ARL-TR-9856 • Dec 2023

DEVCOM
ARMY RESEARCH LABORATORY

A Multi-Timescale Analysis of Reward Functions Learned from Human-Automation Collaboration

by Torin Adamson, Alexander Danvers, Javier Mendoza, Yazied Hansan, Andrew Campbell, Matthias Mehl, Lydia Tapia, and Evan C Carter

NOTICES

Disclaimers

ARL-TR-9856 • DEC 2023

A Multi-Timescale Analysis of Reward Functions Learned from Human-Automation Collaboration

Torin Adamson, Yazied Hansan, and Lydia Tapia
University of New Mexico

Andrew Campbell
Dartmouth College

Javier Mendoza and Matthias Mehl
University of Arizona

Alexander Danvers
Sierra Tucson

Evan C Carter
DEVCOM Army Research Laboratory

REPORT DOCUMENTATION PAGE						Form Approved OMB No. 0704-0188
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**						
1. REPORT DATE *(DD-MM-YYYY)* December 2023	**2. REPORT TYPE** Technical Report					**3. DATES COVERED (From - To)** 1 October 2022 – 30 September 2023
4. TITLE AND SUBTITLE A Multi-Timescale Analysis of Reward Functions Learned from Human–Automation Collaboration						**5a. CONTRACT NUMBER** **5b. GRANT NUMBER** **5c. PROGRAM ELEMENT NUMBER**
6. AUTHOR(S) Torin Adamson, Alexander Danvers, Javier Mendoza, Yazied Hansan, Andrew Campbell, Matthias Mehl, Lydia Tapia, and Evan C Carter						**5d. PROJECT NUMBER** **5e. TASK NUMBER** **5f. WORK UNIT NUMBER**
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) DEVCOM Army Research Laboratory ATTN: FCDD-RLH-FA Aberdeen Proving Ground, MD 21005						**8. PERFORMING ORGANIZATION REPORT NUMBER** ARL-TR-9856
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)						**10. SPONSOR/MONITOR'S ACRONYM(S)** **11. SPONSOR/MONITOR'S REPORT NUMBER(S)**
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited						
13. SUPPLEMENTARY NOTES ORCID ID: Evan C Carter, 0000-0001-7471-8769						
14. ABSTRACT We have developed a system combining mobile research video games with mobile, continuous sensing to study human–autonomy teaming. We propose that two of the primary benefits of such an approach, relative to the more traditional, laboratory-based approach, are that more people can be assessed for longer periods of time. An implication of this proposal is that analyses of the multi-timescale properties of human–autonomy teaming should reveal potentially valuable insights. Therefore, we conducted a preliminary analysis of 45 subjects who have completed a 180-day study. We operationalized human–autonomy teaming in terms of a reward function and assessed the ways in which reward functions varied at multiple timescales and how mobile sensing data streams related to decision-making at different timescales. Our findings clearly indicated the presence of individual differences in the importance of multiple timescales, suggesting that, indeed, the traditional laboratory approach to studying human–autonomy teaming may overlook important information. We argue that future work should find ways to combine the strengths of approaches such as ours with the strengths of more traditional data collection techniques.						
15. SUBJECT TERMS human–autonomy teaming, mobile sensing, adaptive autonomy, mobile video games, ambulatory assessment, multi-timescale analysis, Humans in Complex Systems						
16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**		**19a. NAME OF RESPONSIBLE PERSON** Evan C Carter
a. REPORT Unclassified	**b. ABSTRACT** Unclassified	**c. THIS PAGE** Unclassified	UU	24		**19b. TELEPHONE NUMBER (Include area code)** (240) 478-9295

Standard Form 298 (Rev. 8/98)

Contents

List of Figures

1. Introduction

Human–automation interaction studies in traditional laboratory settings provide robust data on specific phenomena through tight control of the environment and manipulation of experimental parameters. However, human behavior around automation in daily life introduces factors that cannot be captured in conventional single-session studies. Human behavior outside the lab can change over long periods of time, possibly in response to changes in the environment, ambulatory psychophysiology, and other measures of the human context. By combining ubiquitous mobile devices and wearable technology widely available to the consumer public, studies conducted in this new "in the wild" approach are able to compare long-term behavior while continuously recording human factors. This provides context to the observed behavior as participants are able to more freely interact with the game and influence—or be influenced—by their environment. An ongoing study combining "Busy Beeway," a platform for studying human–automation interaction[1] and "StudentLife" for recording human context[2] has had 45 participants complete the study thus far. A total of 5915 daily sessions were recorded—over 16 years worth of data.

