A system for collecting and analyzing experience-sampling data

Simon Dennis^{1,2} · Hyungwook Yim^{1,3} · Paul Garrett⁴ · Vishnu Sreekumar⁵ · Ben Stone^{1,2}

Published online: 25 June 2019
© The Psychonomic Society, Inc. 2019

Abstract

Online and sensor technologies promise to transform many areas of psychological enquiry. However, collecting and analyzing such data are challenging. In this article, we introduce the unforgettable.me experience-sampling platform. Unforgettable.me includes an app that can collect image, Global Positioning System, accelerometry, and audio data in a continuous fashion and upload the data to a server. The data are then automatically augmented by using online databases to identify the address, type of location, and weather conditions, as well as provide street view imagery. In addition, machine-learning classifiers are run to identify aspects of the audio data such as voice and traffic. The augmented data are available to participants in the form of a keyword search interface, as well as via several visualization mechanisms. In addition, Unforgettable Research Services partners with If This Then That (IFTTT), and so can accumulate data from any of over 600 sources, including social media, wearables, and other devices. Through IFTTT, buttons can be added as icons to smartphones to allow participants to register mood conveniently, as well as behaviors and physiological states such as happiness, microaggressions, or illness. Furthermore, unforgettable.me incorporates a mechanism that allows researchers to run experiments and analyze data within an authenticated environment without viewing users' private data.

Keywords Experience sampling · Privacy · Data collection · Data analysis

In the history of psychology, there have (arguably) been three major changes in the methodological basis of enquiry. The first of these was to recognize that psychology is a science, and that one can run experiments and have an expectation that they should replicate. This development is often attributed to Wundt, who established the first psychology laboratory in 1879 (Fancher & Rutherford, 2012). The second major development was the introduction of mathematics as a way of characterizing psychological phenomena. The birth of mathematical psychology occurred at about the same time as that of experimental psychology, when researchers such as Weber and Fechner sought to import the methods of physics into the study of human behavior (Fancher & Rutherford, 2012). The third main development was the introduction of brain-

imaging techniques. Although the origins of neuroimaging can be traced to the late 19th century in the work of Angelo Mosso, the methods gained significant momentum about 40 years ago, with the introduction first of positron emission tomography and then of functional magnetic resonance imaging (Raichle, 2009).¹ These technologies have provided a window inside the functioning brain to elucidate how neural processes implement psychological processes.

We are currently on the verge of two new revolutions in psychological methodology. The use of sensor technologies, including smartphones, wearable devices, and the internet of things, is allowing us, for the first time, to collect dense datasets that characterize human behavior and the contexts in which it occurs. Although the need for ecologically realistic research has long been appreciated (Neisser, 1985), until recently the ability to quantify real-world experience has been limited. Experience-sampling methodologies previously relied on beepers that prompted people to record their current thoughts and actions and focused on the participant's subjective experience (Hurlburt & Akhter, 2006; Larson & Csikszentmihalyi, 2014). Today, however, the broad acceptance of smartphones and wearable devices allows us to

✉ Simon Dennis
simon.dennis@unimelb.edu.au

¹ University of Melbourne, Melbourne, Victoria, Australia

² Unforgettable Research Services Pty Ltd, Melbourne, Victoria, Australia

³ University of Tasmania, Hobart, Tasmania, Australia

⁴ University of Newcastle, Callaghan, New South Wales, Australia

⁵ National Institutes of Health, Washington, DC, USA

¹ If one includes electroencephalography as an imaging technique, then the story starts with Hans Berger in 1924.

augment active experience-sampling methods with passive methods that make possible recording people's behavior at scale. As our refrigerators, washing machines, cars, and so forth become increasingly instrumented, the sources of data will multiply, and the picture of human behavior in the wild that we can construct will increase in fidelity.

For example, Nielson, Smith, Sreekumar, Dennis, and Sederberg (2015) investigated the memory representation of space and time in the hippocampus using the experience-sampling technique. In their study, participants were equipped with a smartphone that took photos of the participant's daily life as well as recording Global Positioning System (GPS) and time information. Later, participants were invited to a test session in which they were presented with these photos while in the MRI scanner. The results showed that the left anterior hippocampus represents time and space up to a scale of a month in time and 30 km in space. These findings not only generalize laboratory experiments that are usually tested on a smaller time scale (e.g., seconds or minutes), but also provide stronger evidence that the anterior hippocampus represents time and space in everyday life.

For a clinical application, an example from our ongoing project with bipolar patients shows that mood states can be predicted from the patients' movement (i.e., accelerometry) and social interactions (e.g., number of phone calls and messages). It is difficult to overstate the significance of these developments. The translational relevance of psychology has been stunted by an inability to understand, predict, and influence real-world behavior at the individual level. That is changing.

