [17] N. Banovic, T. Buzali, F. Chevalier, J. Mankoff, and A. K. Dey, "Modeling and Understanding Human Routine Behavior," *Proc. 2016 CHI Conf. Hum. Factors Comput. Syst. - CHI '16*, pp. 248–260, 2016.
- [18] R. S. Sutton and A. G. Barto, "Reinforcement Learning: An Introduction," 2012.
- [19] L. P. Kaelbling, M. L. Littman, and A. W. Moore, "Reinforcement learning: A survey," *J. Artif. Intell. Res.*, vol. 4, pp. 237–285, 1996.
- [20] T. V. Maia, "Reinforcement learning, conditioning, and the brain: Successes and challenges.," *Cogn. Affect. Behav. Neurosci.*, vol. 9, no. 4, pp. 343–64, 2009.
- [21] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Mach. Learn.*, vol. 8, no. 3–4, pp. 279–292, 1992.
- [22] C. J. C. H. Watkins, "Learning from delayed rewards," University of Cambridge England, 1989.
- [23] A. Geramifard, "A Tutorial on Linear Function Approximators for Dynamic Programming and Reinforcement Learning," *Found. Trends@ Mach. Learn.*, vol. 6, no. 4, pp. 375–451, 2013.
- [24] R. Wang, F. Chen, Z. Chen, and T. Li, "StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones," *Proc. ACM Conf. Ubiquitous Comput.*, 2014.
- [25] S. Cohen and S. D. Pressman, "Positive affect and health," *Curr. Dir. Psychol. Sci.*, vol. 15, no. 3, pp. 122–125, 2006.
- [26] K. Kroenke, R. L. Spitzer, and J. B. W. Williams, "The PHQ-9," *J. Gen. Intern. Med.*, vol. 16, pp. 605–613, 2001.
- [27] American Psychiatric Association, *Diagnostic and Statistical Manual of Mental Disorders, 5th Edition*. 2013.
- [28] J. P. J. Pollak, P. Adams, and G. Gay, "PAM: a photographic affect meter for frequent, in situ measurement of affect," *Proc. SIGCHI Conf. Hum. factors Comput. Syst.*, pp. 725–734, 2011.
- [29] D. Pelleg and A. Moore, "X-means: Extending K-means with efficient estimation of the number of clusters," *Proc. Seventeenth Int. Conf. Mach. Learn. table contents*, pp. 727–734, 2000.
- [30] J. A. Russell, "A circumplex model of affect," *J. Pers. Soc. Psychol.*, vol. 39, no. 6, pp. 1161–1178, 1980.
- [31] F. Rubinsteyn Alex and Sergey, "fancyImpute 0.0.19," 2016. [Online]. Available: <https://pypi.python.org/pypi/fancyimpute>.
- [32] R. Mazumder, T. Hastie, and R. Tibshirani, "Spectral Regularization Algorithms for Learning Large Incomplete Matrices.," *J. Mach. Learn. Res.*, vol. 11, pp. 2287–2322, 2010.
- [33] U.S. Department of Health and Human Services, "2008 Physical activity guidelines for Americans," *Pres. Counc. Phys. Fit. Sport. Res. Dig.*, vol. 9, no. 4, pp. 1–8, 2008.
- [34] K. Kroenke, T. W. Strine, R. L. Spitzer, J. B. W. Williams, J. T. Berry, and A. H. Mokdad, "The PHQ-8 as a measure of current depression in the general population," *J. Affect. Disord.*, vol. 114, no. 1–3, pp. 163–173, 2009.
- [35] A. Bandura, "Social learning theory," *Social Learning Theory*. pp. 1–46, 1971.