One of the primary motivations for developing our mobile data collection system is to produce knowledge that can be generalized to more traditional laboratory settings—that is, finding a way to shore up the limitations of traditional data collection through the strengths of our approach. To make generalization more likely, in this work, we operationalize human–automation collaboration behavior in terms of a model of probability discounting,[3] the parameters of which represent psychological processes that should be general to human decision-making. Specifically, we learn two parameters, h and β, which we collectively refer to as a reward function, that reflect the ways in which uncertainty influences human–automation collaboration.

This technical report presents an initial survey of variability in reward functions over long periods of time (e.g., months) and an investigation into how much of the variability is explained by contextual features. These results represent, to our knowledge, the first assessment of human–automation collaboration of its kind, and suggest future directions for the long-term, out-of-lab approach to studying this important topic.

2. Background

Existing studies seek to validate specific human–automation system designs for vehicle automation in specialized lab environments. Most of these use realistic simulators complete with a physical driver seat and steering wheel.[4–7] These studies validate systems that maintain appropriate levels of trust,[4,8] methods of non-verbal communication,[5] indirect shared control systems,[6] and systems that handle events where the human must take control.[7,9] Some studies use a closed course with test vehicles, thus moving beyond simulators,[8] while others forgo the use of autonomy, having human experimenters pose as automation,[10] for example. These highly controlled lab environments provide detailed data into their respective topics. In contrast, our approach examines changes in human behavior or performance around autonomous agents arising from external factors in their lives. For that, studies on each participant must extend beyond single sessions to capture this evolving behavior.

Conducting studies with many participants over long periods of time is infeasible for conventional approaches. Methods such as crowdsourcing can address this by allowing participants to complete the study using their own computers or mobile devices. For human–automation-related studies, this has been done to study the legibility of autonomous movement,[11] and for a collaborative motion planning Mars landing task.[12] Adapting the study into a video game format, known as "gamification," can make the study more accessible and provide intrinsic motivation to help keep participants involved.[13] Gamification has seen success in protein folding prediction,[14] image labeling for machine learning,[15] and in robotic swarm control.[16] Busy Beeway gamifies human–automation collaboration,[1] and this work adapts it for longitudinal studies by collecting multiple daily sessions from each participant.

Studies that employ daily sessions capture only behavioral trends from participants. However, methods that combine smartphones and consumer-grade wearable devices capture human context data (stress, heart rate, physical activity, sleep duration, etc.)[17] The StudentLife[18] application framework integrates mobile sensing with Garmin smart watches to collect such features. To do so, data streams from the Garmin are streamed directly to the phone and later downloaded for analysis, resulting in high-quality data that would be unavailable through Garmin alone. In previous studies, StudentLife has been used to examine behaviors of first-generation students,[19] job performance,[20,21] life events,[22] and student's life during the COVID-19 pandemic.[23] For example,[20] data collected by StudentLife shows that duration of high stress and average heart rate at the desk changes during job promotions. In our work, we demonstrate how data from StudentLife and Busy Beeway are combined to form a more complete understanding of how and why human interactions with automation change over time.

As mentioned, we draw on the concept of probability discounting to produce a reward function that reflects the behavior of participants. Probability discounting describes the way in which the objective value of an outcome is discounted by the probability of receiving it. When measured in humans, a questionnaire form is often used that repeatedly asks participants to choose between two options, one of which is objectively worth more than the other but associated with some degree of probability of not being received. For example, a choice between \$5.00 for certain or a 75% chance of \$10.00. In these studies, it is common that a discount factor is extracted for each individual, and variation in these values are compared with other constructs of interest, typically related to well-being. For example, a meta-analysis from 2018 synthesized 12 studies and concluded that the tendency to overvalue low-probability gains (or undervalue high-probability losses) was associated with problem gambling behavior.[24] Discount factors, inferred in a variety of ways, have also been associated with cigarette smoking,[25] mood and psychotic disorders,[26] elements of alcohol use disorder,[27] and risky sexual behavior.[28]

Broadly speaking, the logic in the referenced literature is that choices on probability discounting tasks reflect a general psychological process also at play in the real world for decisions of real consequence. Likewise, based on such logic, we hypothesize that reward functions related to probability discounting will meet our purpose of finding a reward function that may be relevant across human–automation collaboration tasks. Of course, many researchers in the area of probability discounting are quick to point out that such associations are sensitive to the specifics of the choice problem, likely to change over time, and vary across individuals. We therefore also expect the same to be true of any discounting parameters that reflect behavior in human–automation collaboration.