At the same time, we are in the midst of a revolution in our ability to understand collective behavior. We are now starting to be able to map the interactions of individuals through social media and other forms of electronic communication. For example, Ugander, Backstrom, Marlow, and Kleinberg (2012) utilized Facebook data to examine social contagion. Traditionally, it has been hypothesized that the probability of an individual being affected by a contagion (e.g., joining a Facebook group) is positively correlated with the number of people in the person's contact who are affected by the contagion (i.e., contact neighborhood). However, the results showed that the size of contact neighborhood was not correlated with the probability of individuals being affected, and when component count was controlled, they were actually negatively correlated. These results could not have been observed without examining large-scale social network data.

Behaviors that underpin health, sustainability, and defense must be understood at the group and population levels, and both state and nonstate actors have begun to exploit the ability to influence behavior using social network data (Cadwalladr, 2017). Both to deepen our understanding of social processes and ward against subversive influences, academia needs to catch up.

The potential of sensor and network data is substantial, but to unlock that potential requires researchers to develop new technologies, training programs, and work processes. Increasingly, data will be collected continuously over months and years. The datasets collected will be large, and researchers will need to use machine-learning and dynamical systems methods that have been developed to analyze these data. These data are also often sensitive, so we researchers will need to change the way we interact with data (see the companion article "Privacy Versus Open Science"). In this article, we describe the *unforgettable.me* system, which we are developing to meet these challenges.

The article begins with a description of the search and visualization mechanisms that participants can use to investigate their data for their own purposes or to know what they are allowing researchers to analyze. Then we will introduce the collection mechanisms of the system, including the *unforgettable.me* app as well as the mechanism for connecting to other services. Next, we will talk about the data marketplace, which is the mechanism by which participants and researchers interact. Finally, we will demonstrate the experimental and analysis platforms that provide the privacy layer that allows researchers to interact with data without seeing the individual data points, such as raw images, GPS coordinates, and so forth.

A memory prosthesis

Unforgettable.me provides an extensive set of facilities to allow users to explore their data. The basic interface consists of a search engine, which users can access when they log in. Keywords can be entered in order to retrieve associated events.

Figure 1 shows the output that is displayed when one of the authors (S.D.) searches for "netherlands" in his account. There were a total of 159 results, of which the first ten are listed. By default, the events are listed in reverse chronological order (although this can be changed with the switch just above the results). By clicking on the cross in the corner of each result, events can be removed to trash and then subsequently deleted.

The first event was an article that appeared on the National Public Radio (NPR) news feed on March 30, 2018, at 7 a.m. about the death of Johan van Hulst, a Dutch school teacher who saved hundreds of Jewish children during the Nazi occupation of the Netherlands. The event includes the title, category, summary, and a permanent URL to the article. These articles provide a way to isolate the time at which personal events occurred when the user can relate them to public news items. For instance, you may not be able to remember the date on which you were in Canberra, but you might recall that it was when President Trump was inaugurated. Capturing the NPR events allows you to search for that time and refine your search.

Fig. 1 Sample search result for the query “netherlands”

The second event was an email received by S.D. from Elsevier at 5 a.m. on the same day. Most of us use our email as a form of memory prosthetic already, but by incorporating them into unforgettable.me, it is possible to relate them to other events in your life easily and to visualize the rate at which you send and receive emails over time.

Figure 2 shows the events returned when the query “netherlands __App__” is issued. The “__App__” tag is added to all events that are created by the unforgettable.me app. In this case, we see an event from June 20, 2017, when S.D. was in Amsterdam. A set of searchable keywords appear at the top. These include common temporal indicators (e.g., “Tuesday, June, 2017, Night, spring”), as well as tags indicating the current weather (e.g., “Clear” in this case) and the moon phase. For some locations, the type of activity (restaurant, place of worship, etc.) is also provided, using Google Places. In addition, unforgettable.me runs a set of machine-learning classifiers that operate on the basis of the audio input, which indicate whether there was voice or traffic in the sample as well as whether the user was in a café. Note that the pen icon next to the tags allows the user to add, delete, and modify tags.

Using the GPS coordinates, the address is sourced from Google Places. Determining the exact address is difficult, since GPS accuracy is not perfect and can be affected when the user moves indoors. Consequently, Google provides a set of possible addresses, and unforgettable.me provides a dropdown with these addresses. The user can then modify the address by selecting an alternative option. Clicking on

the location pin symbol opens a map showing the location of the event (see Fig. 3).

Under the address, there are several Count variables—“Location Count,” “Audio Processed Count,” “Accelerometry Count,” and “Battery Count.” These indicate the number of data samples of each kind that were collected by the user during this hour. There are, for instance, two sets of GPS coordinates. These counts are useful for users to determine what they have collected, but they can also be searched to determine whether a given user meets the selection criteria to participate in a research study (see the section entitled “The Data Marketplace” below).

