The Fallacy of an Airtight Alibi:

Understanding Human Memory for *Where* Using Experience Sampling

Elizabeth Laliberte*, Hyungwook Yim*, Benjamin Stone, and Simon J. Dennis

School of Psychological Sciences, The University of Melbourne, Australia

word count 739 (351+388) /750

this article has been published in *Psychological Science*
<https://doi.org/10.1177/0956797620980752>

Author Note

This research was supported under Australian Research Council's Discovery Projects funding scheme to SJD (DP150100272). Note that SJD is the CEO and BS is the CTO of a startup called Unforgettable Research Services Pty Ltd that specializes in providing privacy-preserving experience sampling collection and analysis services. Part of the study was presented at the 42nd Annual Conference of the Cognitive Science Society, where only the abstract was included in the proceedings. Instructions to access data and the related code can be found at <https://osf.io/7ym9f/>. Correspondence concerning this article should be addressed to Simon J. Dennis or Hyungwook Yim (Melbourne School of Psychological Sciences, The University of Melbourne, Parkville, VIC 3010 Australia. E-mail: simon.dennis@unimelb.edu.au; hyungwook.yim@gmail.com). * *these authors contributed equally to this work as a co-first author.*

Abstract

A primary challenge for alibi generation research is establishing the ground truth of the real world events of interest. In the current study, we used a smartphone app to record data on adult participants ($N = 51$) for a month prior to a memory test. The app captured accelerometry, GPS location, and sound environment every ten minutes. After a week retention interval, we asked participants to identify where they were at a given time from among four alternatives.

Participants were incorrect 36% of the time. Furthermore, our forced choice procedure allowed us to conduct a conditional logit analysis to assess the relative importance of different aspects of the events to the decision process. We found strong evidence that participants confuse days across weeks. In addition, people often confused weeks in general and also hours across days. Similarity of location induced more errors than similarity of sound environments or movement types.

Keywords: episodic memory; memory for where; experience sampling; alibi; autobiographical memory

149/150 words

Statement of Relevance

Project Innocence estimates that between two and five percent of inmates have been falsely incarcerated. While the processes involved in eyewitness identification have received significant attention, the ability of suspects to provide accurate alibis is more poorly understood. The difficulty is mainly rooted in finding an ecologically valid method for capturing one's personal life, which can also be tested later. In the current study, we use passive experience sampling methods utilizing smartphones to overcome this problem. With the ground truth available, we ask where they were at certain times to determine what kind of mistakes people tend to make. The results can be used to alert investigators to the questions they should ask in order to catch the memory errors suspects are likely to make.

125/150 words

The Fallacy of an Airtight Alibi:

Understanding Human Memory for *Where* Using Experience Sampling Technology

In 1985, Ronald Cotton was falsely convicted for charges of rape and burglary and was sentenced to life plus 50 years in prison. Nearly 10 years later, Cotton was exonerated on the basis of DNA evidence. Alongside an error made by the victim in the eyewitness identification, Cotton provided a false alibi that could not be corroborated. Rather than report where he had been at the time of the crime, Cotton mistakenly recalled where he had been the week before. His error probably contributed significantly to the wrongful conviction that he received (Thompson-Cannino, Cotton, & Torneo, 2009).

For memory science to make a more substantive contribution to ameliorating these kinds of issues, we need to develop methods to capture a continuous record of the experiences of people as they engage in their daily activities and develop theories that capture the nature of representation and interference in real world settings. In this paper, we examine how people retrieve where they were at a given time as might happen when a suspect is asked for an alibi. We used a smartphone app to record location, movement and auditory data continually for a month, thus providing the ground truth data to score their retrieval and a characterization of all of their experiences.

Models of memory (Shiffrin & Steyvers, 1997; Clark & Gronlund, 1996; Osth & Dennis, 2015) propose that the similarity between the cue and the memory trace, and between the required target and the memory trace will have substantive effects on performance. In the current work, we test these theories in a real world setting. In our task, the cue is a time and the required target is a location, so the current models would suggest that performance should be heavily

influenced by temporal and spatial similarity and less so by similarity of movement and auditory aspects of the experience.

