Oralytics: A Mobile Health Study to Improve Oral Health Behaviors

Kelly W. Zhang, Anna Trella, Susan Murphy

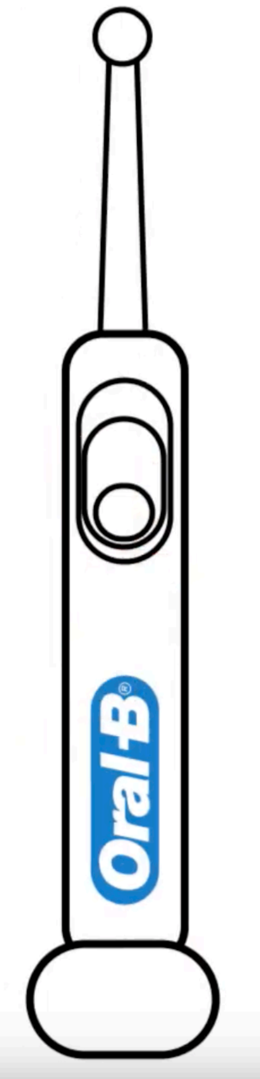

ORALYTICS

Harvard John A. Paulson
School of Engineering
and Applied Sciences

November 10, 2021

Statistical Reinforcement Learning Lab

- Mobile Health Clinical Trials
    - HeartSteps
    - Sense2Stop
    - Oralytics (focus of this talk!)

Statistical Reinforcement Learning Lab

- Mobile Health Clinical Trials
  - HeartSteps
  - Sense2Stop
  - Oralytics (focus of this talk!)
- Developing Learning Algorithms for Personalization
  - Bandit / Reinforcement Learning

ORALYTICS

Statistical Reinforcement Learning Lab

- Mobile Health Clinical Trials
    - HeartSteps
    - Sense2Stop
    - Oralytics (focus of this talk!)
- Developing Learning Algorithms for Personalization
    - Bandit / Reinforcement Learning
- Methods for Statistical Analysis

About this talk…

- Oralytics is a work in progress!
    - We are currently designing the algorithm for the trial which will go in the field spring 2022
    - We welcome and value your feedback!!

About this talk…

- Oralytics is a work in progress!
    - We are currently designing the algorithm for the trial which will go in the field spring 2022
    - We welcome and value your feedback!!
- Goals of this talk
    - Introduce you to using RL / personalizing algorithms in mobile health studies through our Oralytics study
    - Describe and get feedback on the challenges we're facing

Outline

1. **Introduction to Oralytics**
2. Just-In-Time Adaptive Interventions and Trial Overview
3. Design of Personalization / Reinforcement Learning Algorithm

Oralytics Study Collaborators (big team)!

- **Dentists**
    - Led by Vivek Shetty at UCLA
- **Behavioral Scientists**
    - Led by Inbal Nahum-Shani at UMichigan
- **Personalization Algorithm (our team)**
- **Signal Processing Team**
    - Led by Vwani Roychowhury at UCLA
- **App Developers**
    - Led by Cody Diefenthaler

The State of Dental Health

- $108 billon spent annually on dental services in US (2011 CDC report)
    - 95% spent on treating consequences of oral diseases
    - <5% spend on oral health promotion and disease prevention

The State of Dental Health

- $108 billon spent annually on dental services in US (2011 CDC report)
    - 95% spent on treating consequences of oral diseases
    - <5% spend on oral health promotion and disease prevention
- Dental disease linked to poor **Oral Hygiene Behaviors** (2009 UK report)
    - 33% of men brushed only once a day
    - 59% of women regularly skipped brushing at bedtime

The State of Dental Health

- $108 billon spent annually on dental services in US (2011 CDC report)
    - 95% spent on treating consequences of oral diseases
    - <5% spend on oral health promotion and disease prevention
- Dental disease linked to poor **Oral Hygiene Behaviors** (2009 UK report)
    - 33% of men brushed only once a day
    - 59% of women regularly skipped brushing at bedtime
- Dental disease burden falls disproportionately on vulnerable and underserved populations

Oralytics Study Overview

- Mobile health app to improve user's oral health behaviors

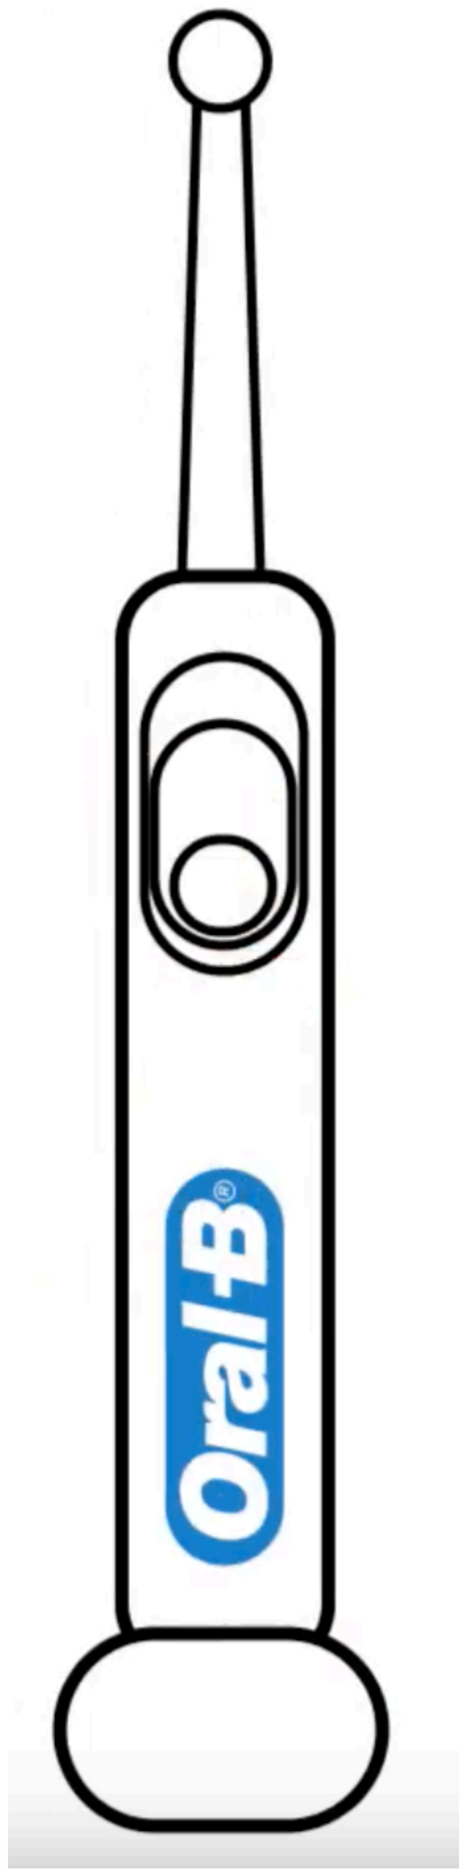

Oralytics Study Overview

- Mobile health app to improve user’s oral health behaviors
- **Digital Oral Health Intervention:** “push” messages to encourage toothbrushing

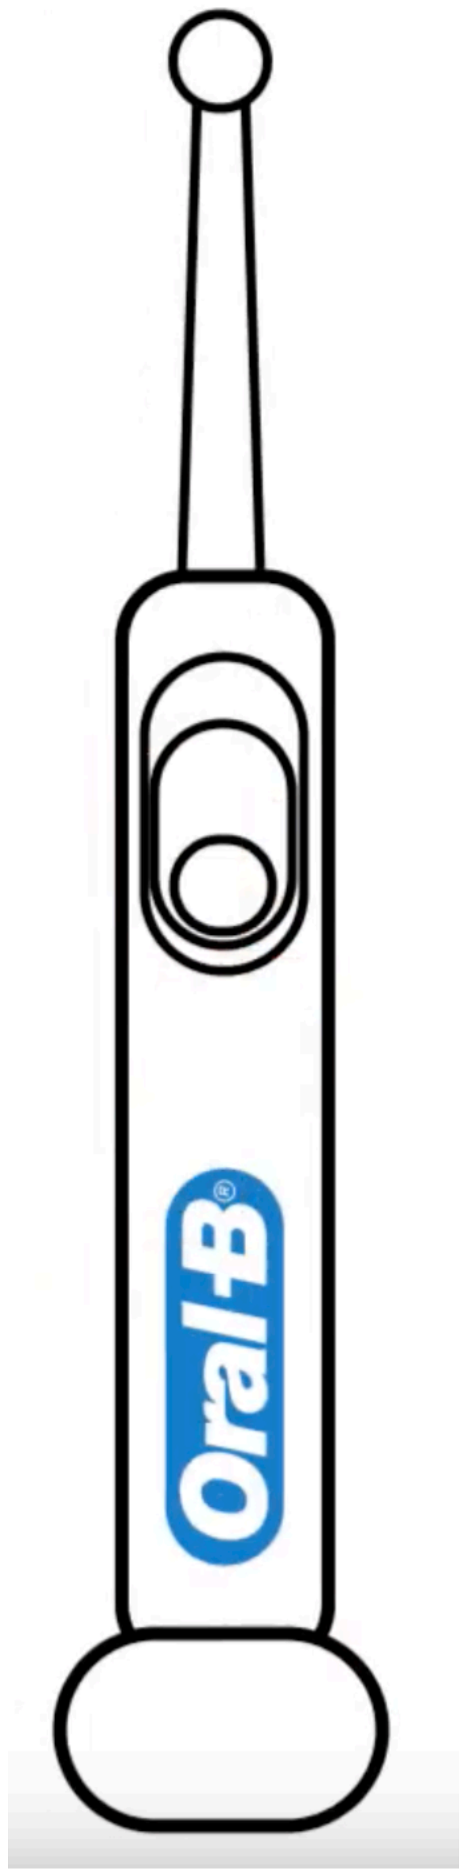

Oralytics Study Overview

- Mobile health app to improve user's oral health behaviors
- **Digital Oral Health Intervention:** "push" messages to encourage toothbrushing
- Sensor data includes mobile phone and bluetooth enabled toothbrush