Managing the battery usage of the app is a critical activity, as excessive use can lead participants to discontinue collection. Including the median battery usage as a variable allows the user to visualize the usage as a function of time. Other variables that are calculated are the moon illumination and age, the weather, the temperature, and the number of kilometers that the participant traveled during this hour. The latter variable is currently being used in the previously mentioned study tracking bipolar patients, to see whether changes in the distance moved are predictive of the participants entering different phases of the disorder.

On the right-hand side of Fig. 2, an image of each location is provided. In this case, the image is derived automatically from Google Streetview. If the user is collecting images or raw audio recording through the app, these will be available here as well.

User

Fig. 2 Diary for the query “netherlands __App__”

Fig. 3 Map view for the query “netherlands __App__”

Fig. 4 Event calendar for the query “Bar Beach”

In addition to the search interface, the system provides a variety of visualizations. For instance, Fig. 4 shows a calendar view when S.D. issues the query “Bar Beach.” S.D. plays tennis at Bar Beach, so the calendar provides a record of his exercise. With some personal knowledge, S.D. is able to provide an interpretation of the graph. July–August 2016 was a period in which his arm was injured, and August 2017 marks the start of a new job.

In addition, S.D. has his personal scales linked to the unforgettable.me system. Whenever he steps on the scales, his weight is automatically uploaded. Figure 5 shows the (pleasing) downward trend during 2017 and early 2018 (notably, after he took the new job).

Unforgettable.me provides a substantial set of visualizations that is continually growing. As we argue in the companion article “Privacy Versus Open Science,” for individuals to be vigilant about their privacy, they have to be engaged with their data. Beyond the usefulness of the search and visualization interfaces when completing transactions with researchers, they are intrinsically motivating. The search tool provides a form of memory prosthesis that people can use to recall what they were doing at any given time. The visualization tools allow users to discover patterns and relationships in their lives about which they might not have been conscious. These kinds of facilities are critical if we are to transform into a more data-aware populace.

In this section, we have provided a sense of the kinds of data that unforgettable.me can capture and visualize. In the next section, we describe how the data are collected.

The unforgettable.me App

The heart of the data collection system is an Android ²app(search for unforgettable.me on Google Play) that is capable of collecting accelerometry, audio, images, and GPS coordinates. Although the rest of the unforgettable.me system can be used by anyone, the app is currently available only for Android users. Figure 6 shows the four tabs that are available when you open the app.

The first tab (Fig. 6a) allows users to log into their account. Users must already be registered on the unforgettable.me site in order to connect the app, and they must use the same username and password. There is also an indicator in the bottom right-hand corner that makes it clear when the user is logged in.

On the second tab (Fig. 6b), there is a large “all stop” button, which allows the user to cease all data collection with a single button press—avoiding any ambiguity about which data streams are active. This capacity is useful when the user is in school, family, or work circumstances in which others may object to data collection, or in secure situations such as border control facilities. On this tab, the user can also see the total number of events that have been collected for each of the data streams since the last upload. The user can monitor these numbers in order to troubleshoot their data collection. If a stream is gray (as the image stream is in the figure), then the user is not collecting that stream. In the bottom right-hand

² There are plans to create an iOS version.

Fig. 5 Chart for the query “weight,” restricted to the last year. Values come from Aria scales linked through a Fitbit

corner is an indicator that signals whether the phone is logging data or whether data are being uploaded currently.

To give users maximum control over their data, the interface allows them to delete data before the data leave the phone. Users can set a “days to review” period, which

indicates how long data will be held before being uploaded. On the third tab, the dates for which data have been collected on the phone are displayed (Fig. 6c top). If the user selects a date, the hours on that date are displayed (Fig. 6c bottom). Selecting the “Delete” button to the right of a time will prevent

Fig. 6 The tabs available in the unforgettable.me app

those data from being uploaded. Selecting the “Restore” button allows users to reverse their decision if they made an error.

Any data that are not marked for deletion will be uploaded after the “days to review” period has elapsed, when the battery level is at least 90% and the phone is connected to wifi. These later stipulations ensure that the phone will not be drained of power when the user attempts to use it and that uploading data will not incur significant expense as a consequence of data charges.

The final tab (Fig. 6d) allows the user to control data collection with greater specificity. Each of the main data streams (accelerometry, GPS, audio, raw audio, and camera) can be switched on or off individually. By default, accelerometry is collected continuously. Depending on the system hardware and the resources being consumed by other applications, accelerometry will be sampled at between 8 and 15 Hz. The other data streams are triggered either at regular time intervals or on the basis of the accumulation of accelerometry—so that the sampling rate increases when you move. Using this method protects battery life, because sampling is minimized when there is likely to have been little change in the prevailing context.

Audio segments are of 500-ms durations. The default audio format is Mel frequency cepstral coefficients (MFCCs). MFCCs are frequency-based representations, which can be used by machine-learning classifiers to extract aspects of the ambient audio environment, such as whether there are voices or traffic noise. If the “raw audio” switch is on, however, raw audio (in 3-s segments) will be collected.