Notably, although longitudinal assessments of decision-making in the form of discount factors exist in the literature (e.g., Olson 2010),[29] they are rare and tend to involve very few measurements through time, often only a single time point and a single follow-up. Our data are unique in that they offer far longer time series, which allows for a novel view of decision-making from the perspective of a complex systems approach. Specifically, our data allow us to investigate whether long-range correlations exist in inferred discount factors, as well as how the influence of contextual time features on discount factors changes over timescale, and importantly, which timescales are most relevant. To our knowledge, although such research questions have been investigated in several human science domains (e.g., time interval estimation,[30] reaction time,[31] and steering behavior[32]), they have never been tested in inferred discount factors, let alone discount factors that characterize human–automation collaboration. Answering these research questions will set the stage for future work to develop models that predict general reward functions in human–automation collaboration, potentially leading to personalized models of individual people that autonomous teammates can use to improve teaming outcomes.

3. Methods

3.1 Study Protocol

The data in this work come from an ongoing human subject study. Sixty participants were recruited at the University of Arizona. Following an initial screening process, individual participants were invited to the lab for onboarding, where the data collection system was explained to the participants and the BusyBeeway and StudentLife applications were loaded onto their personal smartphones. Participants completed a series of self-report items and a short tutorial on the gameplay. Participants were then asked to play once daily when prompted for the next 180 days. Other than troubleshooting, this was the only in-person interaction that the participants had with the study team. Additional details of the study protocol can be found in Adamson et al.[33]

3.2 BusyBeeway

BusyBeeway is a mobile research video game designed for the study of human–automation collaboration for motion planning.[1] The game involves navigation and collision-avoidance in the face of stochastic, moving obstacles. The application version used in this study collects both game play and survey response data and allows the experimenter to push modifications of game parameters to active participants as needed. Participants for this study were given instructions during an onboarding process along with the link to download and install the app via Google Play as a closed beta. In each daily gameplay session, players were assigned one of four AI agents that vary in priorities between quick-level completion and obstacle avoidance. Further detail on gameplay can be found in Adamson et al.[33]

3.3 Data Collection

The game application records all raw screen input and all information about the game objects (position, heading, etc.). Data is recorded at 30 Hz, which matches the per-second frame rate the game maintains. In addition to the data collected through the game application, a participant's context was also measured using a smartwatch (Garmin Vivoactive 4) and a suite of passive sensing software uploaded to their devices. The sensing software is known as StudentLife[2,18] and is used to collect data via a mobile phone's onboard sensors (e.g., GPS). In this study, StudentLife was modified to also coordinate the collection of data from the Garmin, which includes a variety of sensors: an accelerometer, GPS, and a heart-rate monitor. Further details on data collection can be found in Adamson et al.[33,34]

For the analyses discussed here, the relevant data streams are average daily heart rate (beats per minute), average daily respiration (breaths per minute), total steps taken, average daily Garmin stress score (a function of heart rate variability) and average light detected by the phone's camera (where units are a device-dependant scalar). There are many other ways that the contextual data streams could be represented, and indeed, the day-based average is likely the least informative. For example, variables such as activity and heart rate would clearly show multi-timescales properties in that they vary with waking and sleep periods in a 24-h period. For this report, we consider our analyses to be an initial proof of concept, and future analyses will include a more systematic effort at feature engineering.

3.4 Data Analysis

3.4.1 Probability-Discounting-Based Reward Function

Participants make many risk–reward judgments during gameplay, evaluating the distance closer to the next goal the player character could move against the probability of avoiding collisions. When the participant takes control away from the AI agent, they are rejecting (their estimate of) the agent's risk–reward trade-off in favor of their own. This implies a kind of two-alternative forced choice, analogous to the kinds of choices on which probability discounting methods are based.[3] Said differently, the observed behavior of the player when in control can be the revealed preference of the player. In probability discounting, the value V of an option being considered can be expressed in Eq. 1 where A is the reward in terms of moving closer to the goal in game units and θ is the *odds against* receiving this reward in terms of p, the probability of avoiding collision. We define the inferred value of the player's behavior as v_p and the inferred value of the (estimated) agent's behavior as v_a.