Our design also affords an opportunity to consider the nature of the representation of time in memory - in particular we will investigate if the representation of time is continuous in nature or is better captured by a categorical representation (Friedman & Wilkins, 1985; Yim, Garrett, Baker & Dennis, 2018).

Method

In the current study, we used smartphones to collect a continuous record of people's experience as they went about their daily lives. The study examined factors that may influence memory retrieval for locations, including variables such as spatial distance, environmental sounds, accelerometry of movement, and timing of events. Using this method, prior work has investigated how people remember *when* an event happened (Dennis, Yim, Sreekumar, Evans, Garrett & Sederberg, 2017) and the neural representation of autobiographical information (Chow & Rissman, 2017; Nielson, Smith, Sreekumar, Dennis & Sederberg, 2015; Sreekumar, Nielson, Smith, Dennis, & Sederberg, 2018). To the best of our knowledge, this is the first study to use the current methodology to study memory for *where*.

Participants

Fifty one adults participated in the current study. Forty four were recruited through the Unforgettable Research Services (URS, www.unforgettable.me) platform and were compensated AUD\$ 25. These participants can be acquired from the URS system using the key "augmenter_list: __UProjecta46bf991b47afd37abac718416108eed__". Seven were recruited through the university research experience pool for course credit. These participants can be

acquired from the URS system using the key

“augmenter_list: __UProjecte6d90bec11c11c416ac41d8f979d19e4__” (see instructions for accessing the data and the related code at <https://osf.io/7ym9f>). The research was approved by the University of Melbourne Human Research Ethics Committee. As our primary concern is the estimation of conditional logit regression parameters using Bayesian methods, we employed a collect for precision approach. We collected until the widths of the 95% credible intervals were less than 1.5 in all cases. No subjects were excluded from the study.

Procedure and Design

Participants downloaded the URS app (Dennis, Yim, Garrett, Sreekumar, Stone, 2019) on their smartphone and collected data for four weeks. The URS app automatically collected GPS, Accelerometry, and obfuscated audio using Mel Frequency Cepstral Coefficients (MFCC; Mermelstein, 1976). GPS coordinates were collected every 10 minutes (~ 8640 pairs per subject). Accelerometry vectors were collected ~10 times a second continuously (~ 51M data points per subject). These samples were divided into 600 millisecond segments and a Fast Fourier Transform was applied to each coordinate separately. Three second segments of audio were collected in an obfuscated form using MFCC every 10 minutes (~ 8640 samples per subject). The MFCC format strips phase information and only leaves frequency information so that sounds such as conversation can not be reconstructed but ambient sound information of the environment (i.e., talking with people, traffic noise, etc.) can be classified through machine learning algorithms.

The URS app provides several mechanisms to allow participants to control data collection and protect their privacy. A large pause button (Figure 1A) can be used to stop all data

collection at times when the participant feels this appropriate. The same screen allows participants to see which sensors are currently active (the blue labels) and how much data has been collected today in that stream. Participants could determine which types of events would be recorded, how often they were recorded and how long they had to review the data before it left the phone (Figure 1B). They were also able to delete segments from the phone (Figure 1C), and were able to delete data once it reached the server through a search interface on the Unforgettable.me website.

Figure 1. Smartphone application used in the experiment (A~C), and an example of the memory experiment (D).

After the four-week data collection period, there was a one-week retention interval before participants were invited to take an on-line memory test. The memory test consisted of 72 trials, where in each trial, participants were prompted with a time and date cue and asked to select one of four markers labelled A, B, C or D on the google map provided (e.g., “*Click marker where you were Thursday morning 8am on Aug 15, 2019*”; Figure 1D). Participants could interact with the google map, including panning to any location or zooming in and out. After each trial, participants were asked to rate their confidence on a 5-point scale ranging from 1 being ‘not at all confident’ to 5 ‘very confident’.