This tab also contains the Days to Review selector and the Recording Trigger selector (time, movement, or combined). The trigger interval, when using time-based triggering, is located here also. On some phones, autofocus can prevent images from being taken, so a switch is available to prevent these problems. In addition, a Notifications switch can be turned on to help with troubleshooting. Finally, there is a “Copy

Database” button that will transfer the data currently on the phone to an SD memory card. This capability is useful for debugging purposes, but it also provides a mechanism for users to bypass the rest of the unforgettable.me system if they wish to use the app independently of the site.

Connecting additional data sources using If This Then That

Unforgettable Research Services has partnered with If This Then That (IFTTT, ifttt.com) to allow the collection of some 400 sources, including social network data (e.g., Facebook, Twitter, Instagram), wearable data (e.g., Fitbit, Strava, Android wear), and internet-of-things data (e.g., Fitbit scales, GE refrigerators, BMW cars). To collect data in this way, users must first register with IFTTT. After registering, the easiest way to start collecting data is to activate an existing applet provided by the unforgettable.me service (see Fig. 7).

Table 1 lists the applets that Unforgettable Research Services had created at the time of writing. There are a wide variety, and users can add their own (see the unforgettable.me website for instructions). Some of the applets function automatically as they are installed and link one kind of data to unforgettable.me (e.g., the NPR Stories applet automatically adds NPR stories to the server). Other applets are intended for use with the IFTTT button system, through which data are stored when the user presses an icon on their phones (e.g., the Took Meds applet exists as a button and records the event whenever the user presses the button).

The applets marked with a * provide a more private mechanism for collecting short message service (SMS) and call data. A unique identifier derived from the sender or receiver phone number is created. This identifier is sufficient to

Fig. 7 To activate existing applets, log in to your If This Then That (IFTTT) account and search for “unforgettable”

Table 1 List of existing applets, grouped by type**a** Communications

Gmail received	Gmail sent	Upload your Gmail attachments to Unforgettable	Record Android SMS messages received	Record Android SMS messages sent
Record Obfuscated Android SMS messages received*	Record Obfuscated Android SMS messages sent*	Record Android phone calls made	Record Android phone calls received	Record Obfuscated Android phone calls made*
Record Obfuscated Android phone calls received*				

b Calendars and Scheduling

Google calendar events	Record Office 365 emails received	Record Office 365 calendar events
------------------------	-----------------------------------	-----------------------------------

c Photos

iOS photos	Android photos	Flickr Public Photos
------------	----------------	----------------------

d Social

Record Facebook status messages	Record Facebook Shares	Record Twitter Tweets	Record Weibo Texts	Tumblr Post
---------------------------------	------------------------	-----------------------	--------------------	-------------

e Internet of Things and Wearables

Steps from Fitbit	Log weight from Fitbit	Record Strava Activity	Record Misfit Activity Summary
-------------------	------------------------	------------------------	--------------------------------

f News and Information

NPR Stories	Get an email when a new Unforgettable Me Applet is published	Get an email when Unforgettable Me publishes a new trigger or action	Add Note
-------------	--	--	----------

g Buttons

Amused	Angry	Bloated	Can't Sleep	Constipation
Depressed	Diarrhea	Dizziness	End Menstruation	Excited

Fig. 8 Example of obfuscated short message service (SMS) events

Fig. 9 An Android phone home screen showing a set of IFTTT buttons that can be used to register moods and behaviors

recognize that the same person is involved, but it can't be used to determine who the person is. The content of the SMSs is also omitted in these applets (see Fig. 8).

Buttons can be added to the home screen of Android or iOS devices and provide a mechanism for users to record mood, physiological states, and behaviors that are not easily inferred automatically from other data streams. Figure 9 shows a set of these buttons on an Android screen. To record data, the user presses the button. Each press registers as an event in unforgettable.me. Since users are already accessing their phones when they are confronted by the buttons, the buttons provide a mechanism that is sustainable over long periods. There is no interruption to flow, as there is with other forms of ecological momentary assessment. The disadvantage is that the times at which events are recorded are not sampled uniformly, so there can be systematic biases.

The data marketplace

A fundamental part of the Unforgettable Research Services philosophy is that users retain ownership of their data.³ When a researcher wishes to access the user's data, the researcher issues a request on the unforgettable.me data marketplace (see Fig. 10). The request includes a query that specifies

³ This principle is formalized in the user agreement (www.unforgettable.me/terms), the researcher agreement (www.unforgettable.me/researcher-terms), and the privacy policy (www.unforgettable.me/privacy).