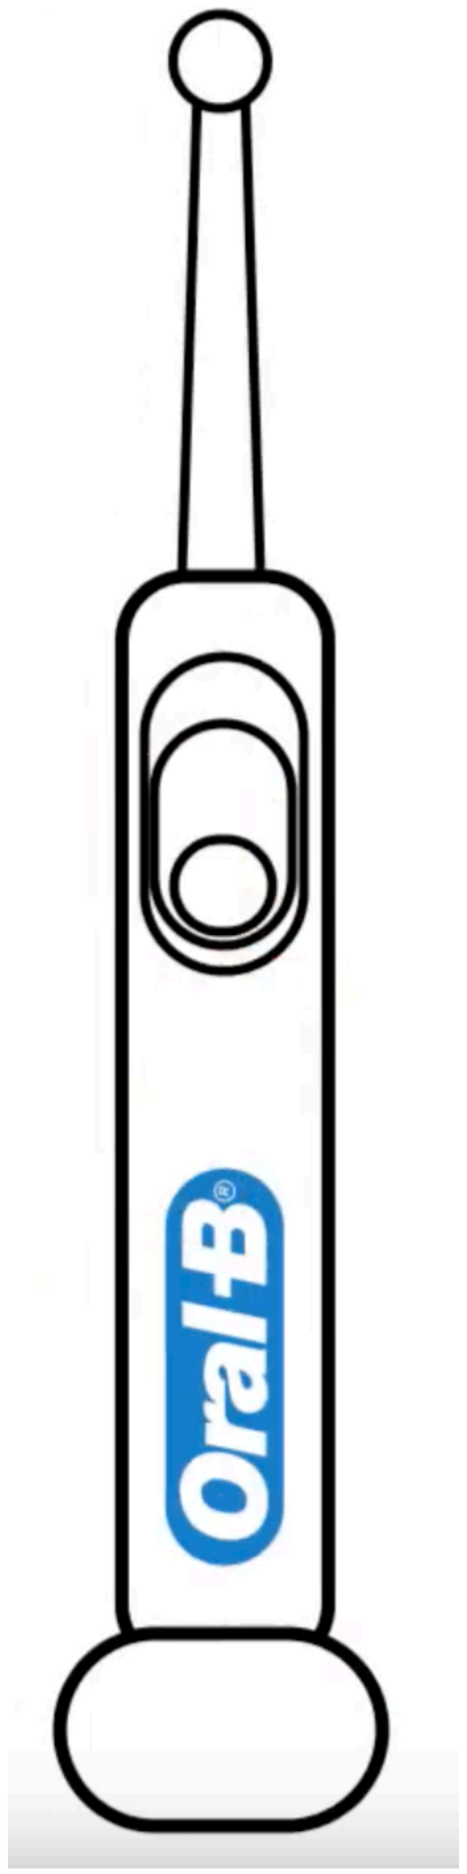

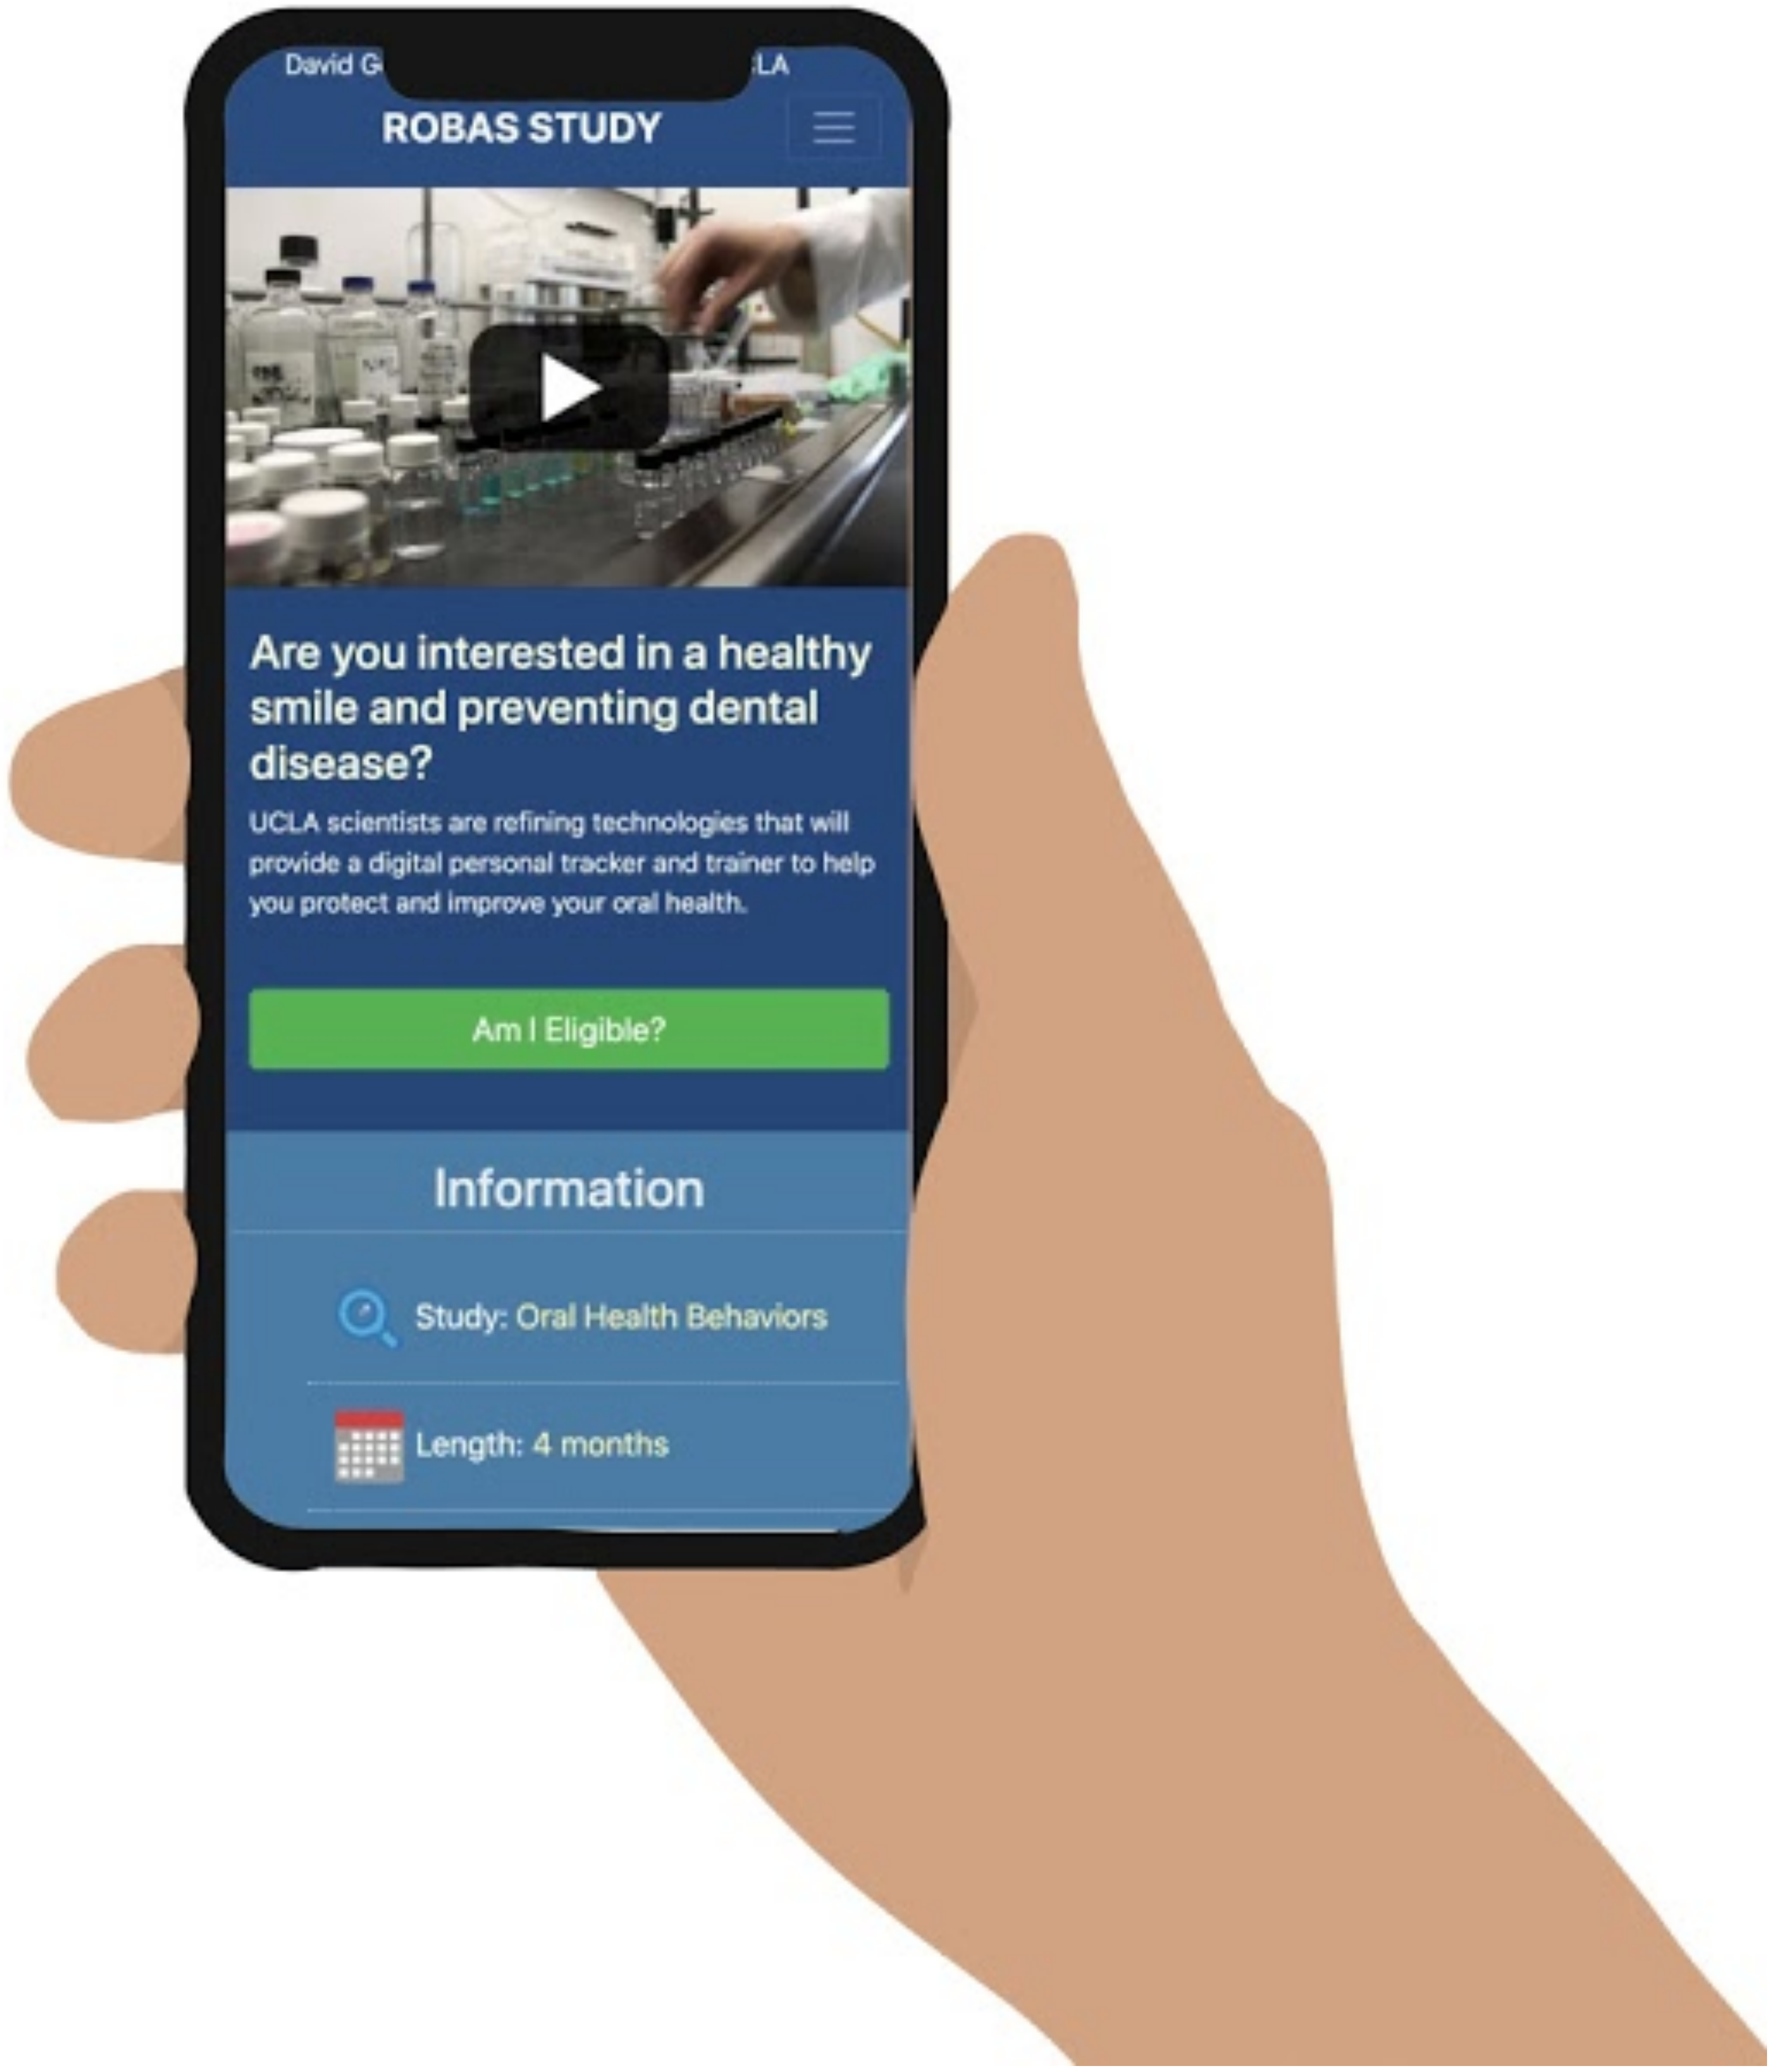

 Link to video: https://www.youtube.com/watch?v=XpI4O0vHYKg&t=5s


ORALYTICS

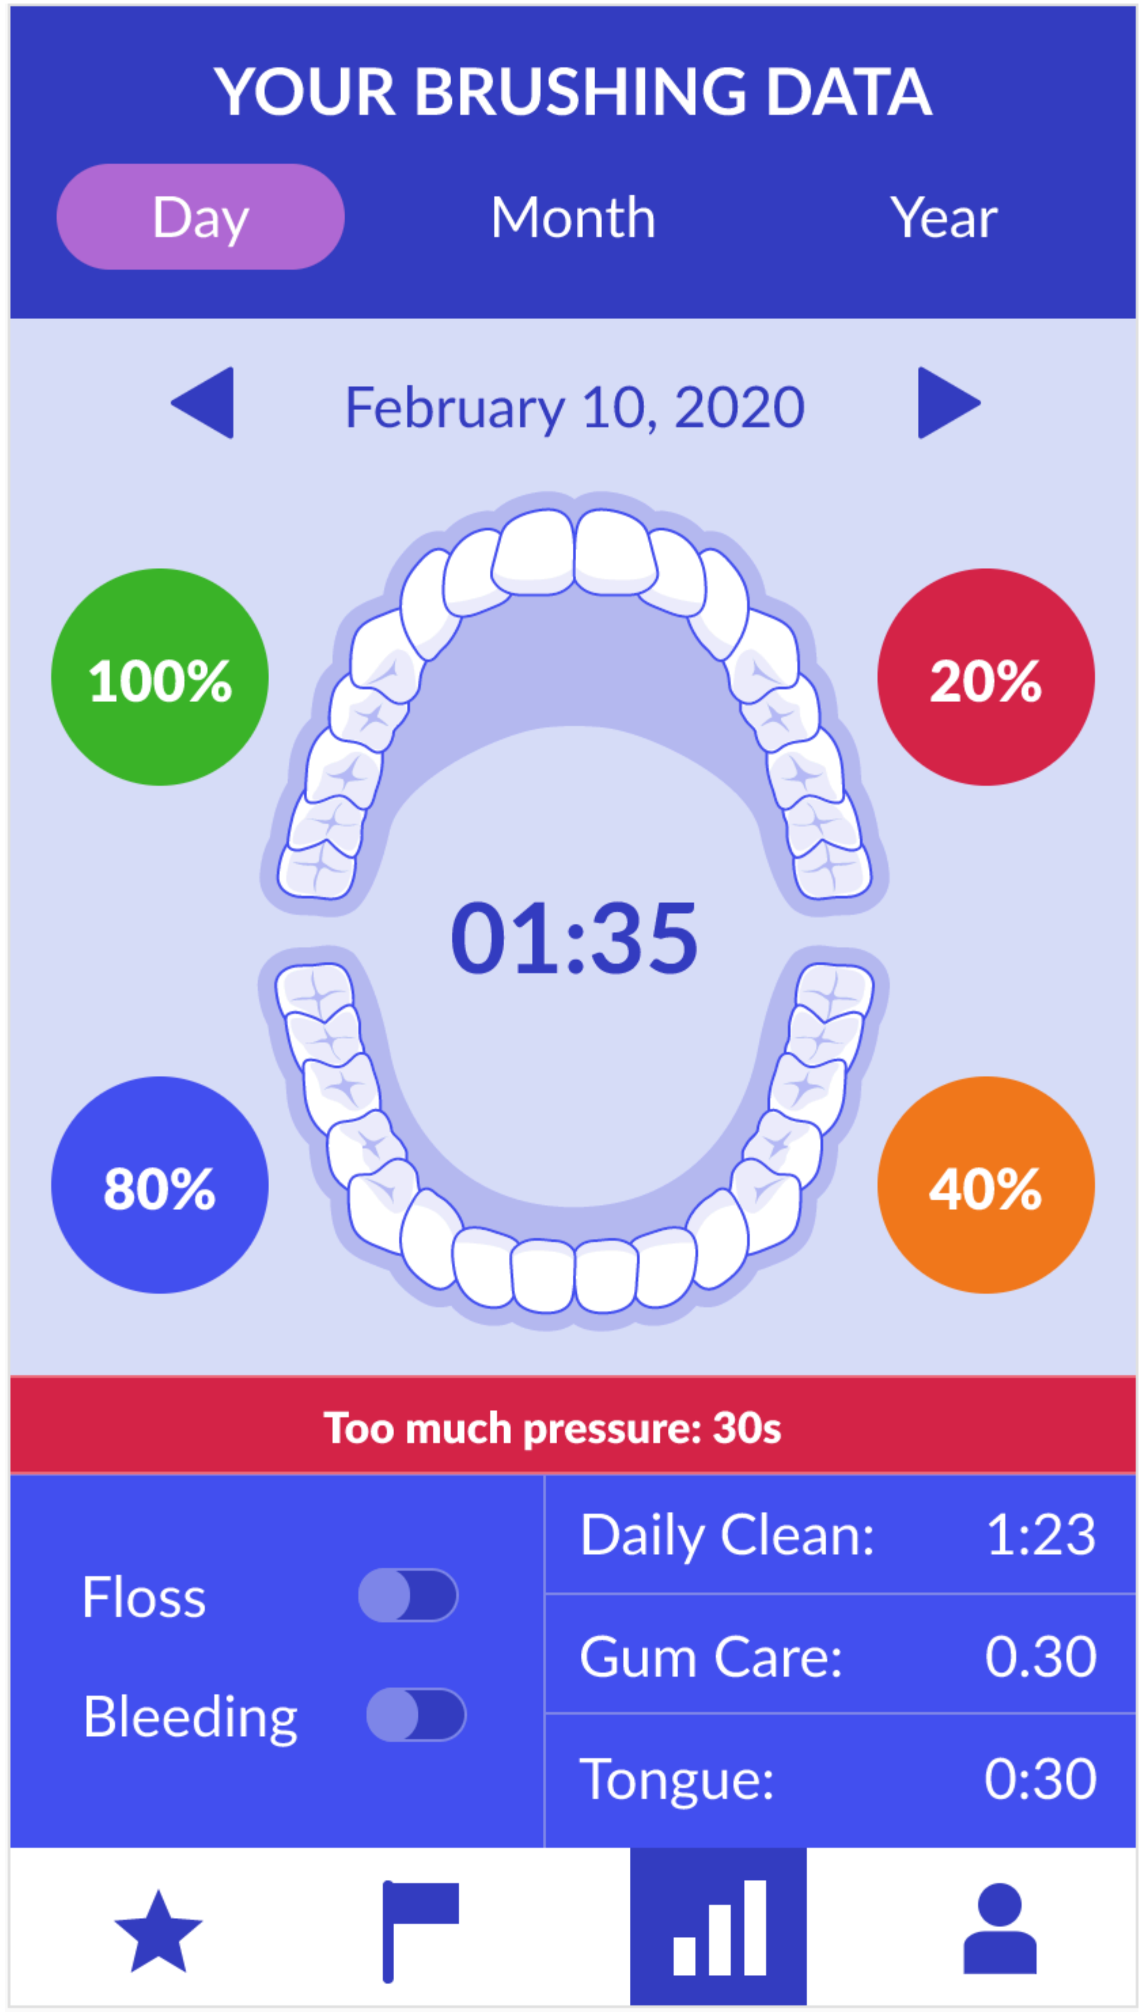
YOUR BRUSHING DATA
Day
Month
Year
February 10, 2020
100%
20%
01:35
80%
40%
Too much pressure: 30s
Floss
Bleeding
Daily Clean: 1:23
Gum Care: 0.30
Tongue: 0:30