Test trials were constructed from each participant’s data by first clustering the median GPS point of each hour time bin. For each hour bin, the GPS point closest to the median of that bin's location points was selected. Hierarchical clustering was performed on these locations to assign each hour bin to clusters. Each hour bin was added to a cluster if it's location was within 100m of the center (average) of that cluster's location points. On average, participants had 47.41 clusters in total ($SD = 32.32$, $Mdn = 40$, $Min = 8$, $Max = 148$), with an average of 9.60 members per cluster ($SD = 45.47$, $Mdn = 1$, $Min = 1$, $Max = 640$). Based on the clustering results, each trial was constructed where the target location was first randomly picked from the smallest cluster in the pool. Then three distractor locations were picked randomly from clusters different from the target cluster with a constraint of being more than a 100m distant from the target location. This ordering was done to maximize the spread of target locations for each participant. This process was repeated until 72 trials were created for a given participant.

Results

In the Unforgettable system, an event is defined as an hour period. Over four weeks the maximum number of events a participant could collect is 672. However, participants have the freedom to turn off the app and to delete events either before or after they leave the phone - a mechanism designed to protect their privacy. As a consequence, the average number of events collected per subject was 458.59 ($SD = 128.98$). Given the sensitive nature of experience sampling data collected for our experiment (e.g. frequent GPS locations of participants), previous methods of participant deidentification are inadequate for this work. Powerful re-identification algorithms can be utilized to induce sensitive personal data (Dennis, Garrett, Yim, Hamm, Osth, Sreekumar & Stone, 2019). Therefore, participant data was analysed using a language called Private that is built into the Unforgettable platform. Private allows data to be analysed whilst maintaining rigorous levels of confidentiality and incorporates methods for blinding researchers to raw participant data, while only allowing access for group-level statistics if they can not be used to reverse engineer individual raw data (see Dennis, Garrett, Yim, et al., 2019 for the detailed method).

Analyses were conducted on the pooled group data, which had 3672 trials. All statistics in the current study were implemented with the Private language, which uses a Bayesian statistics framework. Model parameters were estimated using Hamiltonian Monte Carlo with 8 chains, 1000 tuning samples, and 1500 samples. The models use a prior with a normal distribution with a mean of zero and a standard deviation of 100 for continuous variables, a beta distribution with alpha of 1 and beta of 1 for accuracy, and a half-normal distribution with a standard deviation of 100 for reaction time and confidence interval. All models converged with

Gelman-Rubin $\hat{r}$ values below 1.0. We reported 95% Credible Interval (CI_{95}) of the posterior distribution of the parameters, and interpreted the parameters based on whether they excluded zero or not. For the estimated parameters, we also reported effect size (ES). As there is no agreed method of calculating the effect size in the Bayesian framework, we defined the ES as the mean of the posterior distribution divided by the standard error term. The standard error term was calculated by dividing the half-length of the CI_{95} by 1.96, which is the Z value used when calculating the frequentist 95% confidence interval.

The mean accuracy rate across participants was .64 ($CI_{95} = [.62 .65]$), with a mean reaction time (RT) of 51.04 sec ($CI_{95} = [14.17 85.94]$, $Mdn = 13.33$) and mean confidence rating of 3.93 ($CI_{95} = [3.89 3.97]$). The longer RTs compared to laboratory memory tasks can be accounted for by the interactive nature of the map presented at test (i.e., the ability to zoom and pan the map). On average, 4.77 interactions were made for a trial ($CI_{95} = [4.52 5.04]$, $Mdn = 2$). A Bayesian logistic regression did not show evidence for an effect of the number of interactions on accuracy (mean posterior distribution of the coefficient (M_b) = -.005, $CI_{95} = [-.013 .003]$, $ES = 1.0$, $WAIC = 4717.29$).

Higher confidence ratings predicted better accuracy (Bayesian logistic regression, $M_b = .10$, $CI_{95} = [.09 .11]$, $ES = 10.0$, $WAIC = 4726.31$). Additionally, a confidence-accuracy characteristic (CAC) analysis (Mickes, 2015) was conducted to investigate the confidence-accuracy relationship (see Figure 2A). Results showed that the accuracy increased with confidence (confidence 1: $M = .38$, $CI_{95} = [.31 .45]$, confidence 2: $M = .39$, $CI_{95} = [.34 .45]$, confidence 3: $M = .52$, $CI_{95} = [.47 .56]$, confidence 4: $M = .60$, $CI_{95} = [.57 .63]$, confidence 5: $M = .79$, $CI_{95} = [.77 .81]$).