Add Unforgettable Project

Researchers:

[+ Add Researcher](#)

Title:

Unforgettable Query:

Purpose:

Pay:

Start date:

End date:

Additional tasks:

Inclusion criteria:

Exclusion criteria:

Associated Files

[Upload File\(s\)](#)

Submit

Fig. 10 Researcher sign-up form

unforgettable.me
recollect • reflect • reminisce
BETA

[Getting Started](#)
[Research Services](#)
[Blog](#)
[example@unforgettable.me](#)

01 Jan 1970 to 01 Apr 2018

1 of 1 results

Title: Understanding human memory using experience sampling technologies

Researchers:
[Dr. Simon Dennis](#)
[Dr. Adam Osth](#)
[Dr. Hyungwook Yim](#)
[Dr. Benjamin Stone](#)
[Miss Wan Rong Ong](#)

Pay: AUD 65

Start Date: 2018-02-01

End Date: 2018-03-01

Unforgettable Query: __APP__

1

Fig. 11 User view of marketplace

Fig. 12 Screenshot showing the implementation of a memory experiment. Each trial shows a time and a map with four alternatives from which the participants choose. These alternatives are generated by

the experimental code from the users' data but are not available to researchers directly. Users have the ability to pan and zoom the map

the purpose of the project, any inclusion or exclusion criteria, who is asking for the data, what data the researcher wants from a user's existing collection, what additional tasks the user will be asked to complete, and the compensation that the user will receive. Researchers must already have received ethical approval for their research from their home institution, and the associated documents must be attached to the request.

Unforgettable Research Services can collect gigabytes of data per month on an individual. Currently, over 500 people are using the system. From a researcher's point of view, this is advantageous in that it might not be necessary to enlist new participants and wait for data collection to occur. The data might already exist and be usable in their current form, or when augmented by running an experimental paradigm.

Eligible users who have consented to being contacted will then see the request appear in their Unforgettable Research Services account (see Fig. 11). Periodic emails also make people aware of any experiments in which they may participate. If they chose to participate, they may click on the request and confirm their desire to take part. If the researcher is only asking for access to their data, then the transaction will be completed immediately. If the user must complete additional tasks, such as running experimental protocols or completing surveys, then these will be initiated.

Unforgettable Research Services charges researchers a fee based on the number of participants and the length of time over which the requested data were collected, and thus it is not dependent on government funding sources. In addition,

researchers pay whatever compensation they set in the request to the users (through the system). The objective is to create a sustainable ecosystem of facilities and participants that is usable for a broad range of research purposes.

The experimental and analysis platforms

All experiments (or surveys) on the unforgettable.me system are conducted within a password-protected environment. Because the user is already authenticated, data can be presented to the user without violating their privacy. For instance, in an experiment we are currently running, participants are presented with a map showing four locations and are asked where they were at a given time (see Fig. 12). The experimental code clusters the participant's GPS coordinates for the given period and then selects a target and three distractors for each trial. The participant runs the experiment and makes selections, and the data are added to the user's personal repository. The researcher also has access to these records, but they contain only the event identifiers (random keys) that correspond to the target and distractor coordinates. The researcher does not have access to the GPS coordinates themselves.

At this point, experimental code is written in a combination of JavaScript and Python. Unforgettable.me provides a set of standard experiments that researchers can modify in order to implement their own experiments. In addition, if researchers require capabilities that are not available in our standard

modules, they can write their own code. This code then goes through a review by the company to ensure that it does not expose any raw data outside the authenticated environment. In the future, we hope to provide a specially designed language that does not allow private data to be exposed, in order to remove this review step.

Analyzing data

One of the most challenging issues when constructing a privacy-preserving data analysis system is how to allow researchers to conduct analyses when they are not permitted to see the data. The system must provide protection from potentially malicious researchers, who would attempt to reconstruct an individual's data using the analysis system. One might imagine that it would be sufficient to allow only analyses that provide aggregate statistics—such as the mean, standard deviation, or group-wise inferential tests. However, the motivated attacker can use these statistics to isolate individual values (see the companion article “Privacy Versus Open Science” for an example). Consequently, the results returned to researchers must be random samples, and the system must prevent the coder from gaining access to deterministic values derived from private data. This requirement makes the use of more common data analysis languages such as R or Python inappropriate, because they were designed assuming that the coder should have access to the data. At the same time, the system must be sufficiently flexible that researchers have the ability to conduct a broad range of analyses.

To meet these objectives, Unforgettable Research Services has implemented an augmented version of JAGS (Just Another Gibbs Sampler; Plummer, 2003). JAGS is a close derivative of the BUGS language (Lunn, Spiegelhalter, Thomas, & Best, 2009) and is used to implement Bayesian models. It has several properties that make it well-suited to providing a privacy-preserving analysis platform.⁴

First, although it is quite a simple language, it allows a wide variety of models to be specified in an intuitive form. Several textbooks provide example scripts for standard statistical tests, as well as explaining how to construct more sophisticated models tailored to the question under investigation (Kruschke, 2014; Lee & Wagenmakers, 2013).⁵

Second, the simple dependency graphs that it creates to express models can be analyzed syntactically to establish the privacy status of some variables. In any given model, some variables can be mined to extract personal information,

whereas others are specified at a sufficiently global level that they are safe to release. In privacy-preserving imperative frameworks (e.g., JEEVES; Yang, Yessenov, & Solar-Lezama, 2012), the privacy status of variables can be declared, but whenever a statement is executed involving such a variable, computation must be forked (at least conceptually, and often in practice) in order to account for the different trajectories that the code may have taken based on the protected values. These forks can be computationally expensive and make errorless coding nontrivial. Relying on a declarative language such as JAGS eliminates these difficulties and simplifies automatic analysis significantly (see the example below for an illustration).