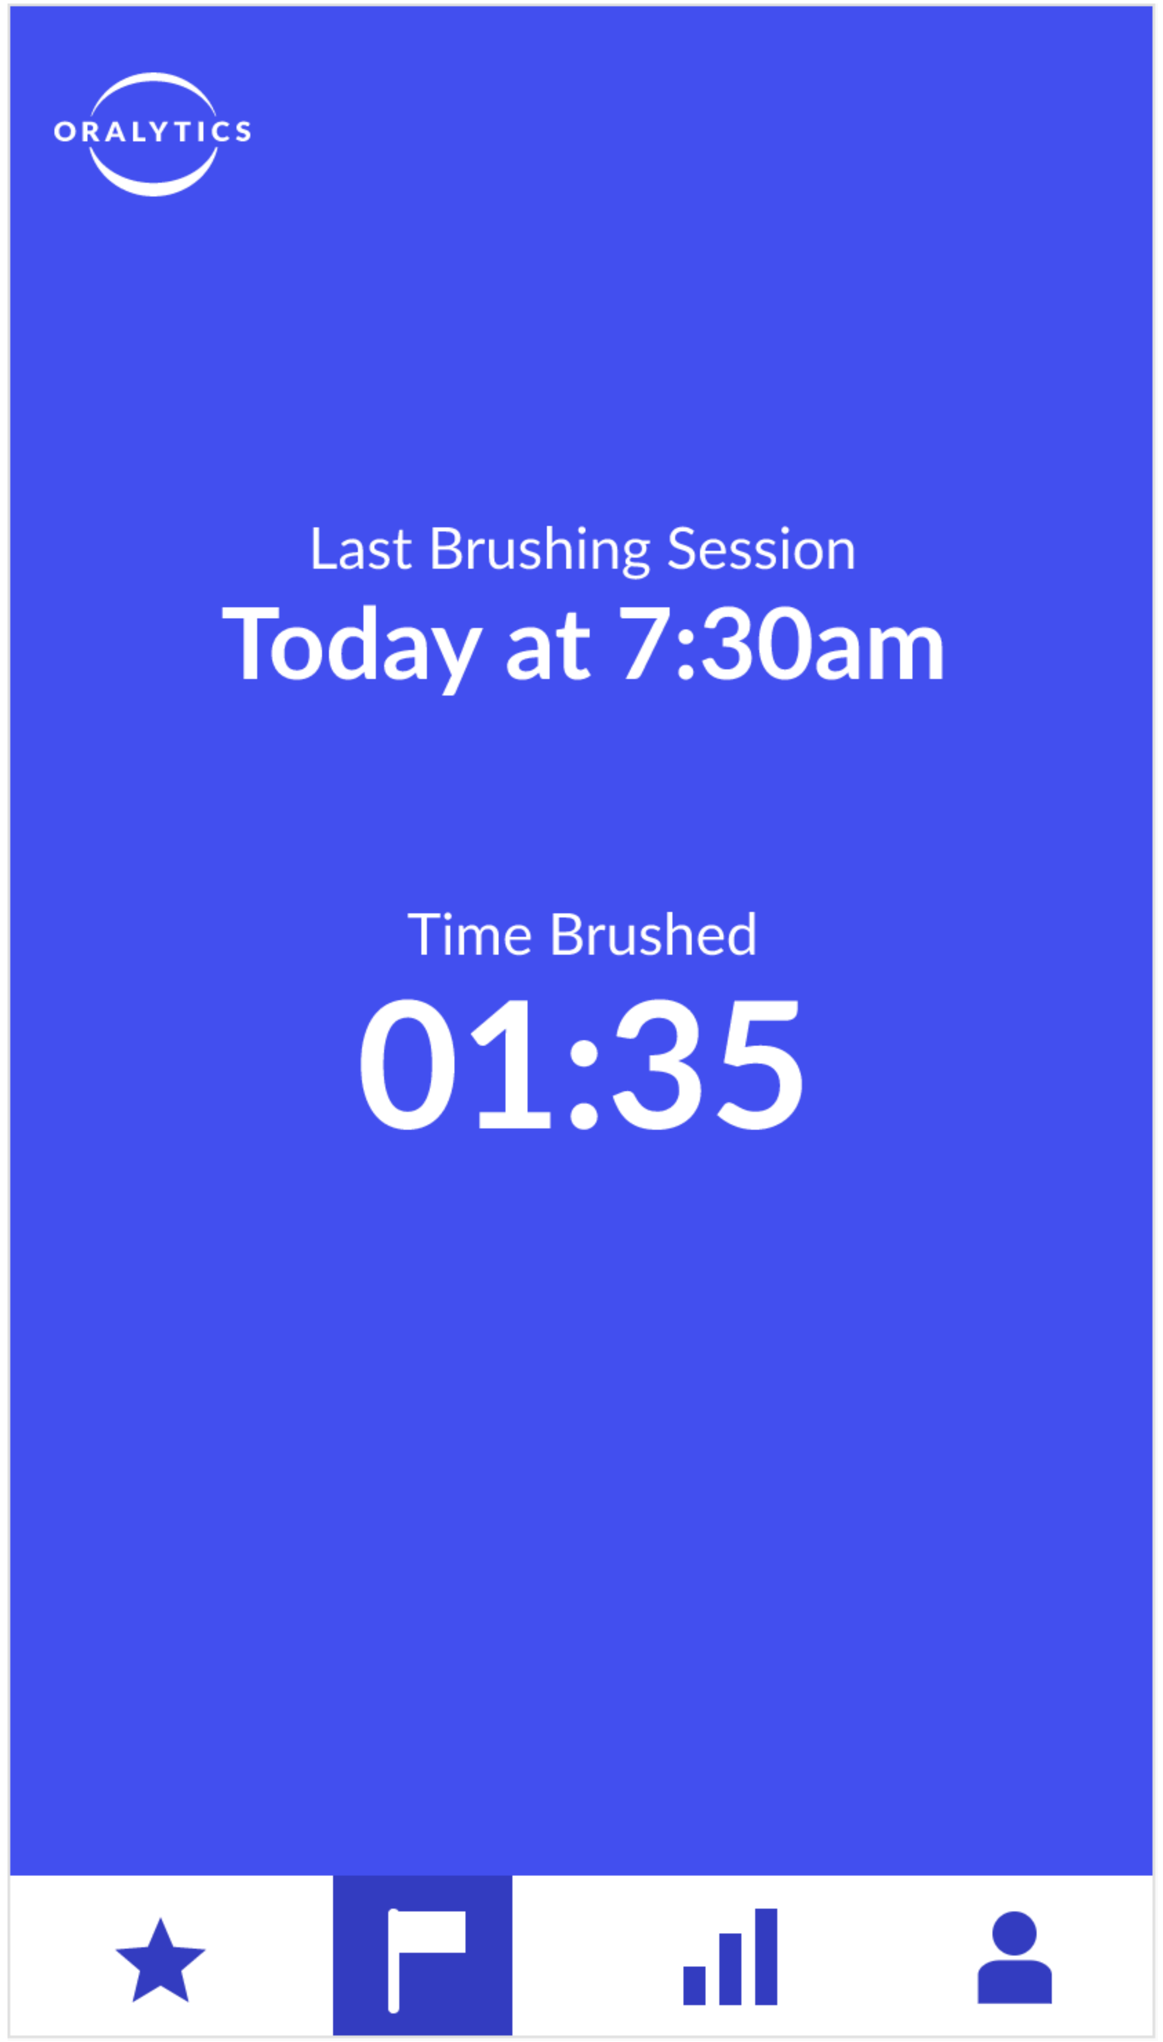
ORALYTICS
Last Brushing Session
Today at 7:30am
Time Brushed
01:35

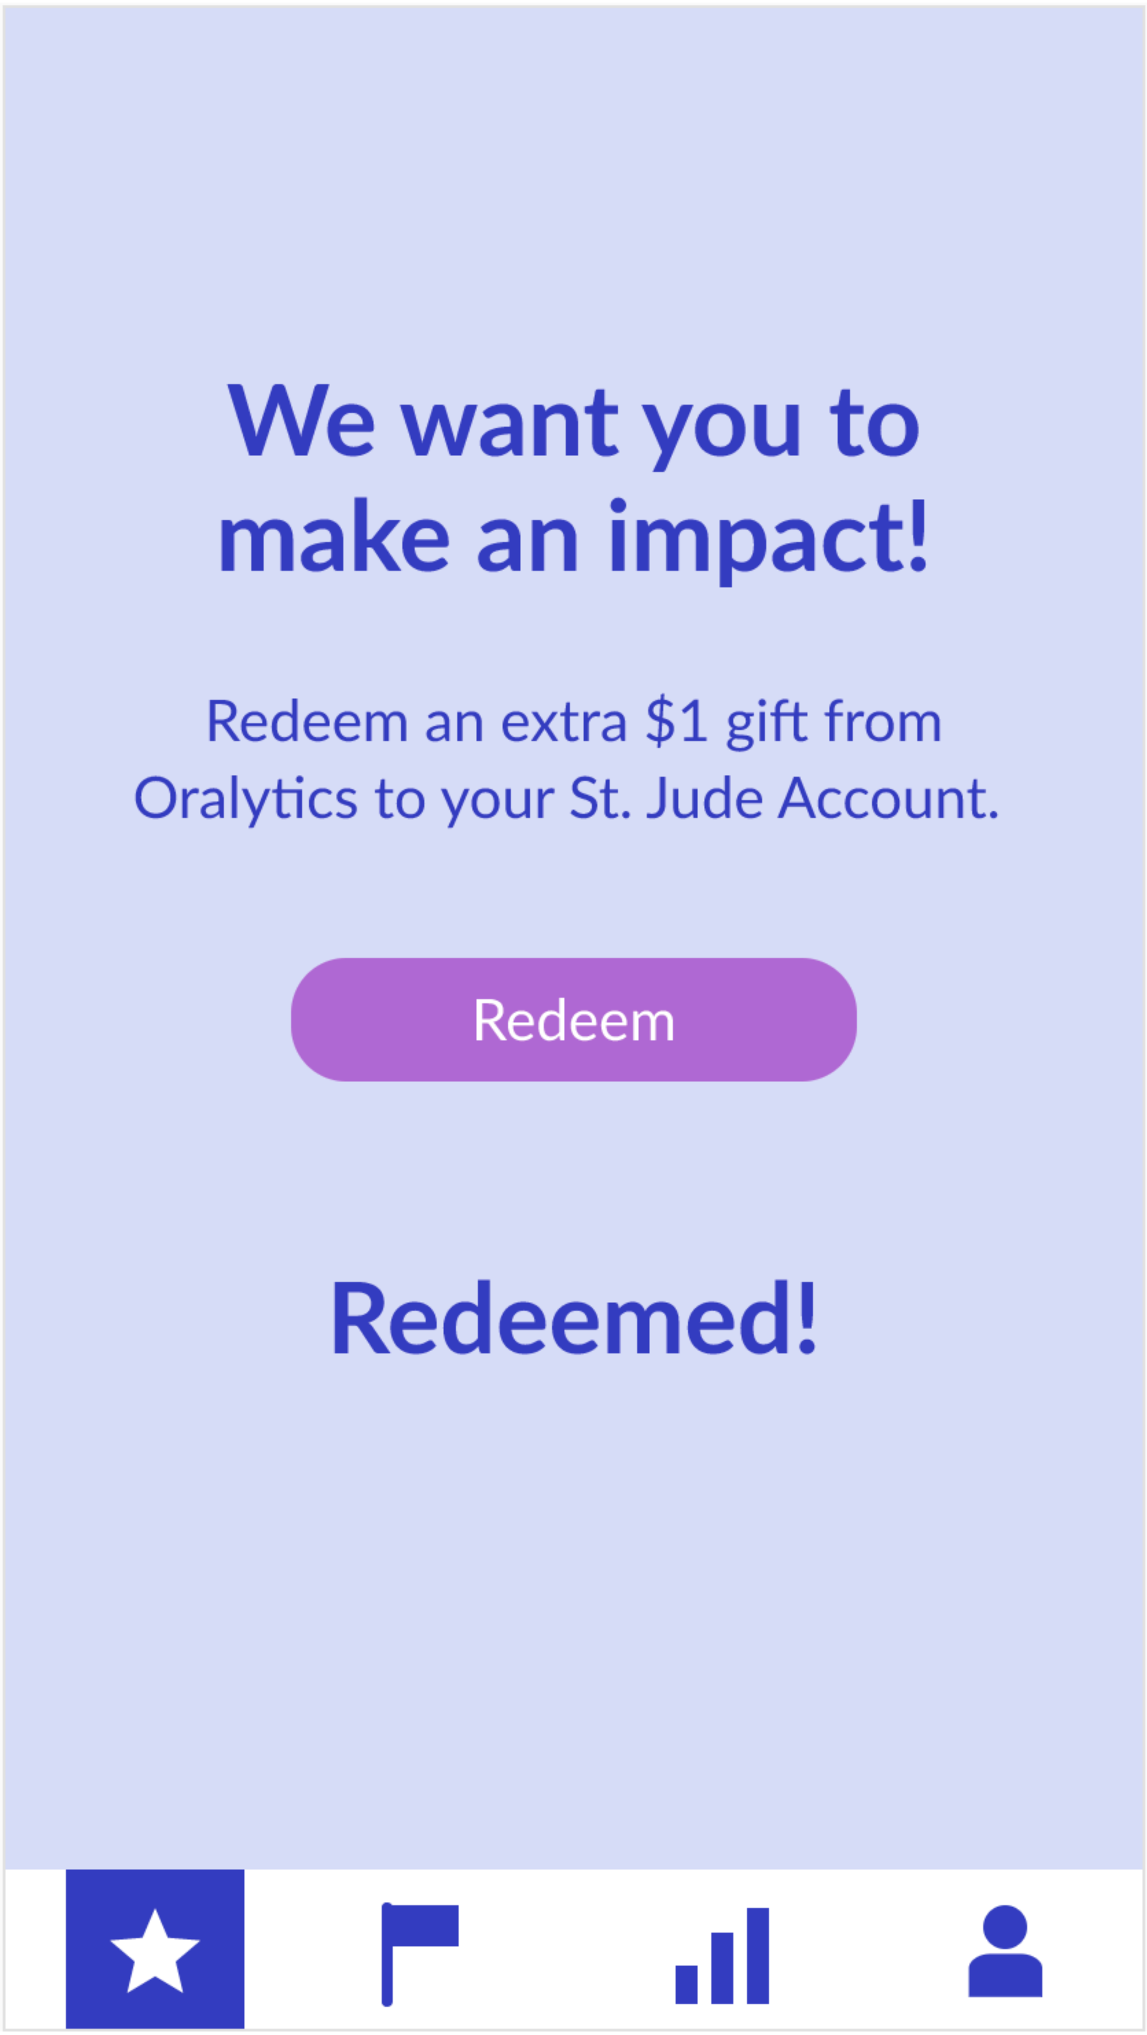
We want you to make an impact!
Redeem an extra $1 gift from Oralytics to your St. Jude Account.
Redeem
Redeemed!