$$V = \frac{A}{1 + h\Theta} \quad \Theta = \frac{1-p}{p} \tag{1}$$

$$\pi_i = \frac{1}{1 + exp(-\frac{V_p - V_a}{\beta})} \tag{2}$$

There are many ways one could infer values of h. Here, we opted for a maximum likelihood approach and adopted the logistic function in Eq. 2 as our likelihood function.[35] This function relates the difference in the value of two options, V_p and V_a, both calculated with Eq. 1 using the same h value, to the probability of picking V_p. As the value of the two options becomes more distinct—that is, one choice becomes more obviously valuable than the other—the probability of picking it should increase. This choice of a likelihood function also introduces the parameter, β, which scales how quickly the probability of picking V_p increases as a function of its relative value to V_b. As β increases, the probability of picking the less valuable option increases, suggesting either a kind of exploration–exploitation trade-off, in which the decision-maker knowingly picks the less valuable option to gain information about it, or a perceptual noise-like process in which discriminating between the two options is difficult, and thus, when values are close, the lower valued option can be picked by mistake.

A value of $h = 1$ would indicate a decision-maker that chooses the option with higher expected value, and therefore, would be considered strictly rational. In the context of BusyBeeway, $h > 0$ indicates a decision-maker choosing to move further from the goal than the AI agent would for the sake of reducing the probability of collision. When $h < 0$, the decision-maker has opted for a path that results in both an increased chance of collision with an obstacle and increase in the closeness to the goal as compared to what the agent would have done. Likewise, as β

increases, the decision-maker appears to be less strictly rational in the sense previously described—either valuing exploration or appearing to have bounds on their rationality related to perception. We refer to these two parameters collectively as the reward function because they relate attributes of an option—that is, the closeness to the goal and the probability of collision for a course of action—to behavior, where higher reward for an option is proportional to higher probability of picking that option.

To estimate the reward function, we calculated the likelihood, π_i, over the range $-20 \leq h \leq 20$ and $1 \leq \beta \leq 200$ for each participant for each day in which we observed player control. The values of h and β that maximized π_i on each day were taken as our estimate of the reward function.

3.4.2 Detrended Fluctuation Analysis and Multi-Timescale Regression Analysis

The resulting time-series data from recording human context measurements and estimating behavioral parameters h and β form a complex noise-like signal. To investigate whether or not effects at different timescales could have an impact on gameplay behavior, multi-timescale analysis is required. This research uses two methods: detrended fluctuation analysis (DFA) to describe the structure of time series h and β and multi-timescale regression analysis (MRA) to measure the strength of the regression coefficients at the different scales.

To perform DFA, the participant's h and β time-series data were converted into a walk-like series by accumulating the difference between successive samples. Then the signal is divided into windows at a particular timescale w in terms of the number of samples per window. Within each window, a fit is found (linear in this study) and subtracted from the signal to detrend it. The root-mean squared (RMS) variation is calculated for each window and then the RMS of those values represent the average fluctuation at the timescale w. This is repeated for all timescales of interest. A linear regression between timescale and RMS is found in logarithmic space, the slope of which yields the Hurst exponent H that describes the fractal structure of the signal. For comparison, white noise has a Hurst exponent of 0.5 and brown noise (random walk) has an exponent of 1.5.

$$\hat{r} = \frac{\sigma_x}{\sigma_y} r \tag{3}$$

The MRA method used in this research is similar; both the time series of measurements x and the behavioral series y (either h or β) are normalized to remove units and then detrended at each timescale as in DFA. Instead of calculating the variance at each window, a linear regression is performed to find the correlation coefficient r. This coefficient is standardized using Eq. 3 where σ_x is the standard deviation for x in the window and σ_y is the standard deviation for y. Because values for $\hat{r}$ can become positively biased in multi-timescale methods, a quadratic local detrending method is used here instead.[36] Finally, the $\hat{r}$ values for each window are averaged to produce an $\hat{r}$ for that timescale. The strength of this regression is reported as R^2, the coefficient of determination for each timescale.

4. Results and Discussion

4.1 Univariate Detrended Fluctuation Analysis

The results from applying DFA to h and β time series are displayed in Fig. 1. The average estimate of the Hurst exponents are 0.71 and 0.65 for h and β, respectively, indicating temporally correlated noise, or a fractal structure. This is consistent with many behavioral time series produced by humans.[32]

Hurst exponent estimates per individual have a mean of 0.685 (standard deviation 0.158) for h and a mean of 0.638 (standard deviation 0.176) for β. This implies a degree of variation, with some individuals seeming to show behavior consistent with anti-correlated noise and some possibly showing behavior consistent with multifractal noise. In general, however, the appearance of long-range temporal correlations is clearly the norm.