Finally, JAGS has no primitives to access the internet, thus eliminating one of the major methods by which the system could be used to compromise privacy.

Although it is beyond the scope of the present article to provide a comprehensive overview of the analysis system, we provide an example to illustrate how the system operates. Figure 13 shows example code (top) and the graphical model (bottom) for an analysis that estimates the mean (μ) and precision ($\tau = 1/\sigma^2$) of the latitudes of a set of participants. To provide access to the experience-sampling data, the JAGS language has been augmented with a set of standard variables. In particular, an `event` is the stream of records associated with the query provided by the researcher to the participants who have been recruited through the data marketplace. Fields of these records are specified with a dot notation. For example, the stream of latitudes associated with the retrieved events is provided as `event.latitude`. If a record does not have an associated latitude, it will be filled with an NA (value not available). Subjects are numbered from 1 to `NumberOfSubjects`, and `subjectids` is used to identify which subject is associated with each event.

In the example JAGS code, latitudes are modeled as draws from a normal distribution with a mean (`mu`) defined hierarchically for each subject and a single precision (`group.tau`) defined over the group. All variables within the subject plate (Fig. 13 bottom) are clearly dependent on private data and cannot be released. These variables can be identified directly by analyzing the code.

The remaining variables, `hier.mu`, `hier.tau`, and `group.tau`, may or may not be sensitive to individual participants' data. To determine which variables can be released, $N+1$ Markov chain Monte Carlo chains are run, where N is the number of subjects. One of the chains is fit to the data from all the participants. The other chains are fit to datasets from which one of the participants has been omitted. Empirical Kullback–Leibler privacy (see the Appendix) is used to estimate the maximum amount of information that will be revealed about a participant by the samples of each variable. If, in any case, the amount of information released exceeds one bit, the samples of that variable are withheld.

⁴ Note that we are not claiming that JAGS is the only language that meets these criteria; other languages, such as STAN or PyMC3, could have been used.

⁵ In the remainder of this description, we will assume a knowledge of Bayesian modelling and the JAGS language. These textbooks provide excellent introductions.

```

model {
  # observable variables
  for (e in 1:length(event.latitude)){
    event.latitude[e] ~ dnorm(mu[subjectids[e]], group.tau)
  }

  # subject variables
  for (subject in 1:NumberOfSubjects){
    mu[subject] ~ dnorm(hier.mu, hier.tau)
  }

  # group variable
  group.tau ~ dgamma(0.001, 0.001)

  # hierarchical variables
  hier.mu ~ dnorm(-30, 0.001)
  hier.tau ~ dgamma(0.001, 0.001)
}

```

Fig. 13 (Top) JAGS script for estimating the mean (hierarchical) and precision (group) of the latitudes of a set of subjects. (Bottom) The corresponding graphical model

To illustrate, the code above was run on a selection of participants. The results were as follows:

```

group.tau: Did not meet privacy threshold.
maxkl = 6.658
hier.tau: Did not meet privacy threshold.
maxkl = 6.289
hier.mu: Did not meet privacy threshold.
maxkl = 5.060

```

The variable `group.tau` is more sensitive to individual data and was withheld, whereas `hier.tau` and `hier.mu` both had maximum Kullback–Leibler values below 1 and were released.

Commonly, it will be the case that adding more participants to the set will decrease the privacy sensitivity of variables, since any one participant's data will have an increasingly small impact upon the general results of the model. However, one cannot assume that this will always be true. In this example, adding another subject generates the following results:

```

group.tau: Did not meet privacy threshold.
maxkl = 6.658
hier.tau: mean = 335.504 5% = 1.994
median = 169.392 95% = 1214.422
hier.mu: mean = -37.829 5% = -38.045
median = -37.841 95% = -37.631

```


The additional participant increases the maximum empirical Kullback–Leibler privacy, preventing the release of `hier.tau`. Note that this result in itself provides some information (we now know that the participant who was added differed from the other participants) and so is not privacy-neutral. However, this method provides a compromise that avoids excessive increases in the computation that is required in order to provide a result, does not involve the adding of extrinsic noise to results, allows us to gain estimates of the variability in our parameters, and can be applied to a wide variety of models (in fact, to any model that can be specified in the JAGS language)—all of which are critical criteria for the system to be usable by researchers.