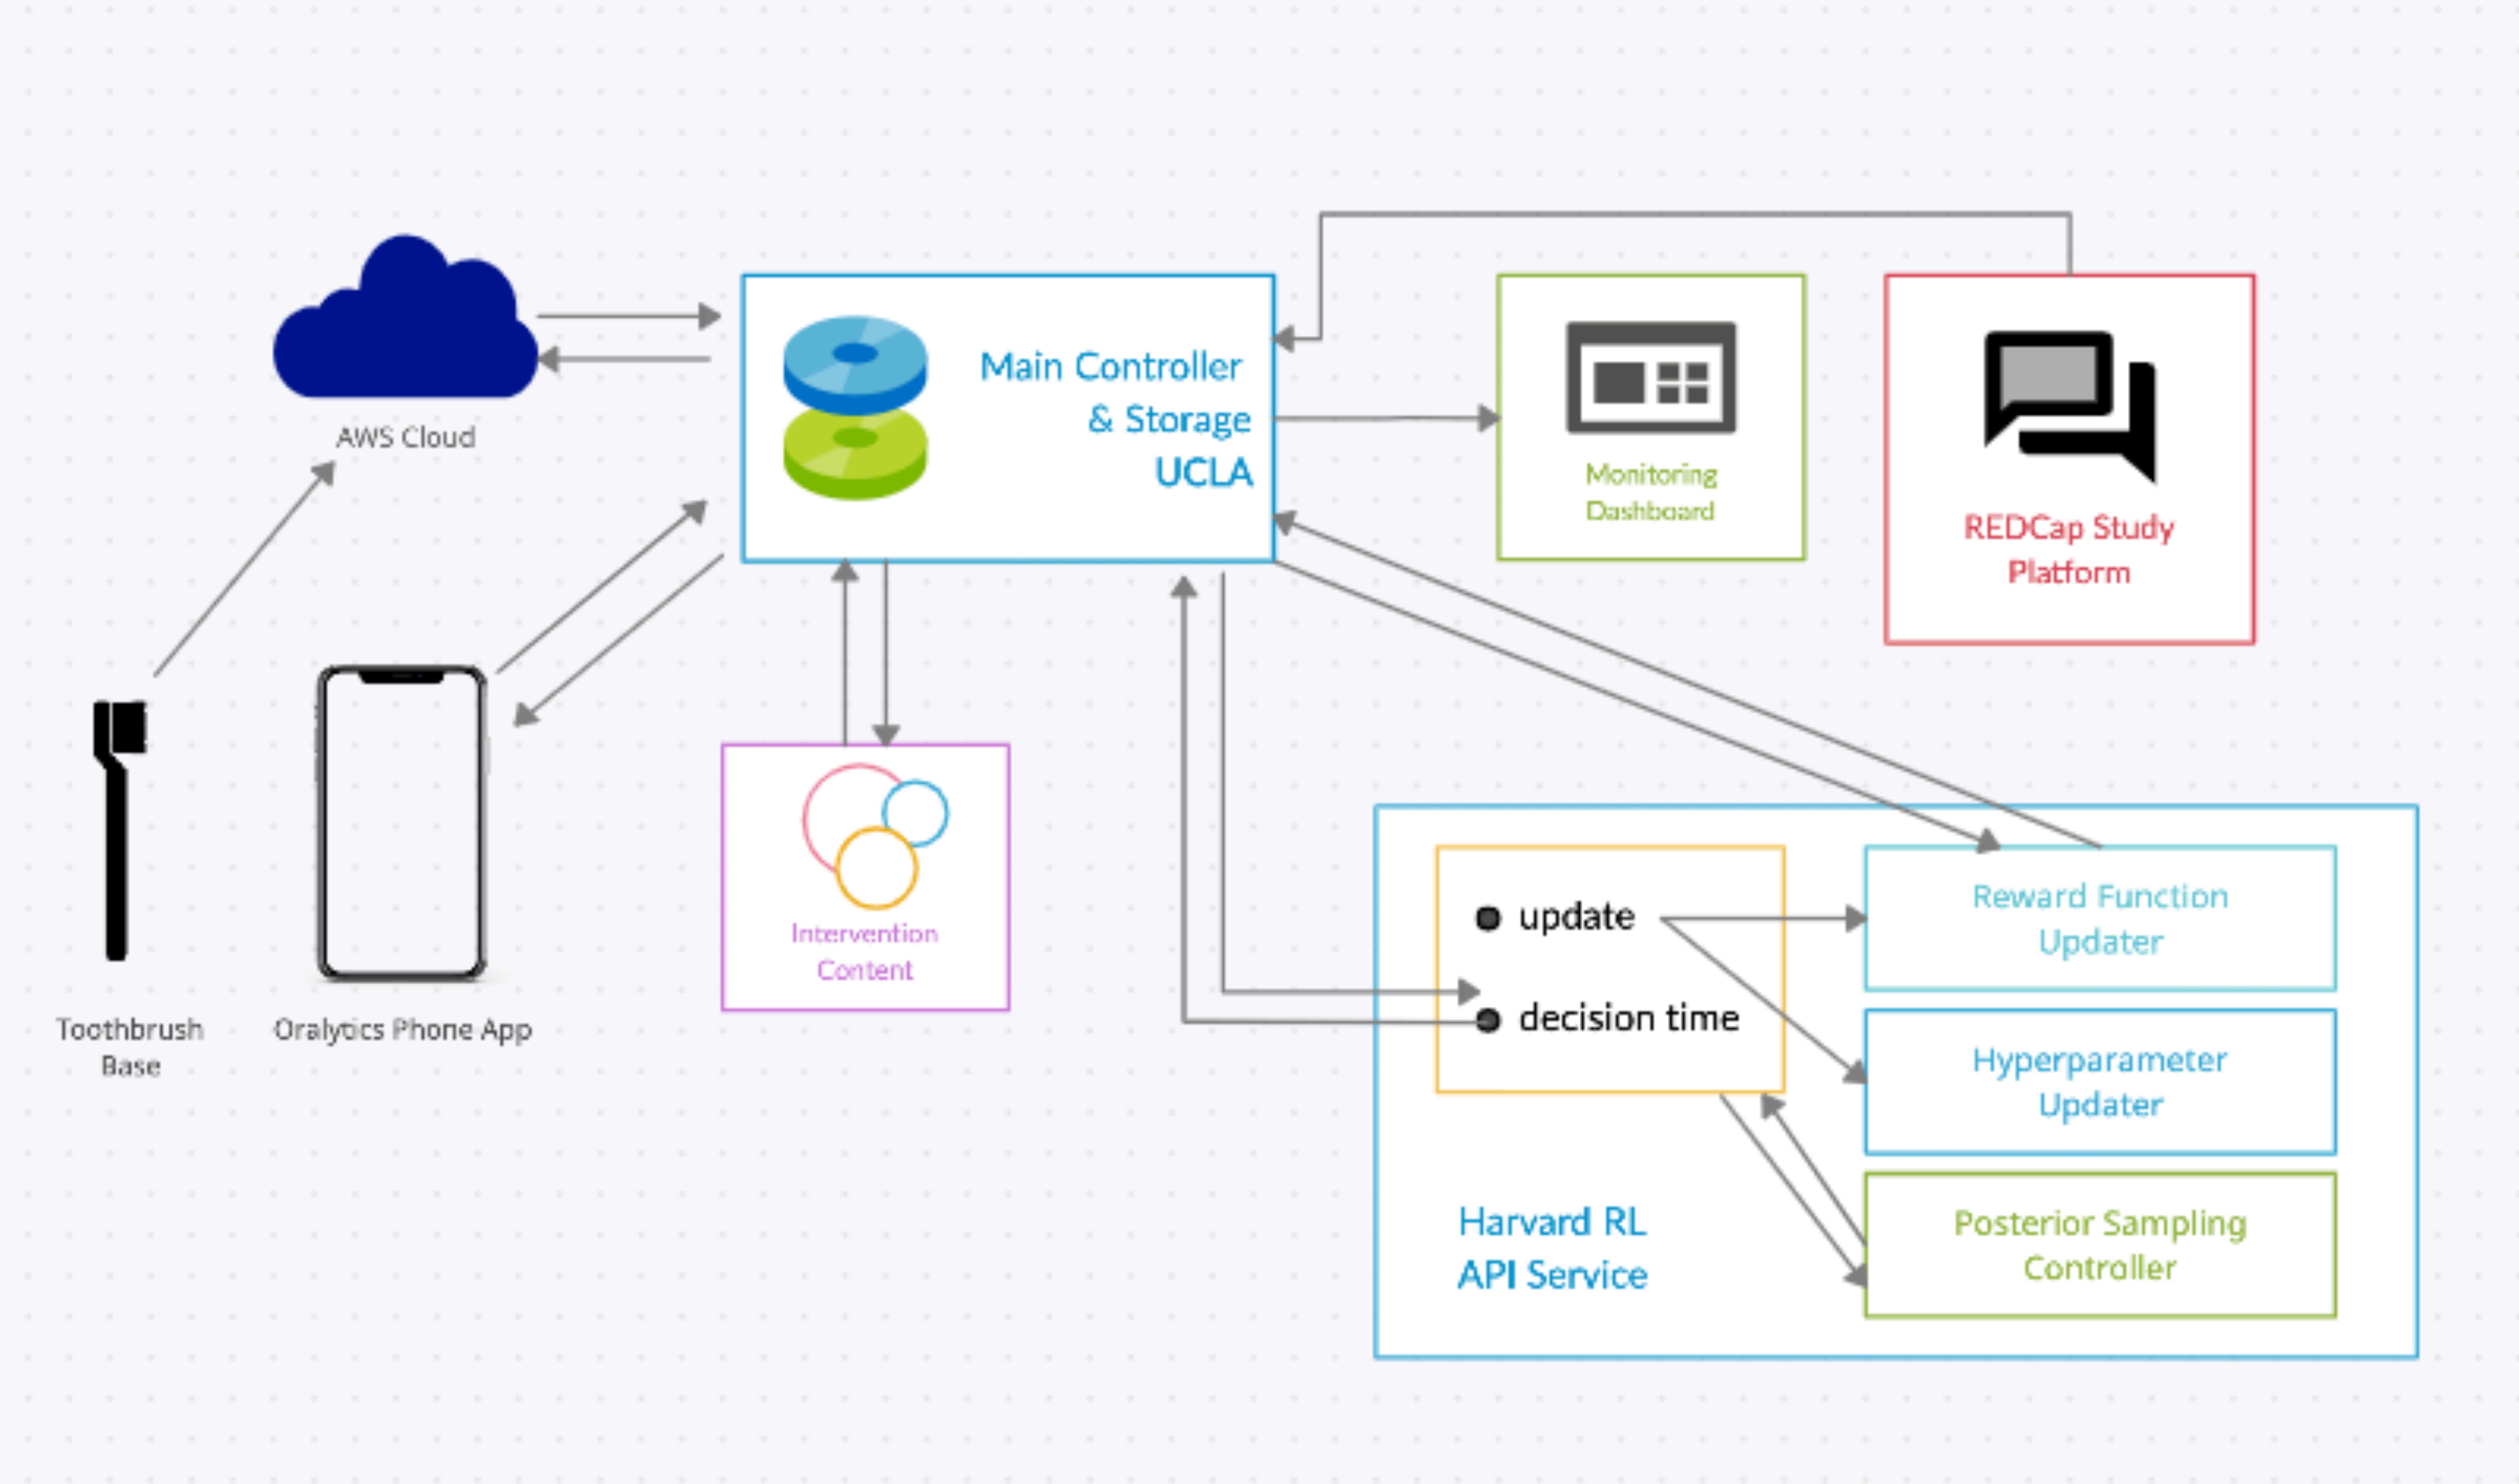
AWS Cloud
Main Controller & Storage UCLA
Monitoring Dashboard
REDCap Study Platform
Toothbrush Base
Oralytics Phone App
Intervention Content
update
decision time
Reward Function Updater
Hyperparameter Updater
Posterior Sampling Controller
Harvard RL API Service

Outline

1. Introduction to Oralytics

2. Just-In-Time Adaptive Interventions and Trial Overview

3. Design of Personalization / Reinforcement Learning Algorithm

Just-In-Time Adaptive Interventions (JITAIs)

- Interventions that are delivered **whenever and wherever needed**

Just-In-Time Adaptive Interventions (JITAIs)

- Interventions that are delivered **whenever and wherever needed**
- Motivation for JITAIs

    (1) Individuals may need support when it is difficult or expensive to provide otherwise

Just-In-Time Adaptive Interventions (JITAIs)

- Interventions that are delivered **whenever and wherever needed**
- Motivation for JITAIs

    (1) Individuals may need support when it is difficult or expensive to provide otherwise

    (2) Individuals are not always aware of when they need support

Just-In-Time Adaptive Interventions (JITAIs)

- Interventions that are delivered **whenever and wherever needed**
- Motivation for JITAIs

  (1) Individuals may need support when it is difficult or expensive to provide otherwise

  (2) Individuals are not always aware of when they need support

  (3) Intervention options may have negative side effects (burden, habituation)

Components of JITAIs

Proximal (Near-Time) and Distal Outcomes

Decision Times

Tailoring / Contextual Variables

Intervention Options / Actions

Decision Rules

Components of JITAIs

Proximal (Near-Time) and Distal Outcomes

Decision Times

Tailoring / Contextual Variables

Intervention Options / Actions

Decision Rules

Distal Outcome:
User's oral health

Proximal (Near-Time) Outcome:
Brushing time at next brushing window

Components of JITAIs

Proximal (Near-Time) and Distal Outcomes

Decision Times

Tailoring / Contextual Variables

Intervention Options / Actions

Decision Rules

Every Morning
Before Brush Time

Every Evening
Before Brush Time

Components of JITAIs

Proximal (Near-Time) and Distal Outcomes

Decision Times

Tailoring / Contextual Variables

Intervention Options / Actions

Decision Rules

Time of Day

Previous Days Brush Times

User Engagement with Oralytics App

Components of JITAIs

Proximal (Near-Time) and Distal Outcomes	No Message
Decision Times	Standard Reciprocity Prompt
Tailoring / Contextual Variables	Reciprocity by Proxy Prompt
Intervention Options / Actions	Q&A Prompt (Morning)
Decision Rules	Personalized Feedback (Evening)

