An Algorithm to Determine Treatment Timing in Mobile Health: Potential, Design and Evaluation

Xiang Meng
xmeng@g.harvard.edu

14th CMStatistics, Advances in statistical methods for mobile health.

18 December 2021

Collaborators

- Harvard University: Peng Liao, Susan Murphy
- University of Michigan: Walter Dempsey, Predrag (Pedja) Klasnja
- University of Washington: Nick Reid

Contents

JITAIs have great potential in health behavior change

- A Just-in-Time Adaptive Intervention, or JITAI is an intervention design that provides the right type (or amount) of support, at the right time.
- Health behavior change is hard, but a JITAI holds enormous potential.
  - Eg: Physical inactivity [1], alcohol use [2], mental illness [3], smoking [4], and obesity [5].

Current issues in JITAIs

- Many JITAIs have been developed with minimal use of empirical evidence, theory, or guidelines
- Gap: Technological capabilities for delivering JITAIs $\Longleftrightarrow$ Research on their development and evaluation.

Design Principles: One opportunity + Two concerns

Design Principles: One opportunity

- One opportunity
  - To capture the right delivery times (risk times) [6].
  - Eg: Send a message to encourage exercise **only when** the user is **sedentary**

Design Principles: Two concerns

- Two concerns

1

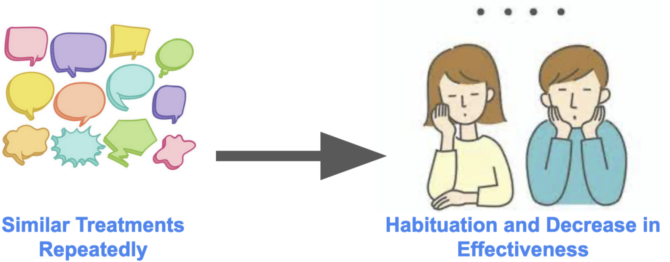

2

We want an algorithm that can capture the opportunity and resolve two concerns

- To capture the opportunity, we only focus on risk times
- To address the concerns
  1. To reduce burden $\longrightarrow$ Constraints on the number of interventions (Goal 1)
  2. To reduce habituation $\longrightarrow$ Probabilistic interventions (instead of deterministic ones)

Evaluation of the algorithm is important

- Is the algorithm good? How to improve?
- Key Question: How often, and in what context should a mobile intervention be provided?
- Algorithm: Uniform sampling on risk times (Goal 2).
- Experimental design: Micro-Randomized Trial (MRT).
- Our task: To design a MRT to address the key question.

Example: MRT to Reduce Sedentary Behavior

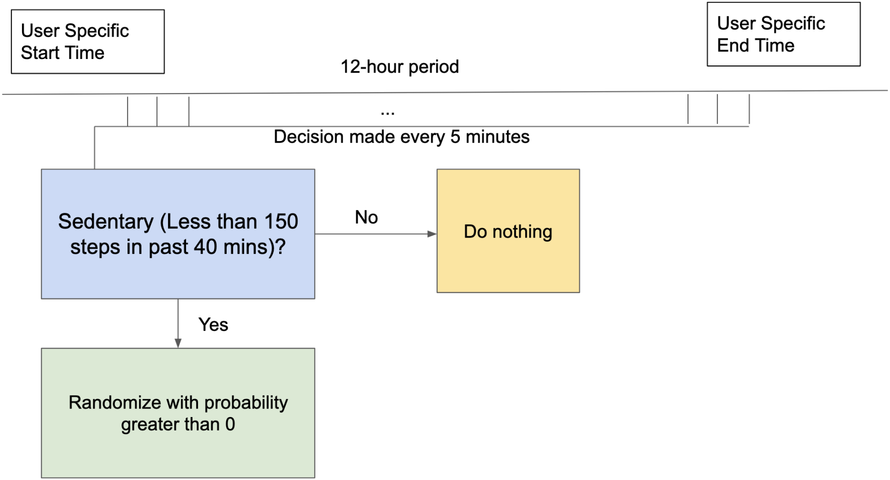

HeartSteps Study: an actual study that implements the MRT

- Aim: Develop a mobile activity assistant for individuals who are newly diagnosed with Stage 1 Hypertension
- Intervention: A push notification to encourage exercise.
- Decision time: Every 5 mins
- Each day is separated into three 4-hour blocks. Denote $\mathcal{B}_{d,k}$ as kth block on day d.
- A_t: intervention; H_t: history; I_t: sedentary indicator; π_t: randomization probability.