To assess the contributions of different modalities, which was captured during the four-week data collection period, to the participant's choices during the memory test, we employed a Bayesian conditional logit model as the choice alternatives change on a trial by trial basis in our experiment. The interpretation of the results from a conditional logistic regression is different from a logistic regression model. In a traditional logistic regression model, the accuracy of the response (correct/incorrect) would be used as dependent variables (DV) and be predicted by the independent variables (IV) which are properties of the stimulus itself. However, a conditional logistic regression model uses the chosen/not-chosen options on each trial as DVs, and the IVs are characteristics of these options (i.e., distance of the choice to the correct answer), which varies trial by trial and option by option. Therefore, in the current study, if the coefficients of the conditional logistic regression is negative, it indicates that participants tend to choose options that have smaller distance to the correct answer (e.g., options that are spatially closer to the correct answer) and vice versa.

For each predictor (GPS, audio, accelerometry, and time), we calculated the distance of that choice to the target. Each event is deemed to have extended for an hour period so that we capture the general experience of the participant at the time, rather than just the moment of the measurement. Therefore, for GPS, audio, and accelerometry, we took the median of the distances between each of the data points in each event. Figure 2B illustrates the distance calculation between two events for the GPS modality (although the same procedure was employed for the audio and accelerometry modalities). In Event 1, there are three points and in Event 2, two points (note that an Event is defined as a one hour bin, e.g., 9:00:00~9:59:59). We calculated the Euclidean distance between all six pairs of points (the green lines) and then took the median of

those distances. One consequence of this method is that the distance from a target event to itself is not typically zero, although in general it will be lower than the distances to distractor events. For time, direct differences between the options and target were used for the distance measure. Once distances had been calculated they were standardized to ensure that the weights derived from different modalities could be compared to each other. A separate model was constructed for each modality.

Figure 2. Results and an illustration of the distance calculation. (A) Confidence-accuracy characteristic (CAC) analysis plot. Error bars represent ± 1 SEM, and the dotted line represents chance level (.25), (B) An example of how distances were calculated using GPS coordinates. Distance for the audio and accelerometry were calculated in the same fashion, (C) Posterior distributions of the beta weights for the GPS, continuous time, audio in MFCC format (Audio), and accelerometry (Accel), (D) Posterior distribution of the regression weights for the three categorical times.

Figure 2C shows results for the four modalities, where the beta weight for all four models show a statistically significant effect on the participants choice (GPS: $M_{\beta} = -2.76$, $CI_{95} = [-3.48$

-2.16], $ES = 4.18$, $WAIC = 9869.25$; Audio: $M_\beta = -.20$, $CI_{95} = [-.26 \ -.14]$, $ES = 3.33$, $WAIC = 7864.59$; Accelerometry: $M_\beta = -.07$, $CI_{95} = [-.13 \ -.002]$, $ES = 1.09$, $WAIC = 8665.16$; Time: $M_\beta = -.81$, $CI_{95} = [-.86 \ -.76]$, $ES = 16.2$, $WAIC = 8879.52$). The weights are negative because small distances indicate a higher probability of confusion. While GPS dominates, the 95% credible interval for the audio, accelerometry, and time modalities do not include zero, which suggests that they also have significant effects on participant's choice performance.