4.2 Multiscale Regression Analysis

The results of our MRA analyses are shown in Figs. 2 and 3 for h and β, respectively. Each panel shows the scale-wise associations of a reward function parameter and a context data stream (i.e., heart rate, respiration, steps, stress, and light). The darker lines represent results from all data and the bands represent the values that contain 1 standard deviation.

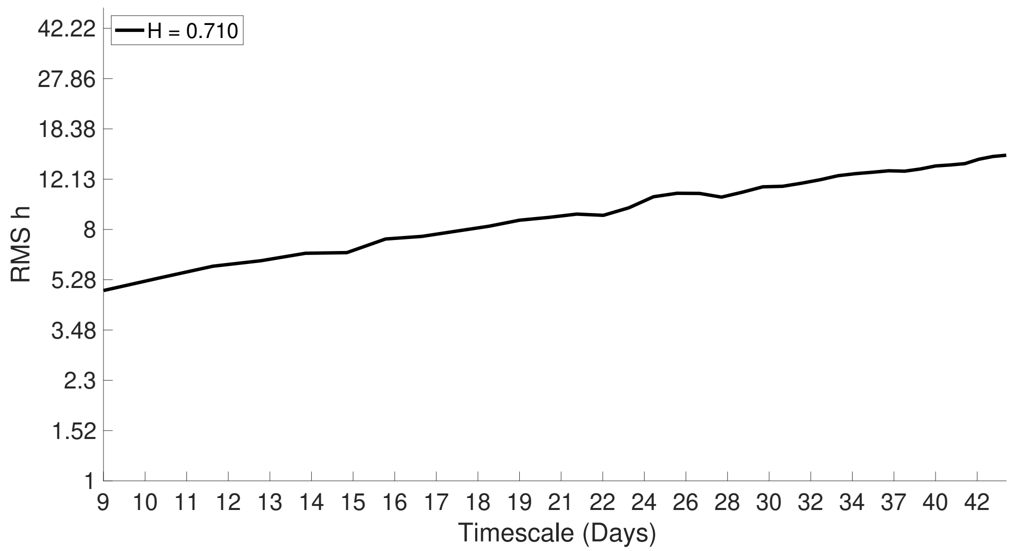

(a) h

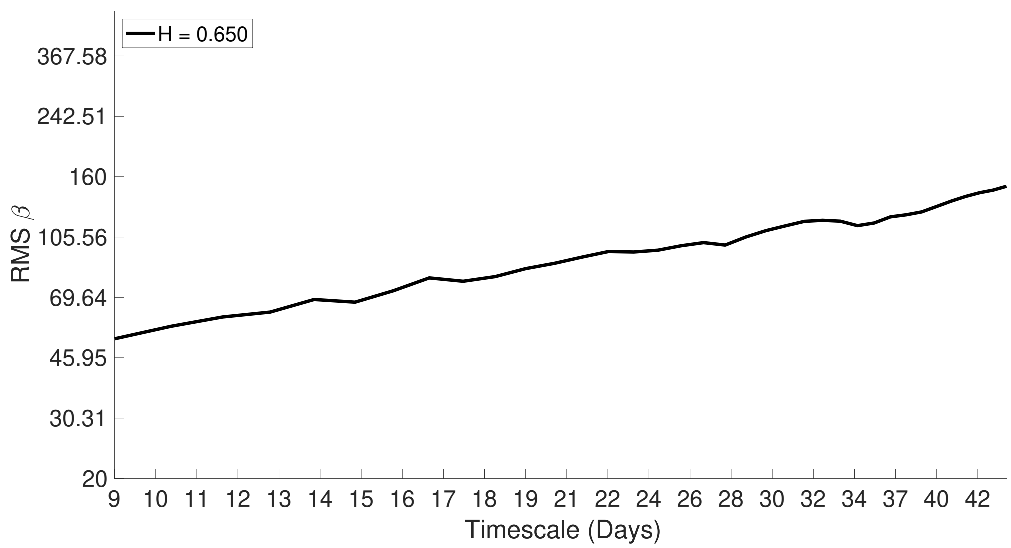

(b) β

Fig. 1 DFA of variables h and β in risk discounting. A base 2 logarithmic scale is applied to both axes and the Hurst exponent H is listed at the top left of each graph. The line indicates the RMS variation in local fluctuations for that timescale.