Components of JITAIs

Proximal (Near-Time) and Distal Outcomes

Decision Times

Tailoring / Contextual Variables

Intervention Options / Actions

Decision Rules

Personalization / Reinforcement Learning Algorithm

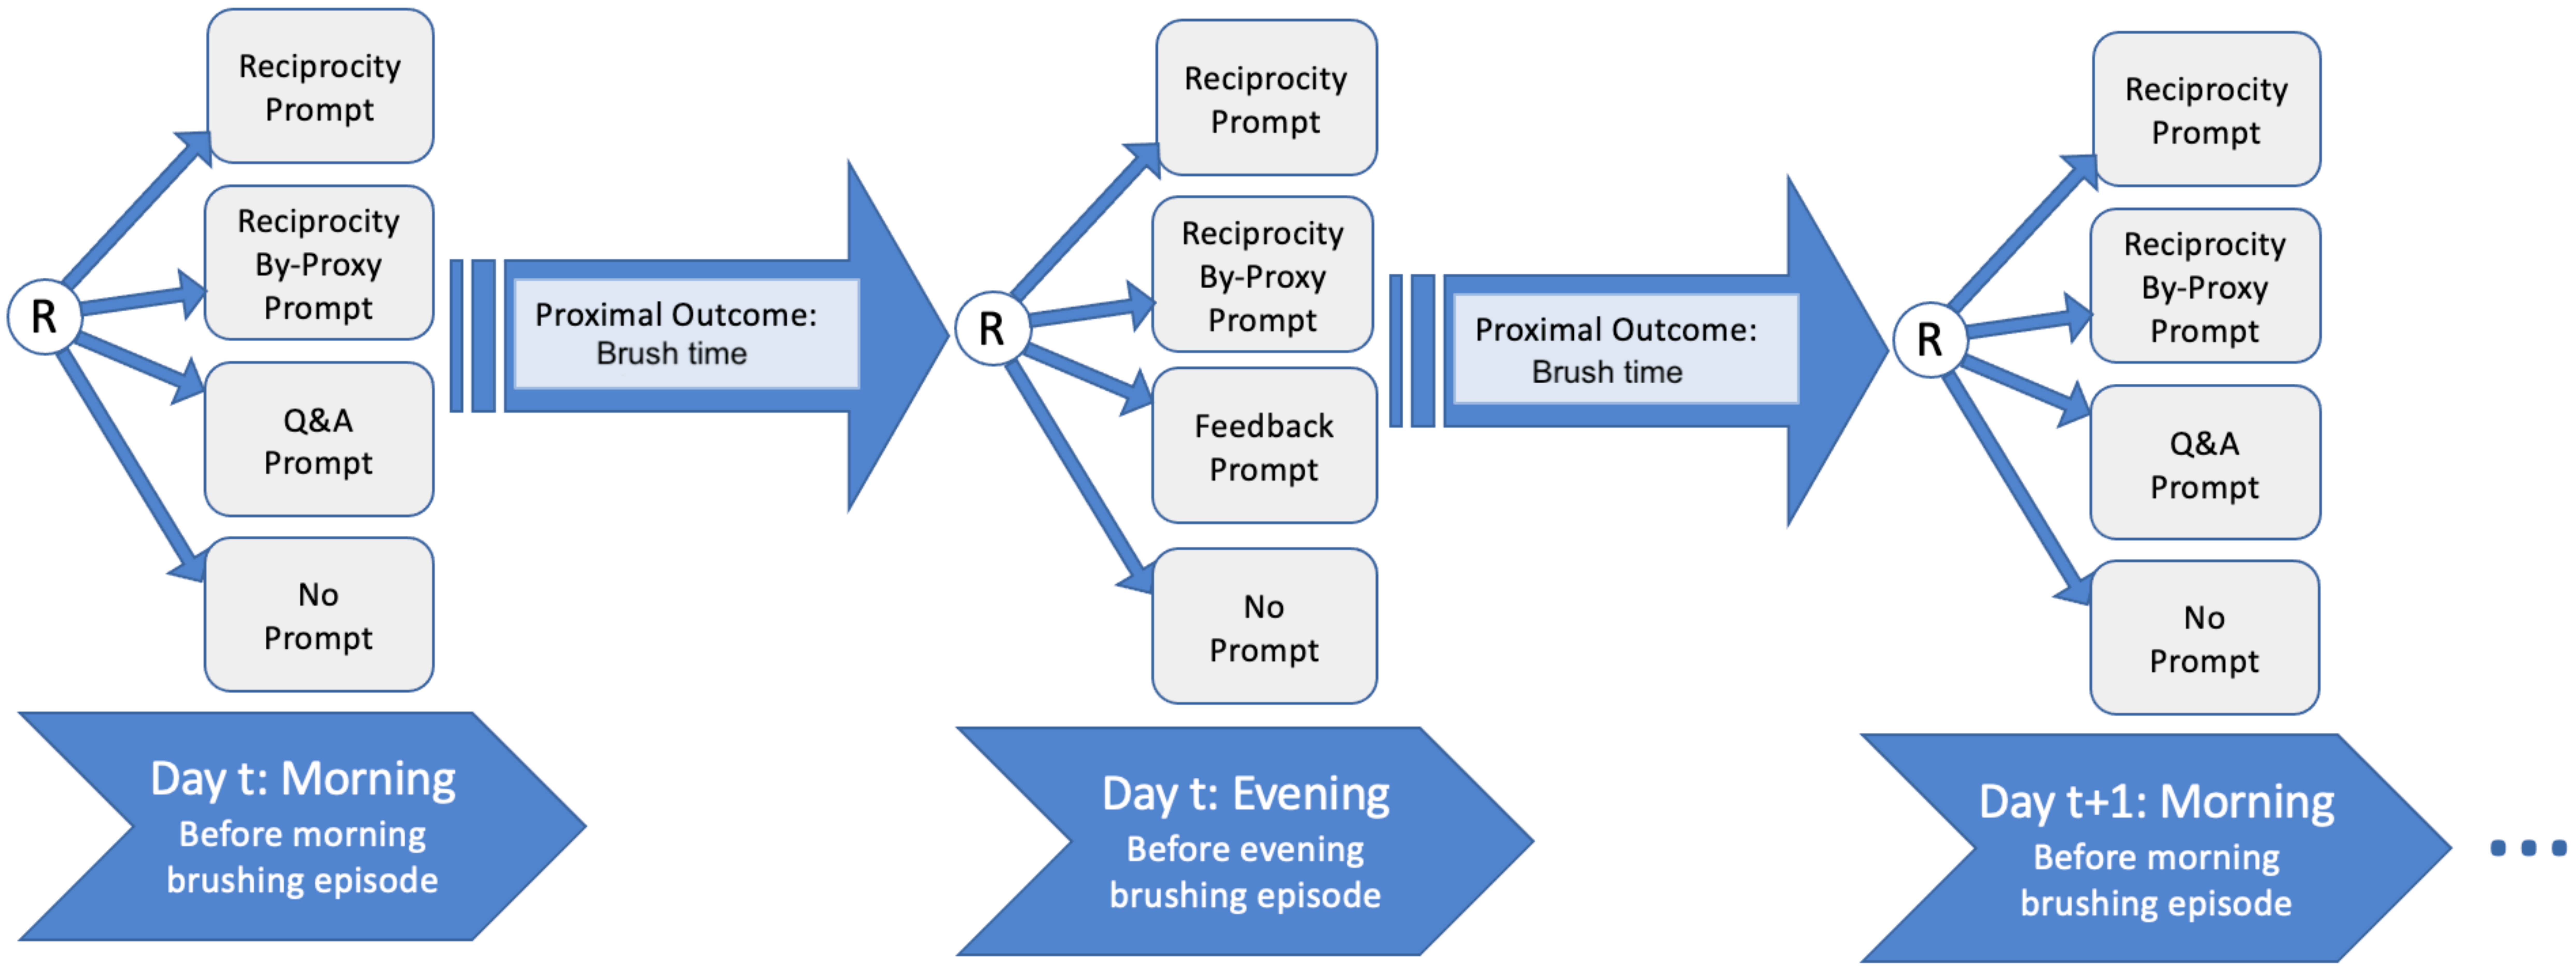

Micro-Randomizations: every day over 10 weeks, day $t = 1, \ldots, 70$

Overview of Study

- **Micro-Randomized Trial (MRT)**
    - 10 week trial with ~75 users
    - Designed to (1) inform design of a JITAI for future use and (2) personalize message delivery (learning for each user)

Overview of Study

- **Micro-Randomized Trial (MRT)**
  - 10 week trial with ~75 users
  - Designed to (1) inform design of a JITAI for future use and (2) personalize message delivery (learning for each user)

- **Standard Clinical Trial will be run after the MRT**
  - Randomize between usual care versus the Digital Oral Health Intervention JITAI
  - Participants recruited from two large “safety net” dental clinical in Los Angeles
  - Evaluate effect on
    (1) Dental plaque
    (2) Gum inflammation

Outline

1. Introduction to Oralytics

2. Just-In-Time Adaptive Interventions and Trial Overview

3. Design of Personalization / Reinforcement Learning Algorithm

Why use an RL algorithm in the MRT?

- Biases micro-randomization probabilities based on the data collected so far

Why use an RL algorithm in the MRT?

- Biases micro-randomization probabilities based on the data collected so far
- Increase probability of sending messages in **contexts for which the data indicates sending a message will increase user brushing**
    - More likely for the intervention to be effective on average

Why use an RL algorithm in the MRT?

- Biases micro-randomization probabilities based on the data collected so far
- Increase probability of sending messages in **contexts for which the data indicates sending a message will increase user brushing**
  - More likely for the intervention to be effective on average
- Decrease the probability of sending messages in contexts the data indicates messages will be less effective
  - Avoid burdening / annoying the user

Approach to Designing the RL Algorithm

- **Develop a simulation environment using**
    - Available prior data
    - Domain expertise

Simulation Environment

- Generates contexts and rewards under each action
- Heterogenous users

Approach to Designing the RL Algorithm

- **Develop a simulation environment using**
  - Available prior data
  - Domain expertise
- **Evaluate RL Algorithms on Simulation Environment**
  - Select algorithm that performs well across simulation variants

Simulation Environment

- Generates contexts and rewards under each action
- Heterogenous users

RL Algorithm

Uses data accrued so far to form action selection probabilities for the current context

ROBAS Dataset

- **Study goal:** Assessing the feasibility of passively tracking brushing behaviors in real world settings
    - 32 participants for 1 month

ROBAS paper: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7380983/

ROBAS Dataset

- **Study goal:** Assessing the feasibility of passively tracking brushing behaviors in real world settings
    - 32 participants for 1 month
- Used a similar sensor suite as will be used in our MRT

ROBAS paper: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7380983/

ROBAS Dataset

- **Study goal:** Assessing the feasibility of passively tracking brushing behaviors in real world settings
    - 32 participants for 1 month
- Used a similar sensor suite as will be used in our MRT
- Differences from our MRT
    - **No intervention**
    - App in ROBAS study only used for data collection

ROBAS paper: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7380983/

ROBAS Dataset

- **Study goal:** Assessing the feasibility of passively tracking brushing behaviors in real world settings
    - 32 participants for 1 month
- Used a similar sensor suite as will be used in our MRT
- Differences from our MRT
    - **No intervention**
    - App in ROBAS study only used for data collection
- **Use data to model outcomes under no intervention**