Ronald Cotton mistakenly retrieved information about the week before the crime. To test for the influence of these kinds of errors, we constructed categorical time predictors. The hour predictor was zero when the two events occurred on the same hour of the day (but potentially on different days). Otherwise, it was one. The day predictor was zero when the two events occurred on the same day of the week (but potentially in different weeks), and the week predictor was zero when an event occurred in the same week without reference to day or hour. Note in this case, we did not standardize the predictors as they are binary values. Figure 2D shows the means of the beta weights along with their 95% credible intervals. The weight associated with the day of week predictor shows a strong effect suggesting the error that Cotton made is not uncommon ($M_\beta = -1.59$, $CI_{95} = [-1.66 \ -1.51]$, $ES = 21.2$, $WAIC = 8213.14$). However, hour of day and week predictors also show substantial effects on the choice behaviour (Week: $M_\beta = -1.49$, $CI_{95} = [-1.57 \ -1.41]$, $ES = 18.63$, $WAIC = 8630.15$; Hour: $M_\beta = -1.48$, $CI_{95} = [-1.56 \ -1.40]$, $ES = 18.5$, $WAIC = 8721.74$).

One caveat of the examining categorical time as a predictor is that the predictors are confounded with continuous time. Therefore, it is unclear if categorical time alone is driving the effect, or whether there is an additional effect from continuous time. To resolve the issue, we

took a model comparison approach and contrasted a model that included all categorical time variables to a model that included all categorical time variables and the continuous time variable. Results for the categorical time only model showed statistically significant effect for all predictors (Week: $M_\beta = -.56$, $CI_{95} = [-.67 \text{ } -.46]$, $ES = 5.33$; Day: $M_\beta = -.92$, $CI_{95} = [-1.02 \text{ } -.81]$, $ES = 8.76$; Hour: $M_\beta = -.59$, $CI_{95} = [-.69 \text{ } -.48]$, $ES = 5.62$, $WAIC = 7864.40$). The model also showed a lower WAIC value than the models with individual predictors, which indicates it is better at predicting participant's choice even after considering there are two more predictors. Results for the categorical time and continuous time model only showed a statistically significant effect for the categorical time predictors (Week: $M_\beta = -.56$, $CI_{95} = [-.71 \text{ } -.41]$, $ES = 3.73$; Day: $M_\beta = -.91$, $CI_{95} = [-1.01 \text{ } -.81]$, $ES = 9.1$; Hour: $M_\beta = -.58$, $CI_{95} = [-.69 \text{ } -.48]$, $ES = 5.52$; Continuous time: $M_\beta = -.004$, $CI_{95} = [-.08 \text{ } .07]$, $ES = .05$, $WAIC = 7866.43$). The WAIC value was also higher than the categorical time only model, which implies that there is no additional contribution from the continuous time that predicts the participant's choice during the task. Therefore, it is more likely that categorical time is driving the effect.

When one considers the errors responses, the probability of choosing the right day of the week, but the wrong week was 0.19 ($CI_{95} = [.16, .22]$) compared with a chance rate of .14. The probability of choosing the right hour of the day, but the wrong day was .08 ($CI_{95} = [.06 \text{ } .09]$) compared with a chance rate of .04.

Discussion

The present study utilized a naturalistic experience sampling technique to examine what information people employ when attempting to access their memory for *where* given a time cue. The first interesting result from this study is that the average error was 36%. Although we

examined participants who mostly had a regular schedule (e.g., university students), this is a significant proportion, particularly considering that in the experimental situation they were given only four alternatives from which to choose. Our confidence rating results (see Figure 2A) also support the growing body of work suggesting that under well controlled testing conditions, people's confidence in the accuracy of their retrievals is reliable (Wixted & Wells, 2017; Wixted, Mickes & Fisher, 2018; Mickes, 2015).

In our task, the cue was a time and the required target was a location. Consistent with current models of memory (Shiffrin & Steyvers, 1997; Clark & Gronlund, 1996; Osth & Dennis, 2015), we found that people were more likely to incorrectly choose alternatives that were similar in time or location. Similarity of the audio environment or movement patterns produced significant effects, but of much smaller magnitude.

Prior literature suggests that people do not retain a continuous measure of time, but rather encode categorical representations of time in memory (Friedman & Wilkins, 1985; Yim, Garrett, Baker & Dennis, 2018). Similarly, we found a categorical model provided a better model fit than a continuous model. Participants mistakenly chose the right day of the week but the wrong week, 19% of the time and mistakenly chose the right hour of the day but the wrong day 8% of the time. Both of these figures are well above chance rates and suggest that the error that Ronald Cotton made might be quite common.