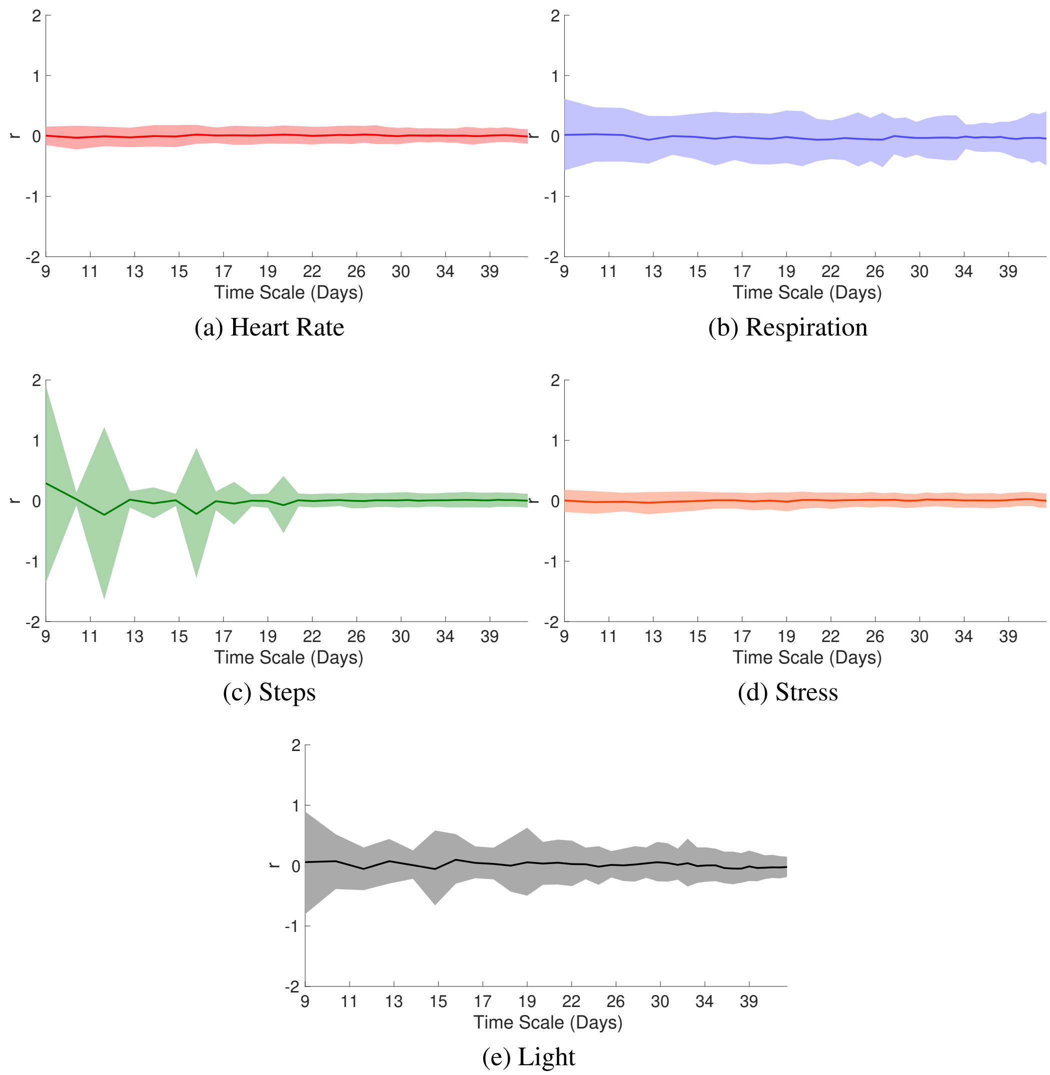

Fig. 2 MRA of several human context measurements and their regression coefficients r against variable h in risk discounting. The measurements are daily measures of (a) average heart rate (in beats per minute), (b) average respiration (in breaths per minute), (c) total steps taken, (d) average stress according to Garmin, and (e) average light detected as a device-dependant scalar. The line indicates the overall average while the shaded region indicates 1 standard deviation assuming symmetry to illustrate magnitude.