Sampling Algorithm: to address the constraint

- Goal 1: Constraint: An average of 0.5 Intervention per Block $\mathcal{B}_{d,k}$
- To achieve Goal 1: We set constraint on

$$N_1^* := E[\sum_{t \in \mathcal{B}_{d,k}, I_t = 1} A_t]$$

Expected/Average number of interventions in block k on day d

Sampling Algorithm: to achieve the uniform distribution

- To achieve the goal of uniform distribution, we calculate
  - The prediction of the number of remaining interventions

$$
\begin{aligned}
h_t(N_1^*) =& N_1^* - E[\sum_{s \in \mathcal{B}_{d,k}, s \leq t-1, I_s = 1} A_s] \\
=& N_1^* - \sum_{s \in \mathcal{B}_{d,k}, s \leq t-1, I_s = 1} \pi_s
\end{aligned}
$$

  - The prediction of the remaining sedentary times

$$
g_t = E[\sum_{s \in \mathcal{B}_{d,k}, s \geq t+1} I_s]
$$

- At time t, if remaining sedentary times g_t is known, by setting $N_1^* = 0.5$ we can sample uniformly with probability

$$
\pi_t = \frac{h_t(N_1^*)}{1 + g_t}
$$

- Truncate to $[0.005, 0.995]$

A practical challenge to the algorithm

The randomization probability:

$$\pi_t = \frac{h_t(N_1^*)}{1 + g_t}$$

- In practice, a message might reduce the number of sedentary times in the future.
- i.e., g_t will be smaller if a message is delivered at $t-1$ (compared to not delivered at $t-1$).

Adjustment to the algorithm in practice

- Solution: We let N_1^* and g_t be two "tuning parameters".
    - Tune using data from the previous batch of the HeartSteps [7].
- Our hope is that with the two tuning parameters, our two goals can be achieved on average.

Evaluation of the algorithm is important

Evaluation of the algorithm is important because we

- Need to check whether 2 goals are met.
- Want to know the direction of improvement.

Evaluation of the algorithm with user study has not been an interest

Two types of related work have not address it.

1. Real user studies that apply online algorithm
   - Eg. To encourage physical activity for people with diabetes [8], [9]; stress management [10]; weight loss support [11]
   - However, their analysis focused on whether the algorithm, as a treatment has a causal effect.
2. Analysis of online algorithms, but all theoretical
   - Eg. Off-Policy RL in mHealth [12], [13]; Efficient policy learning [14], [15]

How many interventions were delivered in each block and on each day?

- Calculation for each user:
  1. Calculate the total number interventions over all days (or over all kth blocks).
  2. Calculate total number of days (or number of the kth blocks).
  3. Calculate the average by dividing 1 by 2
- Then do a summary for all users. Total 81 users are considered.
- Key: What days do we average over?

To determine number of days is non-trivial in practice

- Practical challenges:
  - Users don't wear Fitbit every day.
  - Even if they wear it, they may not connect to server.
  - Some days we encounter bugs...
- There are a lot of days during the study that we cannot do anything at sedentary times.
- Need to detect them based on data we have, carefully.

How many interventions were delivered in each block and on each day?

- We consider days when the server recognizes at least one step.
- Goal: 0.5 per block and 1.5 per day.

	Median	Mean	Std. Dev.
Block 1	0.465	0.421	0.212
Block 2	0.412	0.428	0.248
Block 3	0.412	0.440	0.247
Daily	1.342	1.407	0.602

- Algorithm is not bad at targeting the goal.
- Differences are due to different wearing behaviour and different number of sedentary times.

Is the Algorithm Uniform Across all Sedentary Times?

- Measure: KL divergence **between the actual** treatment probability and the **uniform** distribution defined by 0.5/(number of sedentary times).
- For each user, we calculate mean KL divergence on each (day, block).

A Comparison Using KL Divergence

A baseline algorithm—Block sampling

- For each block k, use the same randomization probability for all users and on all days.