Technologies that allow us to passively track participants' experience continually and comprehensively provide the possibility for a step change in memory science. We are now able to capture the ground truth of people's lives and to test their ability to retain those events. We can start to provide quantitative answers to questions of applied importance by identifying what the

most likely mistakes are - a precursor to understanding how interrogative techniques could be modified to identify when these mistakes have occurred. Most critically, we can develop models so that we can better capture the similarity structure of encoding and how pre-experimental interference determines people's ability to remember.

Author Contributions

All authors contributed to developing the study concept, and study design. Testing and data collection were performed by EL and BS. Data analysis and interpretation was performed by HY and SJD. EL and HY drafted the manuscript and BS and SJD provided critical revisions. All authors approved the final version of the manuscript for submission. HY and EL contributed equally to this work as a co-first author.

Reference

- Chow, T. E., & Rissman, J. (2017). Neurocognitive mechanisms of real-world autobiographical memory retrieval: insights from studies using wearable camera technology. *Annals of the New York Academy of Sciences*, 1396(1), 202-221.
- Clark, S. E., & Gronlund, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review*, 3(1), 37-60.
- 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 G. Bunzelmann, A. Howes, T. Tenbrink, E. Davelaar (Eds.) *Proceedings of the 39th*

- Annual Conference of the Cognitive Science Society*, (pp. 295-300). Austin, TX: Cognitive Science Society.
- Dennis, S., Yim H., Garrett, P., Sreekumar, V., & Stone, B. (2019). A system for collecting active and passive experience sampling data. *Behavior Research Methods*, 51, 1824-1838.
- Dennis, S., Garrett, P., Yim, H., Hamm, J., Osth, A., Sreekumar, V., & Stone, B. (2019). Privacy versus open science. *Behavior Research Methods*, 51, 1839-1848.
- Friedman, W. J., & Wilkins, A. J. (1985). Scale effects in memory for the time of events. *Memory & Cognition*, 13 (2), 168–175.
- Mermelstein, P. (1976). Distance measures for speech recognition, psychological and instrumental. In C. H. Chen (Ed.) *Pattern Recognition and Artificial Intelligence*, (pp. 374-388). New York: Academic.
- Mickes, L. (2015). Receiver operating characteristic analysis and confidence–accuracy characteristic analysis in investigations of system variables and estimator variables that affect eyewitness memory. *Journal of Applied Research in Memory and Cognition*, 4(2), 93–102.
- Nielson, D. M., Smith, T. A., Sreekumar, V., Dennis, S., & Sederberg, P. (2015). Human hippocampus represents space and time during retrieval of real-world memories. *Proceedings of the National Academy of Sciences of the United States of America*, 112(35), 11078-11083.
- Osth, A. F., & Dennis, S. (2015). Sources of interference in item and associative recognition memory. *Psychological Review*, 122(2), 260.

- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4(2), 145-166.
- Sreekumar, V., Nielson, D. M., Smith, T. A., Dennis, S. J., & Sederberg, P. B. (2018). The experience of vivid autobiographical reminiscence is supported by subjective content representations in the precuneus. *Scientific Reports*, 8(1), 14899.
- Thompson-Cannino, J., Cotton, R., & Torneo, E. (2009). *Picking Cotton: Our memoir of injustice and redemption*. New York: St. Martin's Press.
- Wixted, J. T., Mickes, L. & Fisher, R. P. (2018). Rethinking the Reliability of Eyewitness Memory. *Perspectives on Psychological Science*, 13, 324-335.
- Wixted, J. T. & Wells, G. L. (2017). The relationship between eyewitness confidence and identification accuracy: A new synthesis. *Psychological Science in the Public Interest*, 18, 10-65.
- Yim, H., Garrett, P. M., Baker, M., & Dennis, S. J. (2018). Examining the Independence of Scales in Episodic Memory using Experience Sampling Data, In Kalish, C., Rau, M., Zhu, J., & Rogers, T. T. (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 2726-2731). Austin, TX: Cognitive Science Society.