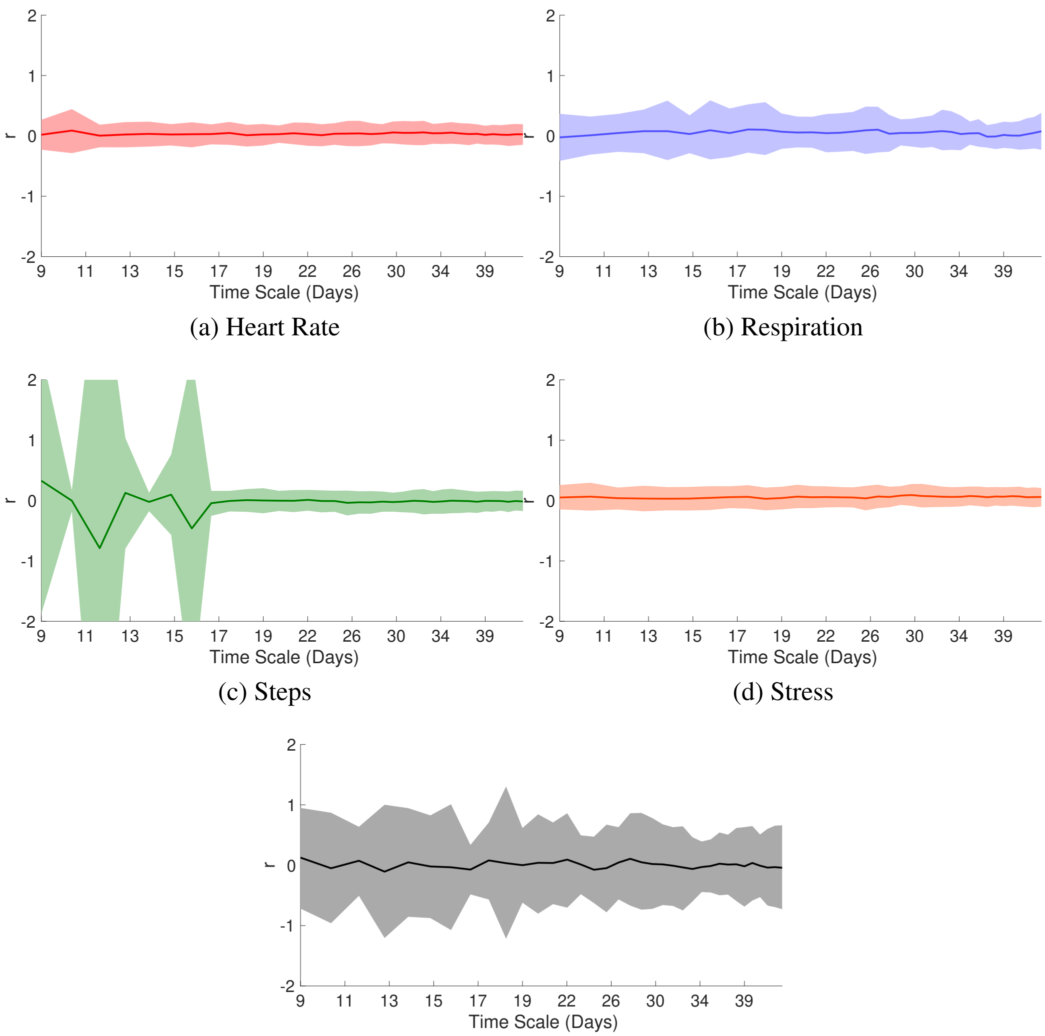

(a) Heart Rate

(b) Respiration

(c) Steps

(d) Stress

(e) Light

Fig. 3 MRA of several human context measurements and their regression coefficients r against variable β in risk discounting. The measurements are daily measures of (a) average heart rate (in beats per minute), (b) average respiration (in breaths per minute), (c) total steps taken, (d) average stress according to Garmin, and (e) average light detected as a device-dependant scalar. The line indicates the overall average while the shaded region indicates 1 standard deviation assuming symmetry to illustrate magnitude.

Figures 2 and 3 demonstrate how the overall association of each context variable with each reward function parameter generally stays around zero, regardless of timescale. However, for some signals, there is potentially interesting variation among individuals. For example, the association between steps and both h and β has several peaks in variance at around 9, 12, and 16 days. This indicates that, for some people, assessing longer timescales could lead to capturing important connections between context and behavior that would be overlooked by traditional modeling efforts that focus only on single timescales or averages of individuals.

5. Conclusions

The results from our application of DFA indicate that, much like many other human phenomena, the reward functions we inferred from BusyBeeway appear to show fractal behavior. Although controversial, the typical interpretation of such results in the literature would be that fractal behavior suggests a complex system, comprising multiple interacting processes, that is showing the potential to respond to constraints by striking a balance between order and disorder (what is often referred to as flexibility). Notably, as described by Likens et al. (2025),[32] "Behaviors heavily constrained by either task or intention show weakened fractal properties, while those left free to vary do not" (p. 2815). From this perspective, the individual differences in fractility observed in our data may suggest that participants differed in the degree to which their preferences remained consistent over time. It would seem reasonable that such variation could be traced back to variation in context.