Conclusions

Experience-sampling approaches to psychological phenomena are likely to be transformational. To realize their potential, however, we must create scalable collection and analysis platforms that preserve the privacy of participants. In this article, we have described the user retrieval and visualization engine, the data collection pipelines, the data marketplace that links participants to researchers, and the experimental and analysis systems of the *unforgettable.me* system. The system provides a way to conduct ecologically realistic research and has the potential to be used in studies that examine collective behavior. Multiple studies that have benefited from the present system (e.g., Dennis et al., 2017; Nielson et al., 2015; Sreekumar, Dennis, & Doxas, 2017; Sreekumar, Dennis, Doxas, Zhuang, & Belkin, 2014) and numerous ongoing studies show the potential for the system (see the companion article “Privacy Versus Open Science” for more examples).

Author note This research was supported by the Australian Government through the Australian Research Council’s Discovery Projects funding scheme (project DP150100272). The views expressed herein are those of the authors and are not necessarily those of the Australian Government or the Australian Research Council.

Compliance with ethical standards

Conflict of interest A critical part of the open science movement involves the open disclosure of both the explicit and potential implicit motivations for scientific work. In that spirit, the reader should be aware that S.D. is the CEO of a startup called *Unforgettable Research Services Pty Ltd* (URS) that specializes in providing privacy-preserving experience-sampling collection and analysis services. B.S. is the chief technical officer of URS.

Open practices statements No experiment was involved in the present article, and none of the data reported here are available. However, an example dataset through the system is accessible at <https://unforgettable.me>.

Appendix: Empirical on-average Kullback–Leibler privacy

In this appendix, we situate the notion of empirical Kullback–Leibler (KL) privacy that we use to decide how many samples from a program can be released within the privacy literature.

Differential privacy (Dwork, 2008) is a probabilistic definition of the maximum privacy loss an individual can experience when their data are used in the calculation of a function. It is defined as

$$p(f(D) \in h) \leq e^\epsilon p(f(D^-) \in h) \forall D, D' d(D, D') \leq 1 \forall h \subseteq R,$$

where f is a probabilistic function of the data that we wish to release, D^- is a dataset that contains the same data as D except those contributed by one subject, h is a measurable subset of R (the range of the function f), and ϵ is the privacy bound.

Provided that f is a density, we can rewrite the definition of differential privacy as follows (Wang, Lei, & Fienberg, 2016):

$$\sup_{D, D': d(D, D') \leq 1} \sup_{h \in H} \log \frac{f(h, D)}{f(h, D^-)} \leq \epsilon.$$

The log expression is the KL distance between the distributions over h . If we employ log base 2 rather than the natural log, then this term can be thought of directly as the information gained by a potential adversary with respect to h , in bits.

Differential privacy has become a de facto standard in the privacy literature, but it is often too restrictive to be practical. As a consequence, many alternative definitions have been proposed (see Y. X. Wang et al., 2016, for a survey). Here we develop the line of successive approximations that leads to our definition.

Wang et al. (2016) pointed out that KL privacy as defined by Barber and Duchi (2014) is equivalent to differential privacy with the second *sup* function replaced by an expectation over $r \in R$:

$$\sup_{D, D': d(D, D') \leq 1} \int_{r \in R} r \log \frac{f(r, D)}{f(r, D^-)} dr \leq \epsilon.$$

Instead of insisting that every point must lie within the bound, the definition is weakened to being an expectation that all points lie within the bound. Y. X. Wang et al. then defined the on-average KL privacy as

$$\frac{1}{N} \sum_{D, D': d(D, D') \leq 1} \int_{r \in R} r \log \frac{f(r, D)}{f(r, D^-)} dr \leq \epsilon,$$

where N is the number of subjects. The on-average KL privacy thus calculates expectancies over both datasets and the range of the function f .

In the context of constructing a privacy-preserving language, the problem with most definitions of privacy is that they require one to prove properties of the function to be computed and the scope of the possible datasets. Because

the function to be computed is provided by the coder, in our case, it is unknown a priori. Similarly, it is a strong condition to protect against all possible datasets, when in practice we have only particular datasets that we wish to protect. Another issue in our context is that algorithms for preserving privacy often involve the addition of extrinsic noise of unknown magnitude, making interpretation difficult.

Bayesian methods that rely on taking samples from a posterior distribution already posit a probabilistic function of the data and are well-established. Critically, the variability that they propose is intrinsic to the statistical inference problem they are trying to solve, and is therefore irreducible. Our objective, then, is to calculate how much information a potential adversary is likely to gain from a set of samples that we propose to release. The number of samples to release can then be determined by a privacy bound.

We define the empirical KL privacy as follows:

$$\sup_{i=1..N} \sum_{s \in S_R} \log \frac{\hat{f}(s, D)}{\hat{f}(s, D^{-i})} \leq \epsilon,$$

where N is the number of subjects, D is the data of all subjects, D^{-i} is the data from all subjects except subject i , S_R is the candidate set of samples for release, and $\hat{f}$ is an empirical estimate of f .