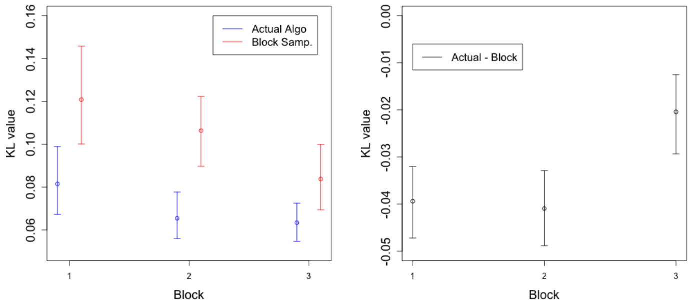

Figure: Left: KL between the actual probability(blue)/block sampling(red) and the uniform probability. Right: CI of their difference

Other Challenges in Practice

- Imputation for missing data in step count/heart rate for the proximal outcome.
- Deciding to drop which users.
- Deciding which one to use when we have one variable from different data sources.
  - Eg. step count from the server and Fitbit

Thanks for your attention!

Any questions?

References I

[1] Sunny Consolvo, David W McDonald, Tammy Toscos, Mike Y Chen, Jon Froehlich, Beverly Harrison, Predrag Klasnja, Anthony LaMarca, Louis LeGrand, Ryan Libby, et al. Activity sensing in the wild: a field trial of ubifit garden. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1797–1806, 2008.

[2] David H Gustafson, Fiona M McTavish, Ming-Yuan Chih, Amy K Atwood, Roberta A Johnson, Michael G Boyle, Michael S Levy, Hilary Driscoll, Steven M Chisholm, Lisa Dillenburg, et al. A smartphone application to support recovery from alcoholism: a randomized clinical trial. *JAMA psychiatry*, 71(5):566–572, 2014.

[3] Dror Ben-Zeev, Susan M Kaiser, Christopher J Brenner, Mark Begale, Jennifer Duffecy, and David C Mohr. Development and usability testing of focus: A smartphone system for self-management of schizophrenia. *Psychiatric rehabilitation journal*, 36(4):289, 2013.

References II

[4] William Riley, Jami Obermayer, and Jersino Jean-Mary. Internet and mobile phone text messaging intervention for college smokers. *Journal of American College Health*, 57(2):245–248, 2008.

[5] Kevin Patrick, Fred Raab, Marc Adams, Lindsay Dillon, Marion Zabinski, Cheryl Rock, William Griswold, and Gregory Norman. A text message-based intervention for weight loss: randomized controlled trial. *Journal of medical Internet research*, 11(1):e1, 2009.

[6] Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. Just-in-time adaptive interventions (jitais) in mobile health: key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine*, 52(6):446–462, 2018.

References III

[7] Predrag Klasnja, Shawna Smith, Nicholas J Seewald, Andy Lee, Kelly Hall, Brook Luers, Eric B Hekler, and Susan A Murphy. Efficacy of contextually tailored suggestions for physical activity: a micro-randomized optimization trial of heartsteps. *Annals of Behavioral Medicine*, 53(6):573–582, 2019.

[8] Irit Hochberg, Guy Feraru, Mark Kozdoba, Shie Mannor, Moshe Tennenholtz, and Elad Yom-Tov. Encouraging physical activity in patients with diabetes through automatic personalized feedback via reinforcement learning improves glycemic control. *Diabetes care*, 39(4):e59–e60, 2016.

[9] Elad Yom-Tov, Guy Feraru, Mark Kozdoba, Shie Mannor, Moshe Tennenholtz, and Irit Hochberg. Encouraging physical activity in patients with diabetes: intervention using a reinforcement learning system. *Journal of medical Internet research*, 19(10):e338, 2017.

References IV

[10] Pablo Paredes, Ran Gilad-Bachrach, Mary Czerwinski, Asta Roseway, Kael Rowan, and Javier Hernandez. Poptherapy: Coping with stress through pop-culture. In *Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare*, pages 109–117, 2014.

[11] Evan M Forman, Stephanie G Kerrigan, Meghan L Butryn, Adrienne S Juarascio, Stephanie M Manasse, Santiago Ontañón, Diane H Dallal, Rebecca J Crochiere, and Danielle Moskow. Can the artificial intelligence technique of reinforcement learning use continuously-monitored digital data to optimize treatment for weight loss? *Journal of behavioral medicine*, 42(2):276–290, 2019.

References V

[12] Peng Liao, Walter Dempsey, Hillol Sarker, Syed Monowar Hossain, Mustafa Al'Absi, Predrag Klasnja, and Susan Murphy. Just-in-time but not too much: determining treatment timing in mobile health. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies*, 2(4):1–21, 2018.

[13] Daniel J Luckett, Eric B Laber, Anna R Kahkoska, David M Maahs, Elizabeth Mayer-Davis, and Michael R Kosorok. Estimating dynamic treatment regimes in mobile health using v-learning. *Journal of the American Statistical Association*, 2019.