ROBAS paper: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7380983/

Observation 1: Zero-Inflated Brush Time Outcomes

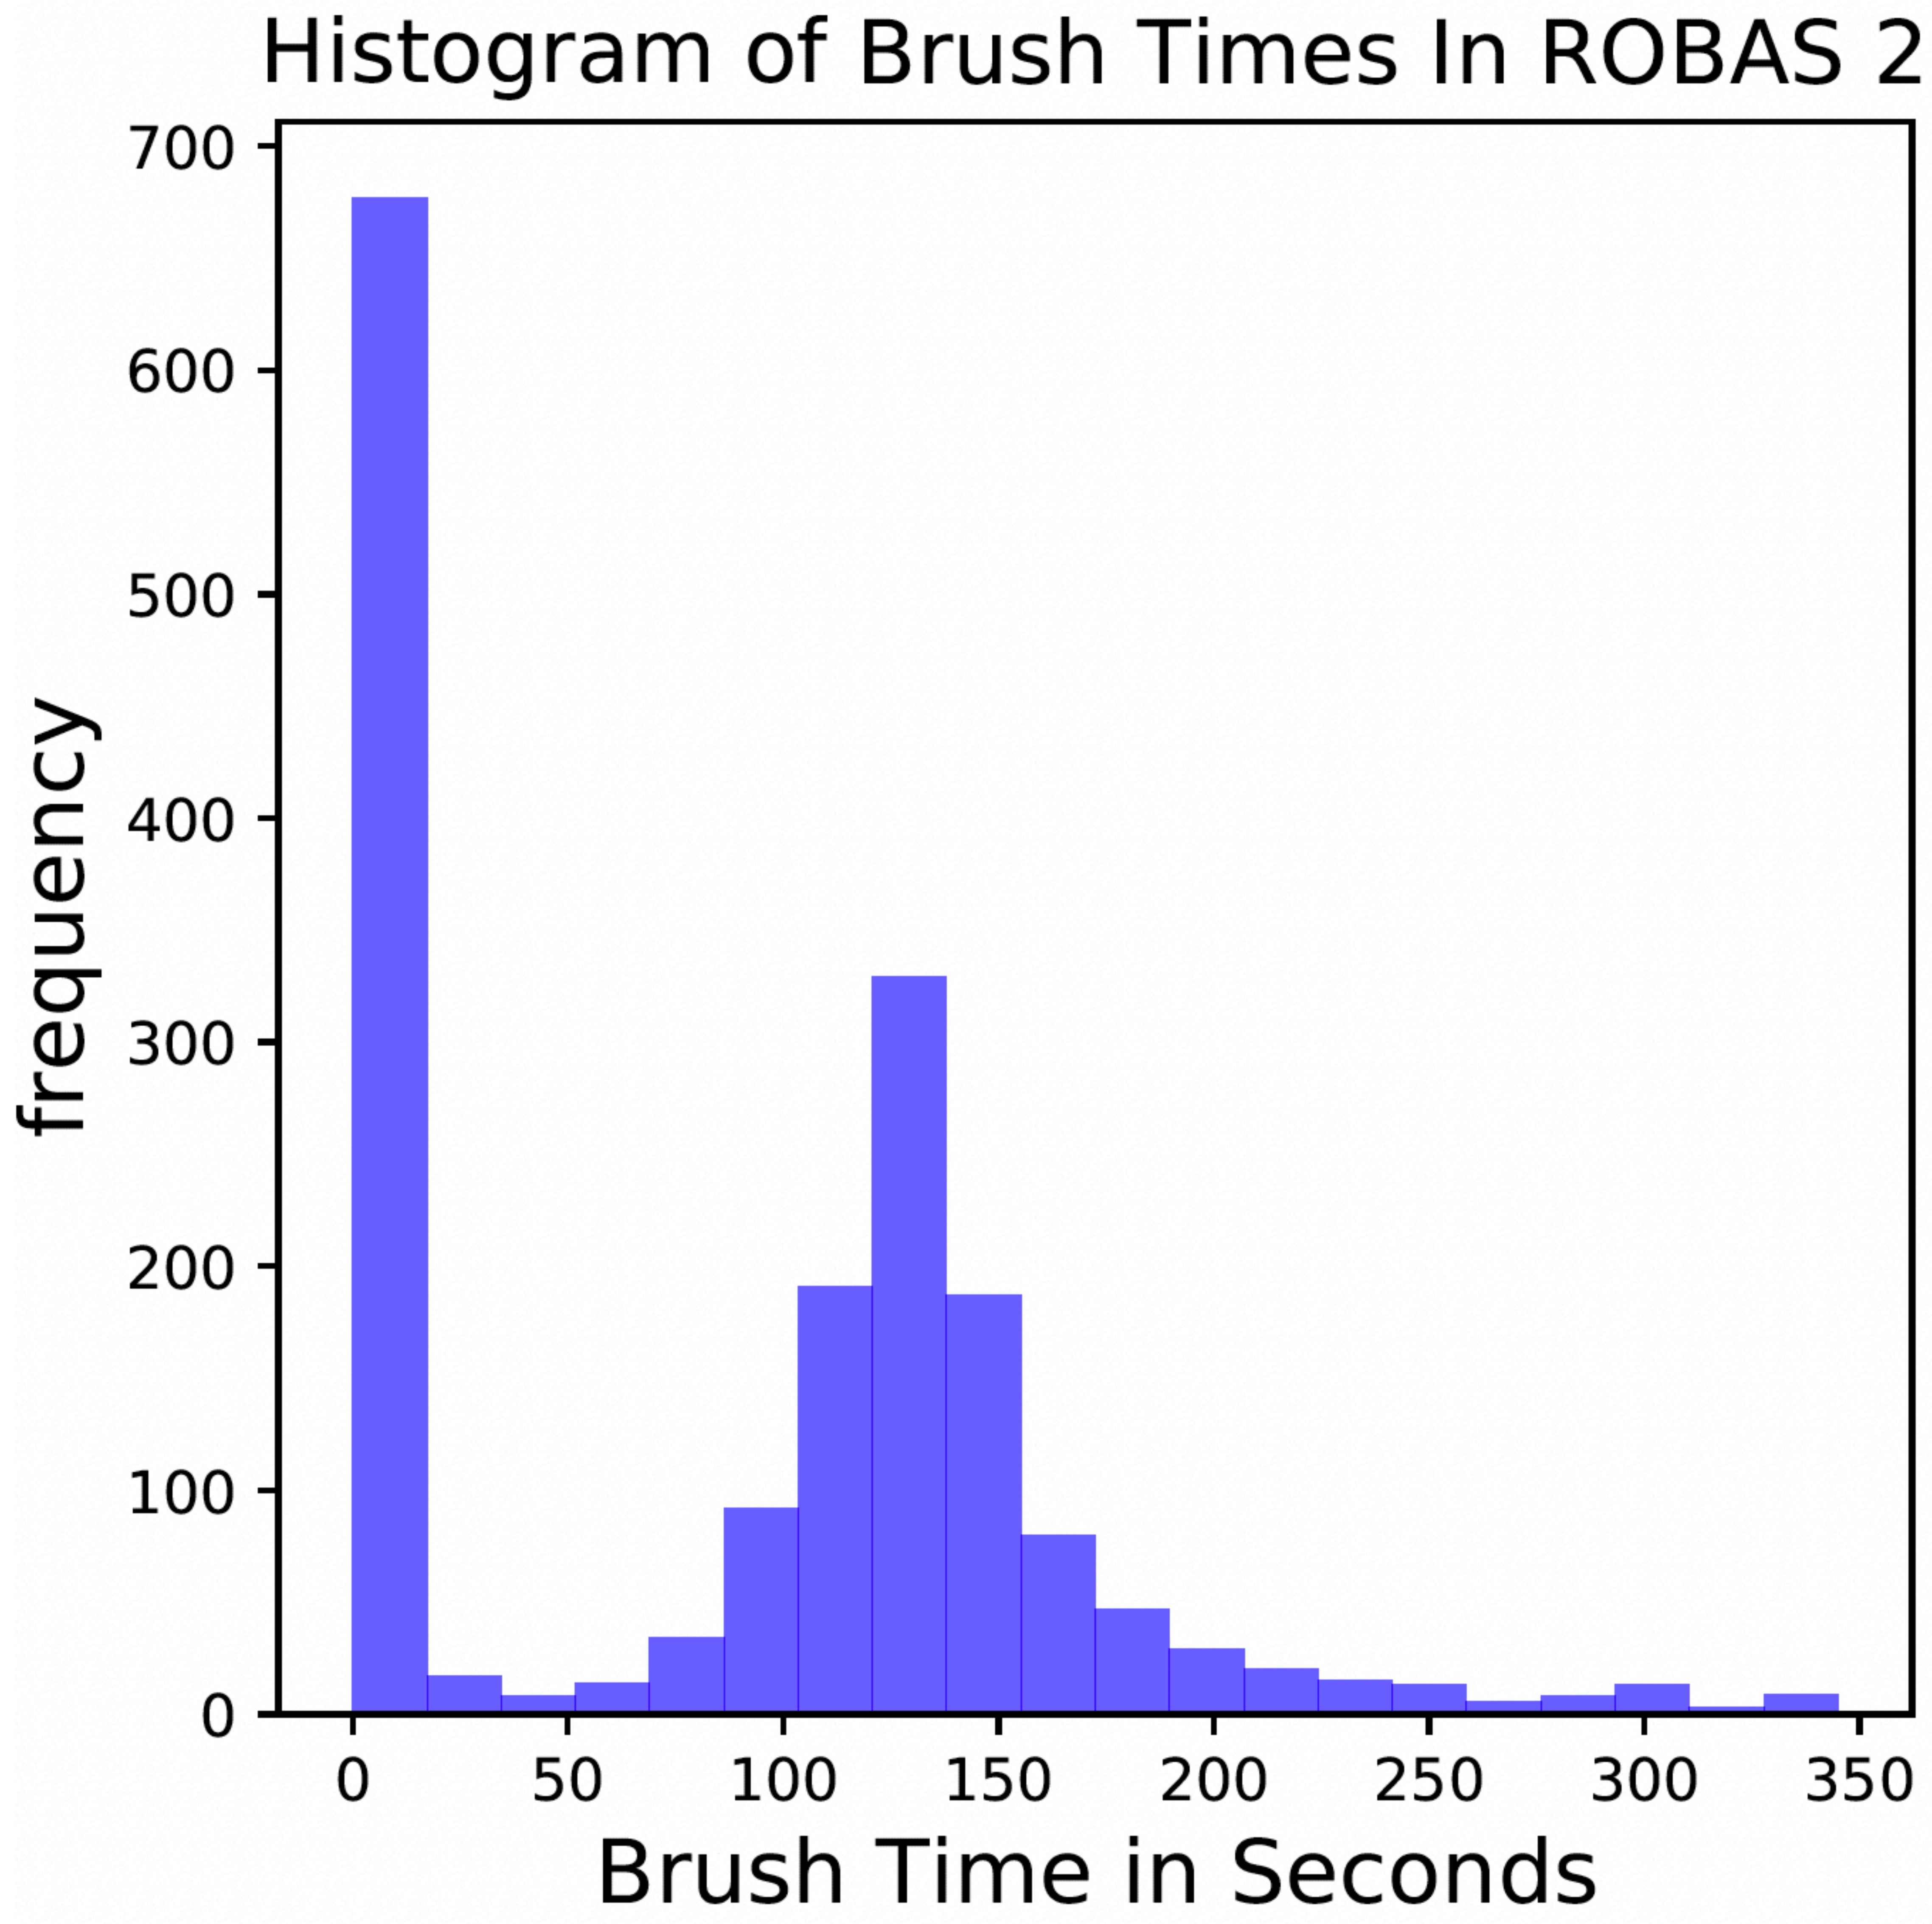

Observation 2: User Heterogeneity

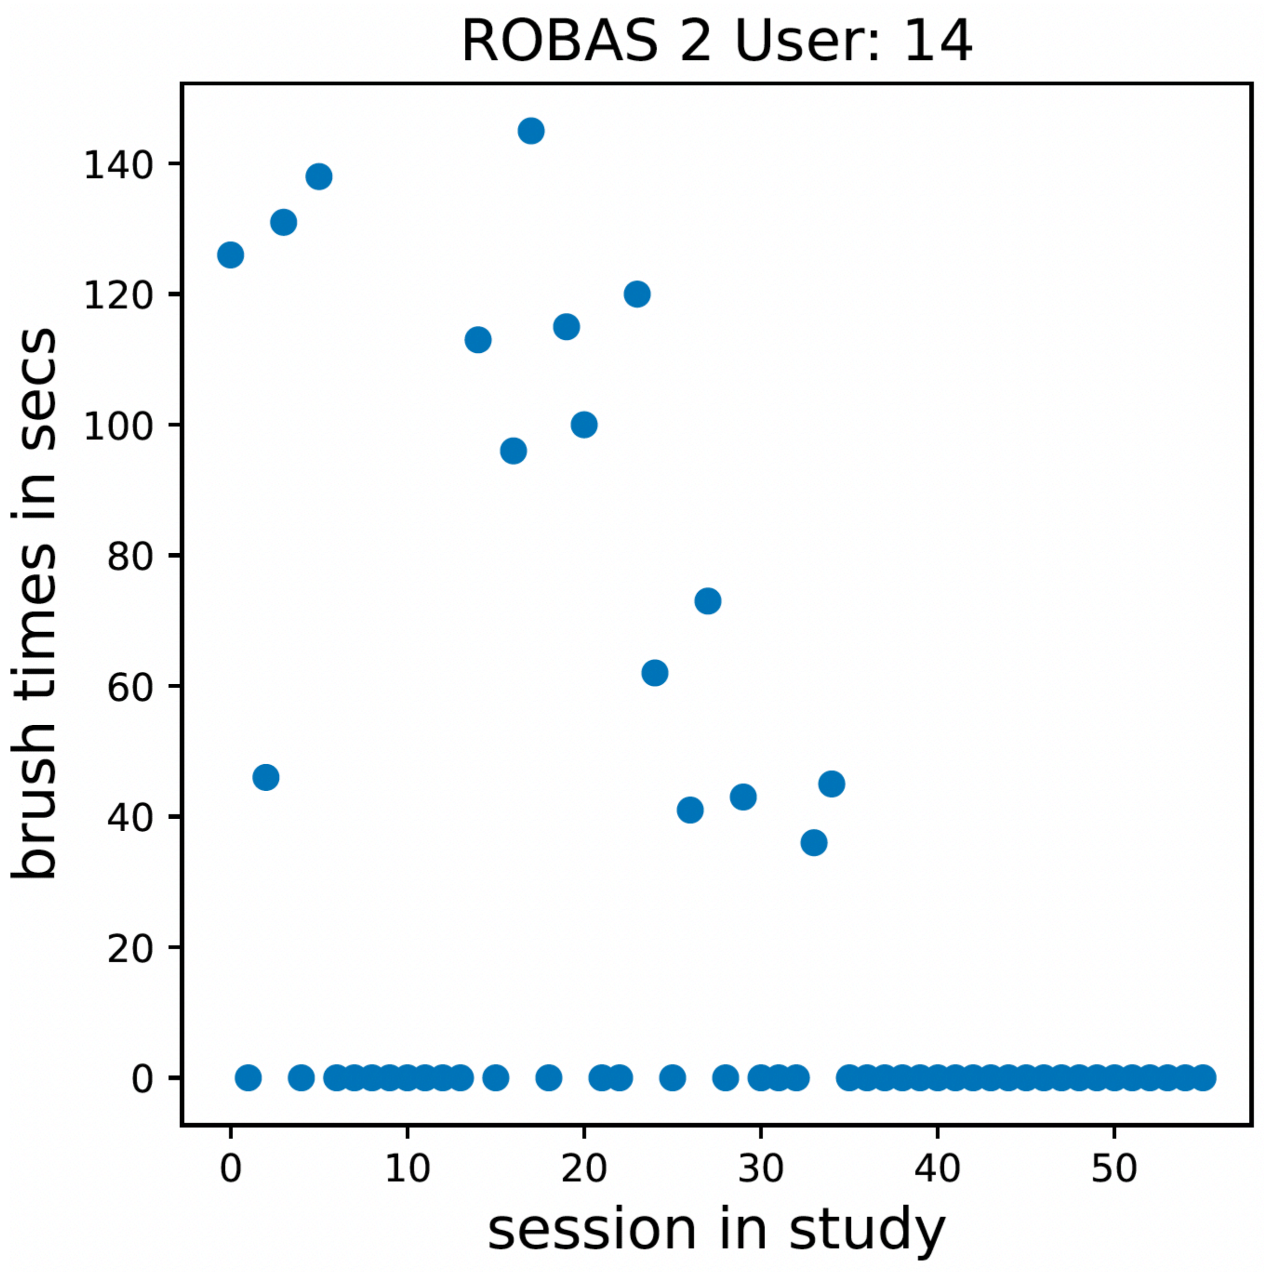

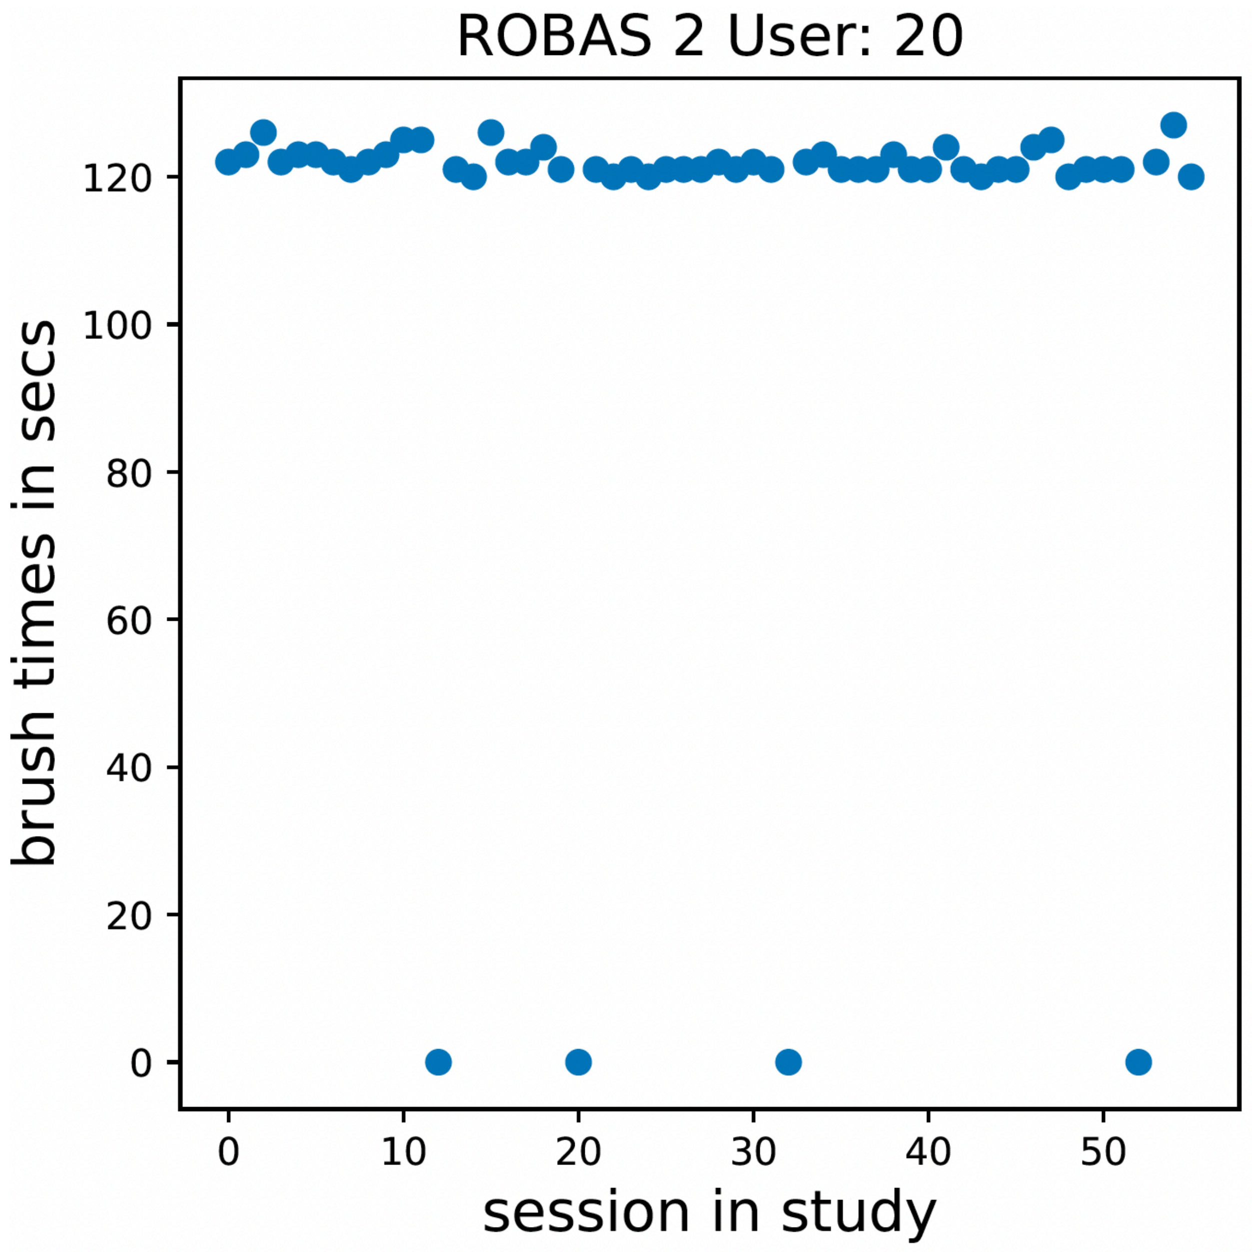

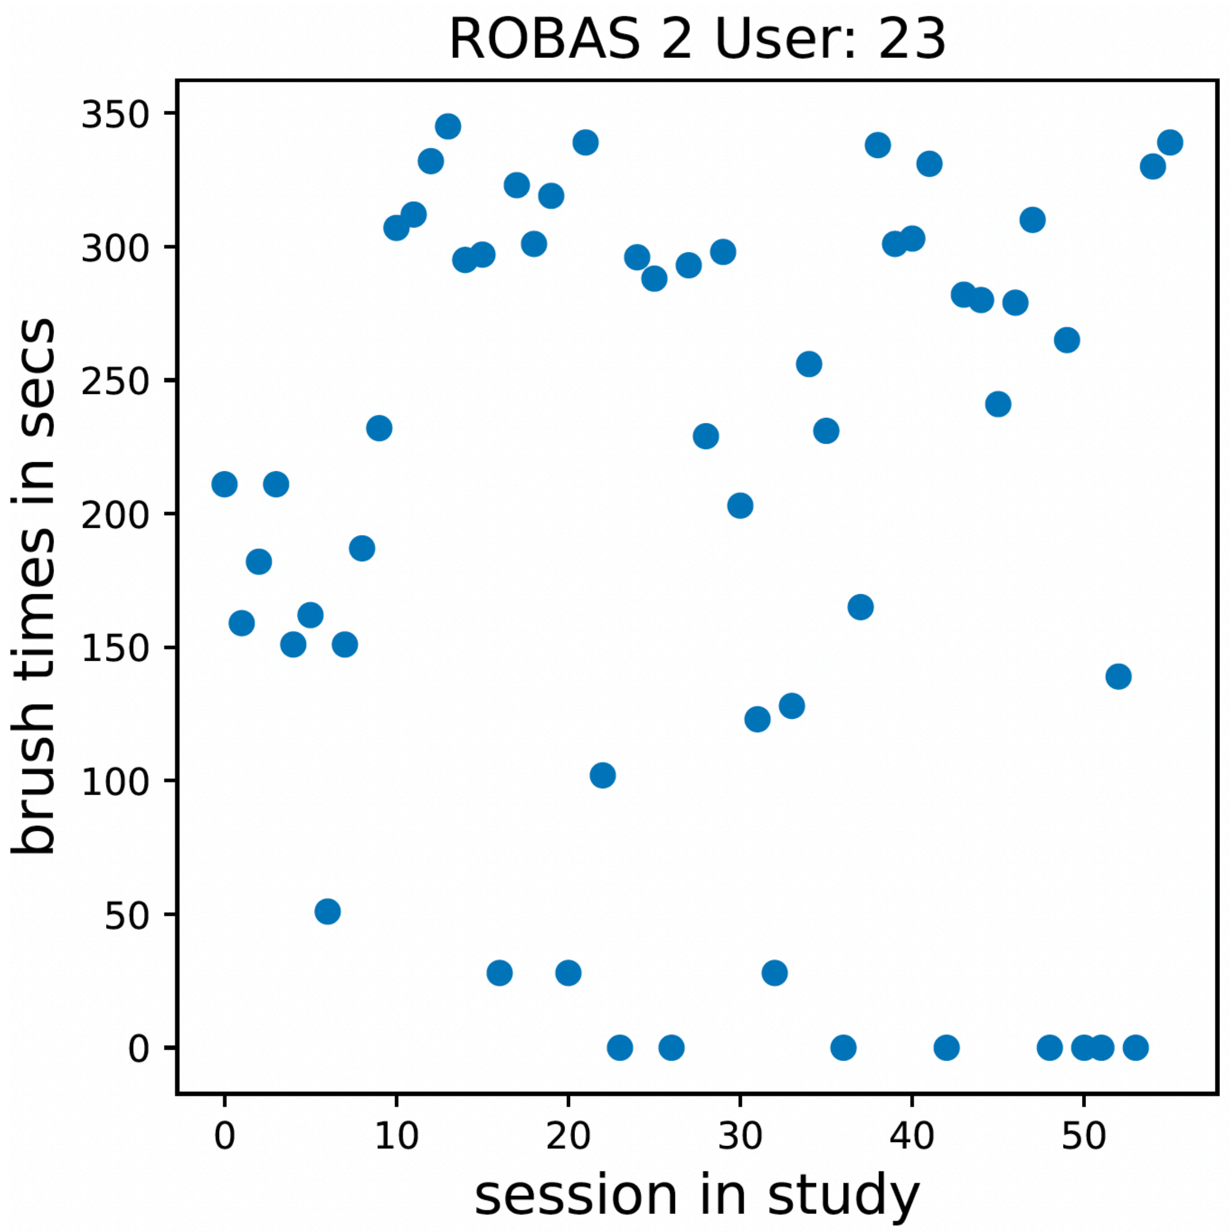

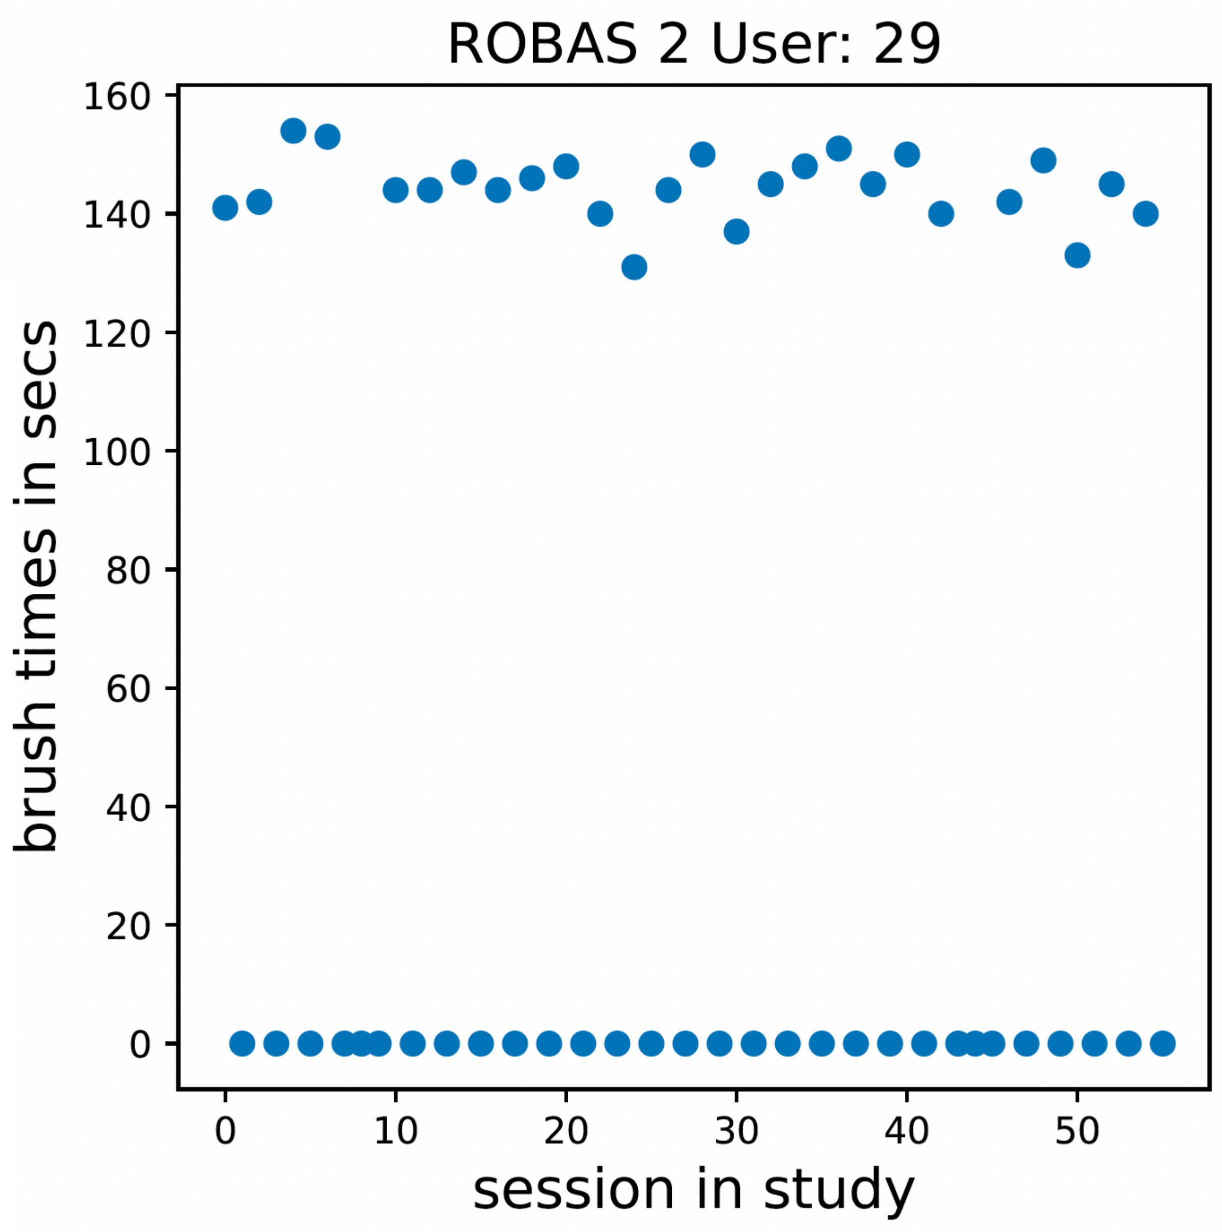

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome

Forming Treatment Effects

- Use ROBAS data to model **brush time in seconds under no message**

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome | **Forming Treatment Effects**

- Use ROBAS data to model **brush time in seconds under no message**
- **Zero-Inflated Poisson model**
    - Model **whether user intends to brush** as a **binary outcome**
    - Model **brush time in seconds** when user intends to brush as a **Poisson outcome**

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome | **Forming Treatment Effects**

- Use ROBAS data to model **brush time in seconds under no message**
- **Zero-Inflated Poisson model**
    - Model **whether user intends to brush** as a **binary outcome**
    - Model **brush time in seconds** when user intends to brush as a **Poisson outcome**
- One model per user (32 total) to accommodate user heterogeneity