Similarly, we define the empirical on-average KL privacy as

$$\frac{1}{N} \sum_{i=1..N} \sum_{s \in S} \log \frac{\hat{f}(s, D)}{\hat{f}(s, D^{-i})} \leq \epsilon.$$

Rather than estimate f , we propose to derive an estimate of the pointwise KL distance from an independent set of samples using a k -nearest-neighbor procedure inspired by Q. Wang, Kulkarni, and Verdú (2009).

Let x_j be samples drawn from D , and x_j^{-i} be samples from D^{-i} . Also, let x_{NN} be the k th nearest neighbor of s among the x_j , and x_{NN}^{-i} be the k th nearest neighbor of s among the x_j^{-i} . We then estimate the pointwise KL divergence as

$$\log \frac{\hat{f}(s, D)}{\hat{f}(s, D^{-i})} = \log \left| \{x_j \leq \min(x_{NN}, x_{NN}^{-i})\} \right| - \log \left| \{x_j^{-i} \leq \min(x_{NN}, x_{NN}^{-i})\} \right|.$$

Note that either $|\{x_j \leq \min(x_{NN}, x_{NN}^{-i})\}|$ or $|\{x_j^{-i} \leq \min(x_{NN}, x_{NN}^{-i})\}|$ will be equal to k .

References

- Barber, R. F., & Duchi, J. C. (2014). *Privacy and statistical risk: Formalisms and minimax bounds*. arXiv preprint. arXiv:1412.4451
- Cadwalladr, C. (2017). The great British Brexit robbery: How our democracy was hijacked. *The Guardian*, 20.
- Dennis, S. J., Yim, H., Sreekumar, V., Evans, N. J., Garrett, P., & Sederberg, P. (2017). A hierarchical Bayesian model of “memory for when” based on experience sampling data. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.) *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 295–300). Austin, TX: Cognitive Science Society.
- Dwork, C. (2008). Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation* (pp. 1–19). Berlin, Germany: Springer.
- Fancher, R. E., & Rutherford, A. (2012). *Pioneers of psychology*. New York, NY: Norton.
- Hurlburt, R. T., & Akhter, S. A. (2006). The descriptive experience sampling method. *Phenomenology and the Cognitive Sciences*, 5, 271–301.
- Kruschke, J. K. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2nd ed.). Boston, MA: Academic Press.
- Larson, R., & Csikszentmihalyi, M. (2014). The experience sampling method. In *Flow and the foundations of positive psychology* (pp. 21–34). Berlin, Germany: Springer.
- Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian cognitive modeling: A practical course*. Cambridge, UK: Cambridge University Press.
- Lunn, D., Spiegelhalter, D., Thomas, A., & Best, N. (2009). The BUGS project: Evolution, critique and future directions. *Statistics in Medicine*, 28, 3049–3067.
- Neisser, U. (1985). The role of theory in the ecological study of memory. *Journal of Experimental Psychology: General*, 114, 272–276. <https://doi.org/10.1037/0096-3445.114.2.272>
- Nielson, D. M., Smith, T. A., Sreekumar, V., Dennis, S. J., & Sederberg, P. B. (2015). Human hippocampus represents space and time during retrieval of real-world memories. *Proceedings of the National Academy of Sciences*, 112, 201507104.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing* (pp. 1–10). Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.r-project.org/conferences/DSC-2003/Proceedings/Plummer.pdf>
- Raichle, M. E. (2009). A brief history of human brain mapping. *Trends in Neurosciences*, 32, 118–126. <https://doi.org/10.1016/j.tins.2008.11.001>
- Sreekumar, V., Dennis, S., & Doxas, I. (2017). The episodic nature of experience: A dynamical systems analysis. *Cognitive Science*, 41, 1377–1393.
- Sreekumar, V., Dennis, S., Doxas, I., Zhuang, Y., & Belkin, M. (2014). The geometry and dynamics of lifelogs: Discovering the organizational principles of human experience. *PLoS ONE*, 9, e97166. <https://doi.org/10.1371/journal.pone.0097166>
- Ugander, J., Backstrom, L., Marlow, C., & Kleinberg, J. (2012). Structural diversity in social contagion. *Proceedings of the National Academy of Sciences*, 109, 5962–5966.
- Wang, Q., Kulkarni, S. R., & Verdú, S. (2009). Divergence estimation for multidimensional densities via k -nearest-neighbor distances. *IEEE Transactions on Information Theory*, 55, 2392–2405.
- Wang, Y. X., Lei, J., & Fienberg, S. E. (2016). On-average kl-privacy and its equivalence to generalization for max-entropy mechanisms. In *International Conference on Privacy in Statistical Databases* (pp. 121–134). Cham, Switzerland: Springer.
- Yang, J., Yessenov, K., & Solar-Lezama, A. (2012). A language for automatically enforcing privacy policies. In *ACM SIGPLAN notices* (Vol. 47, No. 1, pp. 85–96). New York, NY: ACM Press.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.