[14] Susan Athey, Stefan Wager, et al. Efficient policy learning. *arXiv preprint arXiv:1702.02896*, 78, 2017.

[15] Toru Kitagawa and Aleksey Tetenov. Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616, 2018.

Model for Prediction of the Remaining Sedentary Times g_t

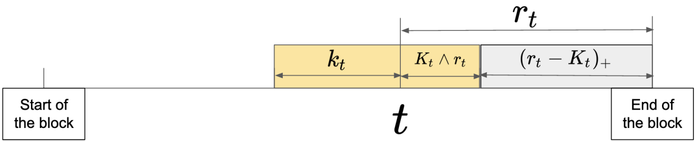

- Notation:
  - k_t: run length.
  - K_t: remaining run length of the current run
  - r_t: remaining time in the current block
  - F_t: the fraction of sedentary times outside the run.
- Estimate for g_t:

$$\hat{g}_t = \hat{E}[K_t \wedge r_t] + \hat{E}[F_t]\hat{E}[(r_t - K_t)_+]$$

- $\hat{E}[K_t \wedge r_t], \hat{E}[F_t], \hat{E}[(r_t - K_t)_+]$ are non-parametric averages from prior data.

Selection of N_1^* Given the Model for $g(1|H_t)$

Algorithm 1: Tune N_1^*

```
Input: Grid to be searched G_N; g(1|H_t), number of users n
Output: Tuned N_1^*
for N_1 ∈ G_N do
    for i ← 1 to n do
        for j ← 1 to 1000 do
            for t ∈ T_i do
                Set H_{i,t} = {A_{i,t-1}, π_{i,t-1}, I_{i,t}}
                Set {π_{i,t}} from the algorithm. Sample A_{i,t} ~ Bern(π_{i,t}).
            end
        end
        Store H_{i,j}(N_1)
    end
    for k ← 1 to 3 do
        Compute the average number of treatment over each block k
         as F_k(N_1)
    end
end
```

Why KL?

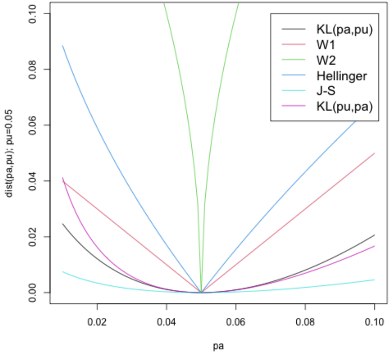

- We want to look at KL and Hellinger becauce they are two extremes but not as extreme as J-S and W2
- No large difference between KL(pa, pu) and KL(pu, pa); focus analysis on KL(pa, pu) because it looks more symmetrical
- We will replicate our analysis for KL(pu, pa) and Hellinger

Is the Algorithm Uniform Across all Sedentary Times?

- Notation: $N_{u,d,k}$ as the number of sedentary times per (user, day, 4-hour block).
- Measure: KL divergence **between the actual** treatment probability p_a and the **uniform** distribution defined by $p_u = 0.5/N_{u,d,k}$.
- We calculate mean KL divergence on each (user, day, 4-hour block) by

$$\frac{1}{N_{u,d,k}} \sum_{t=1}^{N_{u,d,k}} KL(p_{a,t}, p_{u,t})$$

where $KL(p, q) = p \log(p/q) + (1 - p) \log((1 - p)/(1 - q))$

A Comparison Using KL Divergence

A baseline algorithm—Block sampling

- First, from prior data, calculate average number of sedentary times across all (user, day) in block k: $\hat{M}_k, k = 1, 2, 3$.
- Then set the randomisation probability to be $0.5/\hat{M}_k$ for sedentary time in block k.

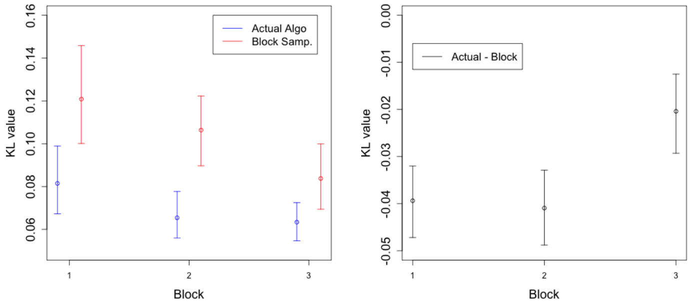

Figure: Left: KL between the actual probability(blue)/block sampling(red) and the uniform probability. Right: CI of their difference