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome

Forming Treatment Effects

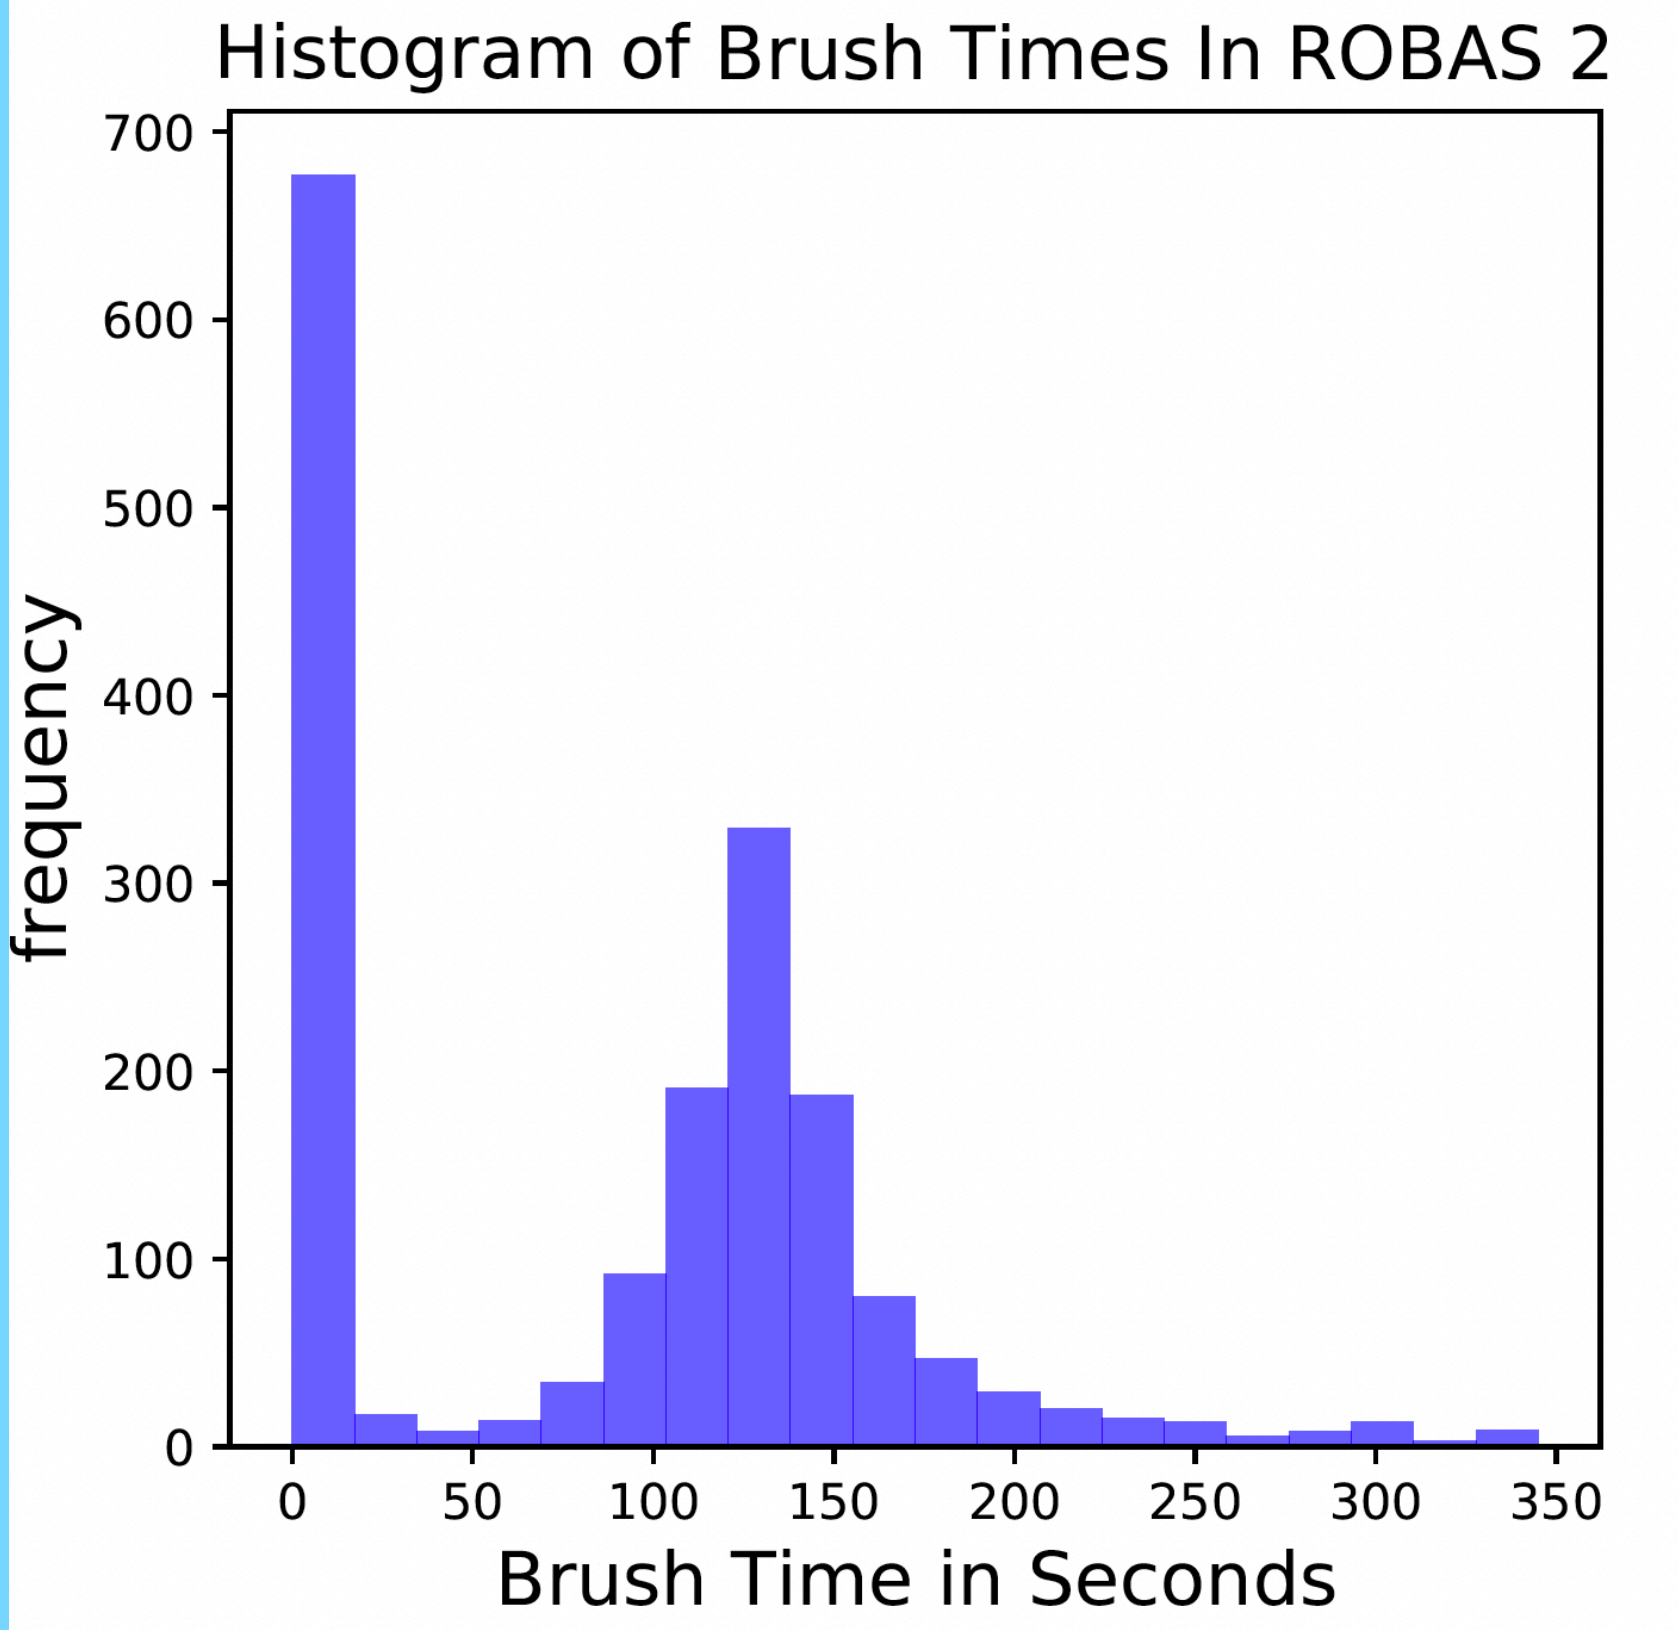

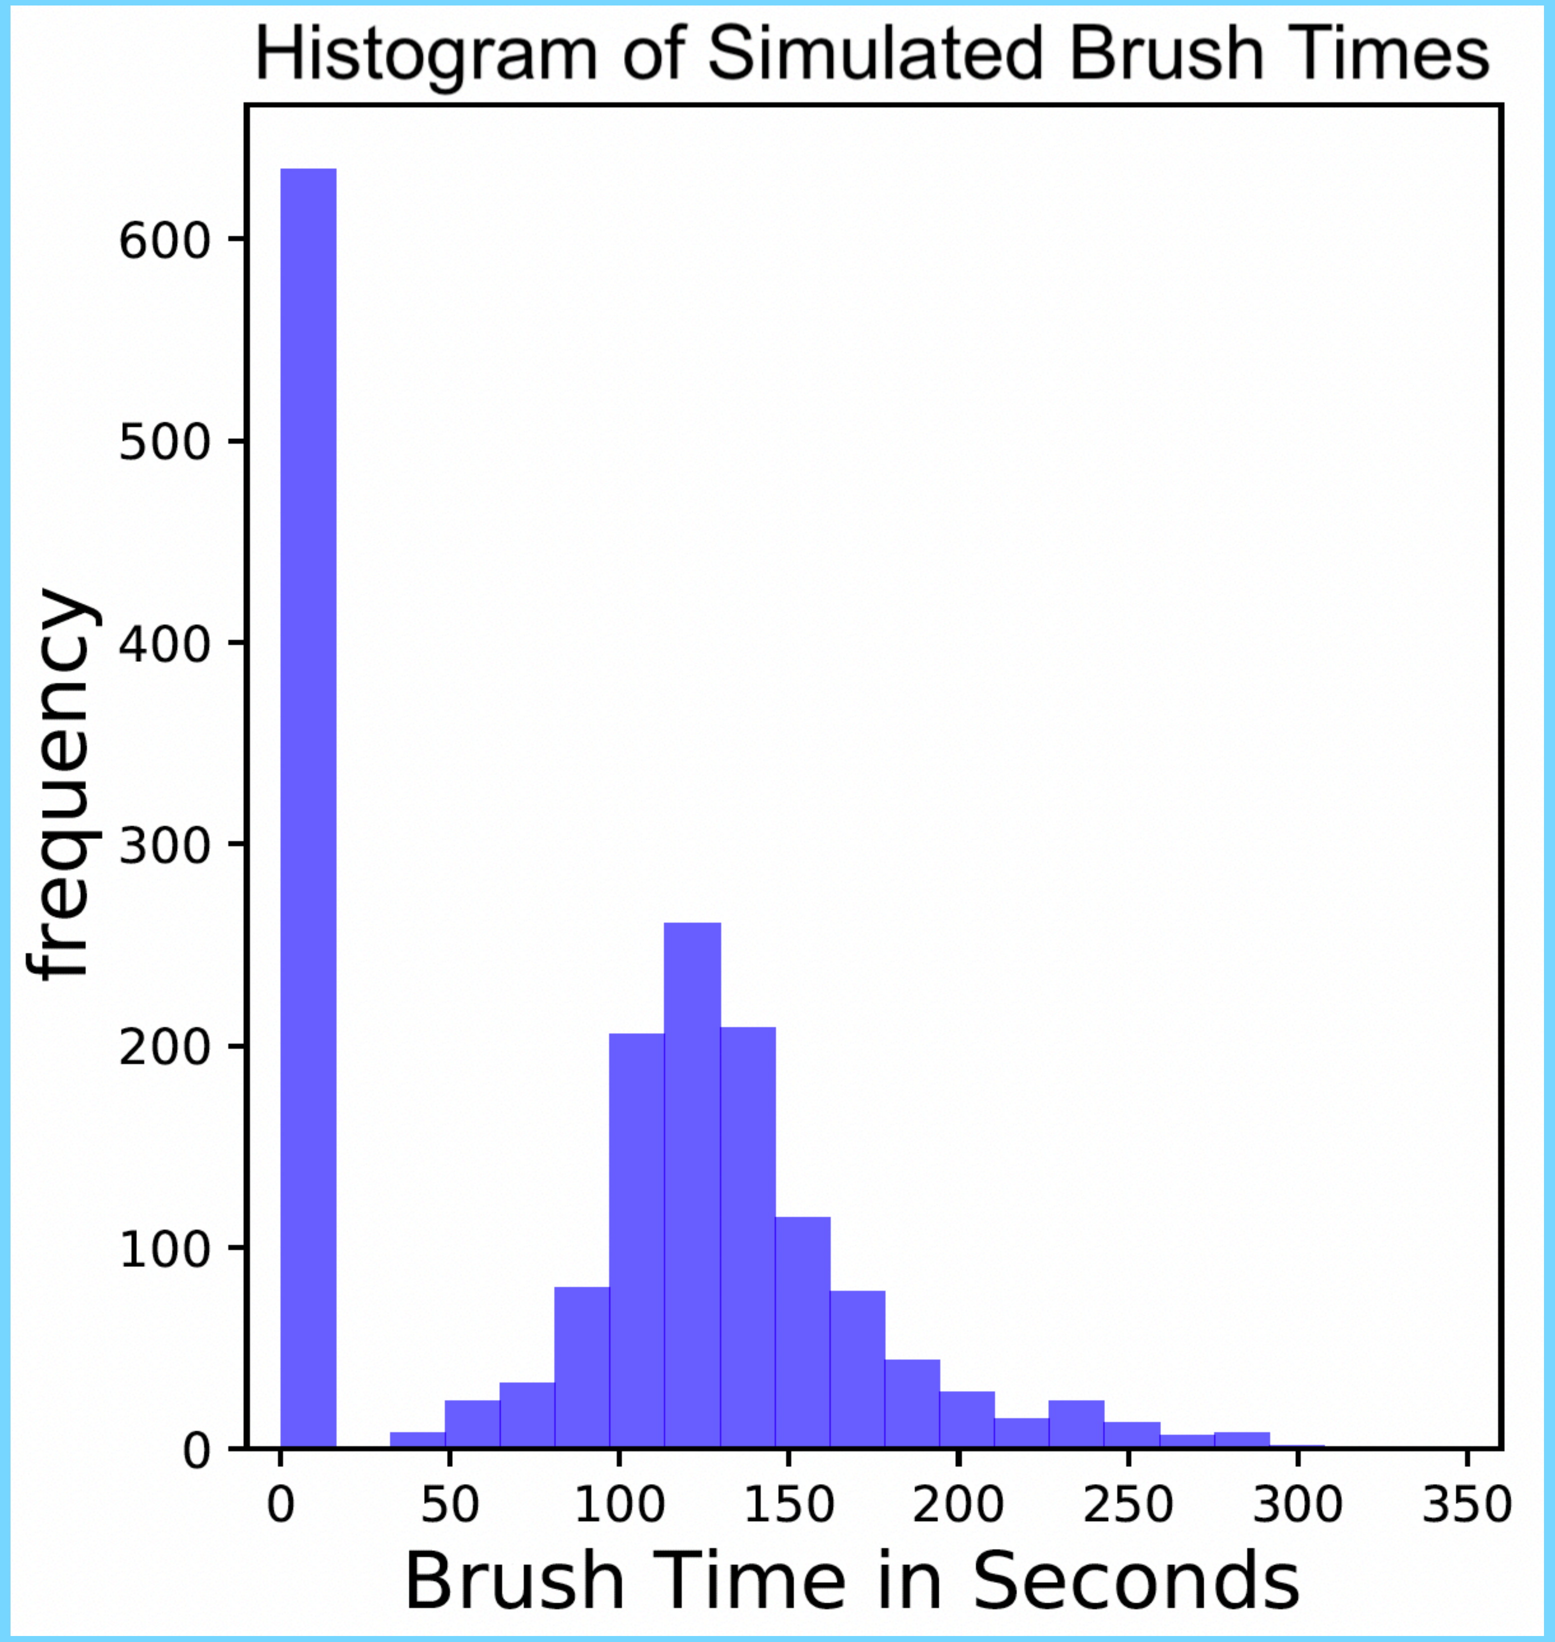

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome

Forming Treatment Effects

```
Count model coefficients (poisson with log link):
                                         Estimate Std. Error
(Intercept)                              4.913283   0.006276
Time.of.Day                              0.081599   0.005468
Prior.Day.Brush.Time.norm                0.138481   0.004236
Proportion.Brushed.In.Past.7.Days       -0.128480   0.009166
Day.in.Study.norm                        0.026660   0.005710
Time.of.Day:Prior.Day.Brush.Time.norm   -0.036234   0.005453

Zero-inflation model coefficients (binomial with logit link):
                                    Estimate Std. Error
(Intercept)                          0.17941    0.10494
Prior.Day.Brush.Time.norm           -0.63728    0.05852
Proportion.Brushed.In.Past.7.Days   -1.65646    0.20382
Day.in.Study.norm                    0.91647    0.11126
```

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome

Forming Treatment Effects

- Need to specify a model for the **brush time in seconds when a message is sent**

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome

Forming Treatment Effects

- Need to specify a model for the **brush time in seconds when a message is sent**
- Construct **plausible treatment effects** with domain scientists
    - Want to ensure our algorithm performs well in such settings

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome

Forming Treatment Effects

- Need to specify a model for the **brush time in seconds when a message is sent**
- Construct **plausible treatment effects** with domain scientists
  - Want to ensure our algorithm performs well in such settings
- Conjecture: Messages will **increase the probability of intending to brush** rather than **increase brush times when intending to brush**

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome

Forming Treatment Effects

- Need to specify a model for the **brush time in seconds when a message is sent**
- Construct **plausible treatment effects** with domain scientists
  - Want to ensure our algorithm performs well in such settings
- Conjecture: Messages will **increase the probability of intending to brush** rather than **increase brush times when intending to brush**
- How to construct treatment effects? Context-dependent effects?

Decisions: The Simulation Environment

Modeling Zero-Inflated Outcome

Forming Treatment Effects

- Our Current Approach: **Population Treatment Effect**
  - Logistic regression model for whether the user brushes
- Our Current Approach: **User-Specific Treatment Effect**
  - Sample the treatment effect size for each user from $\mathcal{N}(0.2, \sigma^2)$

Two Components of RL Algorithms

Fit Statistical Model

Action Selection Strategy

Two Components of RL Algorithms

Fit Statistical Model

Fit model for expected reward conditional on context and action using all data collected so far (actions, context, rewards)

Action Selection Strategy

Two Components of RL Algorithms

Fit Statistical Model

Fit model for expected reward conditional on context and action using all data collected so far (actions, context, rewards)

Action Selection Strategy

Use most recently **fitted statistical model** and **current context information** to form **micro-randomization probabilities**

Two Components of RL Algorithms

Fit Statistical Model

Fit model for expected reward conditional on context and action using all data collected so far (actions, context, rewards)

Action Selection Strategy

Use most recently **fitted statistical model** and **current context information** to form **micro-randomization probabilities**

Updated every night

Requires all newly collected data

Two Components of RL Algorithms

Fit Statistical Model

<u>Fit model for expected reward conditional on context and action</u> using all data collected so far (actions, context, rewards)

Action Selection Strategy

Use most recently **<u>fitted statistical model</u>** and **<u>current context information</u>** to form **<u>micro-randomization probabilities</u>**

Updated every night
Requires all newly collected data

Make Decisions every morning
Needs yesterday's outcome data

Two Components of RL Algorithms

Fit Statistical Model

<u>Fit model for expected reward conditional on context and action</u> using all data collected so far (actions, context, rewards)

Action Selection Strategy

Use most recently **<u>fitted statistical model</u>** and **<u>current context information</u>** to form **<u>micro-randomization probabilities</u>**

Bayesian Model

"Posterior Sampling"

Posterior Sampling Algorithm Example

- Consider one user with decision times
  $t \in [1 : 70]$

Posterior Sampling Algorithm Example

- Consider one user with decision times $t \in [1:70]$
- $A_t \in \{0,1\}$ binary action
    - 0 = don't send a message
    - 1 = send a message

Posterior Sampling Algorithm Example

- Consider one user with decision times $t \in [1:70]$
- $A_t \in \{0,1\}$ binary action
  - 0 = don't send a message
  - 1 = send a message
- X_t context vector
  - Intercept, time of day, previous day brush time, Oralytics app engagement

Posterior Sampling Algorithm Example

- Consider one user with decision times $t \in [1:70]$
- $A_t \in \{0,1\}$ binary action
  - 0 = don't send a message
  - 1 = send a message
- X_t context vector
  - Intercept, time of day, previous day brush time, Oralytics app engagement
- R_t brush time in seconds

Posterior Sampling Algorithm Example

- Consider one user with decision times $t \in [1:70]$
- $A_t \in \{0,1\}$ binary action
    - 0 = don't send a message
    - 1 = send a message
- X_t context vector
    - Intercept, time of day, previous day brush time, Oralytics app engagement
- R_t brush time in seconds
- $H_{t-1} = \left\{X_t, A_t, R_t\right\}_{s=1}^{t-1}$ all data seen so far

Posterior Sampling Algorithm Example

- Consider one user with decision times $t \in [1:70]$
- $A_t \in \{0,1\}$ binary action
    - 0 = don't send a message
    - 1 = send a message
- X_t context vector
    - Intercept, time of day, previous day brush time, Oralytics app engagement
- R_t brush time in seconds
- $H_{t-1} = \left\{X_t, A_t, R_t\right\}_{s=1}^{t-1}$ all data seen so far

Fit Statistical Model

- $\mathbb{E}[R_t \mid X_t, A_t] = \theta_0^\top X_t + A_t \theta_1^\top X_t$
- **Informative prior** $[\tilde{\theta}_0, \tilde{\theta}_1] \sim \mathcal{N}(\mu, \Sigma)$
- **Model rewards as** $R_t \mid X_t, A_t \sim \mathcal{N}\left(\tilde{\theta}_0^\top X_t + A_t \tilde{\theta}_1^\top X_t, \sigma^2\right)$
- **Update posterior distribution of** $\tilde{\theta}_0, \tilde{\theta}_1$

Posterior Sampling Algorithm Example

- Consider one user with decision times $t \in [1:70]$
- $A_t \in \{0,1\}$ binary action
  - 0 = don't send a message
  - 1 = send a message
- X_t context vector
  - Intercept, time of day, previous day brush time, Oralytics app engagement
- R_t brush time in seconds
- $H_{t-1} = \{X_t, A_t, R_t\}_{s=1}^{t-1}$ all data seen so far

Fit Statistical Model

- $\mathbb{E}[R_t \,|\, X_t, A_t] = \theta_0^\top X_t + A_t \theta_1^\top X_t$
- **Informative prior** $[\tilde{\theta}_0, \tilde{\theta}_1] \sim \mathcal{N}(\mu, \Sigma)$
- **Model rewards as** $R_t \,|\, X_t, A_t \sim \mathcal{N}\left(\tilde{\theta}_0^\top X_t + A_t \tilde{\theta}_1^\top X_t, \sigma^2\right)$
- **Update posterior distribution of** $\tilde{\theta}_0, \tilde{\theta}_1$

Action Selection Strategy

Posterior probability that treatment effect is positive

$$\mathbb{P}(A_t = 1 \,|\, X_t, H_{t-1}) = \mathbb{P}\left(\tilde{\theta}_1{}^\top X_t > 0 \,\middle|\, X_t, H_{t-1}\right)$$

Your Thoughts? Designing the RL Algorithm

Bayesian Statistical Model | **Pooling in Clusters**

(1) Model Zero-Inflated Brush Time Outcome?

- No closed-form posterior for zero-inflated Poisson model

Your Thoughts? Designing the RL Algorithm

Bayesian Statistical Model

Pooling in Clusters

(1) Model Zero-Inflated Brush Time Outcome?

- No closed-form posterior for zero-inflated Poisson model

(2) Model Brush Time as Continuous Outcome?

- Ignores the zero-inflated nature of data

Your Thoughts? Designing the RL Algorithm

Bayesian Statistical Model

Pooling in Clusters

(1) Model Zero-Inflated Brush Time Outcome?

- No closed-form posterior for zero-inflated Poisson model

(2) Model Brush Time as Continuous Outcome?

- Ignores the zero-inflated nature of data

(3) Model Brushing as Binary Outcome?

- Only model binary outcome of whether the user brushed

Your Thoughts? Designing the RL Algorithm

Bayesian Statistical Model

Pooling in Clusters

- The RL algorithm will **pool the data of users in the same cluster together** to fit one statistical model per cluster
    - Increase the speed of RL algorithm learning / reduce noise

Your Thoughts? Designing the RL Algorithm

Bayesian Statistical Model

Pooling in Clusters

- The RL algorithm will **pool the data of users in the same cluster together** to fit one statistical model per cluster
    - Increase the speed of RL algorithm learning / reduce noise
- Planning to cluster users by entry date into the study
    - Incremental recruitment — limited by recruitment rate

Your Thoughts? Designing the RL Algorithm

Bayesian Statistical Model | **Pooling in Clusters**

- The RL algorithm will **pool the data of users in the same cluster together** to fit one statistical model per cluster
    - Increase the speed of RL algorithm learning / reduce noise
- Planning to cluster users by entry date into the study
    - Incremental recruitment — limited by recruitment rate
- How to choose the cluster size?

Thank you!!

Back-Up Slides!!

“Micro-Randomized Trials [are] used to optimize JITAI decision rules, with the ultimate goal of developing an effective and efficient JITAI.” (Qian et al., 2021)

The Micro-Randomized Trial for Developing Digital Interventions: Experimental Design and Data Analysis Considerations: https://arxiv.org/abs/2107.03544

MRT Questions to Inform JITAI Design

- Which time-varying context is it best to send an intervention?

MRT Questions to Inform JITAI Design

- Which time-varying context is it best to send an intervention?
- Which time-varying context is it best **not** to send an intervention?

MRT Questions to Inform JITAI Design

- Which time-varying context is it best to send an intervention?
- Which time-varying context is it best **not** to send an intervention?
- Do the intervention options differentially impact the proximal outcome?

MRT Questions to Inform JITAI Design

- Which time-varying context is it best to send an intervention?
- Which time-varying context is it best **not** to send an intervention?
- Do the intervention options differentially impact the proximal outcome?
- Does the main effect deteriorate with time (day since beginning intervention)?

Oralytics MRT Scientific Questions

- **Primary Analysis:** “On average across time, does delivering messages increase the brushing quality at the next brushing window after the message is delivered, compared to no message?”

Oralytics MRT Scientific Questions

- **Primary Analysis:** “On average across time, does delivering messages increase the brushing quality at the next brushing window after the message is delivered, compared to no message?”
- **Secondary Analysis:** Effect moderation
  - Previous day brush time at same time of day
  - Time of day
  - Engagement binary indicator prior to decision time.
  - Gender
  - Age
  - Ethnicity