Our assessment of how reward function parameters relate to contextual data streams further highlights the importance of individual differences. Specifically, when averaged over all subjects, the association between all context variables and both reward function parameters tends to remain near zero, suggesting minimal, if any influence. However, for some signals for some people, it appears that the association shows up more clearly, but only at certain timescales (i.e., short- to mid-range timescales such as 9, 12, and 16 days for steps).

In total, our findings support the argument that the traditional laboratory-based study potentially loses out on critical information by being limited in both the length of time data can be collected and the number of people that can be sampled. The question then becomes whether large data sets that span long timescales, such as ours, can be productively used in conjunction with the more specialized, traditional data sets associated with laboratory-based studies of human–automation collaboration. This is the direction our work will take next.

6. References

1. Adamson T, Oishi M, Chiang HTL, Tapia L. Busy Beeway: a game for testing human-automation collaboration for navigation. In: ACM Proc Intl Conf Motion Games (MIG). 2017 Nov. Barcelona, Spain. 2017. p. 9:1–9:6.

2. Wang R, Chen F, Chen Z, Li T, Harari G, Tignor S, Zhou X, Ben-Zeev D, Campbell AT. StudentLife: using smartphones to assess mental health and academic performance of college students. In: Mobile Health. Springer; c2017. p. 7–33.

3. Shead NW, Hodgins DC. Probability discounting of gains and losses: implications for risk attitudes and impulsivity. J Exp Anal Behav. 2009;92(1):1–16.

4. Yin W, Chai C, Zhou Z, Li C, Lu Y, Shi X. Effects of trust in human-automation shared control: a human-in-the-loop driving simulation study. In: IEEE Intl Int Transp Sys Conf. 2021. p. 1147–1154.

5. Cutlip S, Wan Y, Sarter N, Gillespie RB. The effects of haptic feedback and transition type on transfer of control between drivers and vehicle automation. IEEE Trans Hum-Mach Sys. 2021;51(6):613–621.

6. Li R, Li Y, Li SE, Zhang C, Burdet E, Cheng B. Indirect shared control for cooperative driving between driver and automation in steer-by-wire vehicles. IEEE Trans Intl Transp Sys. 2021;22(12):7826–7836.

7. Huang C, Lv C, Hang P, Hu Z, Xing Y. Human–machine adaptive shared control for safe driving under automation degradation. IEEE Int Transport Sys Magazine. 2022;14(2):53–66.

8. Xu A, Dudek G. Maintaining efficient collaboration with trust-seeking robots. In: Proc IEEE Intl Conf Intel Rob Syst (IROS). 2016. p. 3312–3319.

9. de Waard D, van der Hulst M, Hoedemaeker M, Brookhuis KA. Driver behavior in an emergency situation in the automated highway system. Trans Hum Fact. 1999;1(1):67–82.

10. Johns M, Mok B, Sirkin DM, Gowda NM, Smith CA, Talamonti Jr WJ, Ju W. Exploring shared control in automated driving. In: ACM Intl Conf Hum Rob Interaction (HRI). 2016 Mar. Christchurch, New Zealand. 2016. p. 91–98.

11. Dragan AD, Lee KC, Srinivasa SS. Legibility and predictability of robot motion. In: ACM Intl Conf Hum Rob Interaction (HRI). 2013. p. 301–308.

12. Srivastava DK, Lilly JM, Feigh KM. Improving human situation awareness in AI-advised decision making. In: Improving Human Situation Awareness in AI-Advised Decision Making. 2022. p. 1–6.

13. Eveleigh A, Jennett C, Blandford A, Brohan P, Cox AL. Designing for dabblers and deterring drop-outs in citizen science. In: Proc SIGCHI Conf Hum Factors Comput Syst (CHI). 2014 Apr. Toronto, Ontario, Canada. 2014. p. 2985–2994.

14. Cooper S, Khatib F, Treuille A, Barbero J, Lee J, Beenen M, Leaver-Fay A, Baker D, Popovic Z, Players F. Predicting protein structures with a multiplayer online game. Nature. 2010;466:756–760.

15. von Ahn L, Dabbish L. Labeling images with a computer game. In: Proc SIGCHI Conf Hum Factors Comput Syst (CHI). SIGCHI; c2004. 2004. p. 319–326.

16. Becker A, Ertel C, McLurkin J. Crowdsourcing swarm manipulation experiments: a massive online user study with large swarms of simple robots. In: Proc IEEE Intl Conf Rob Auto (ICRA). 2014. p. 2825–2830.

17. Lane N, Rabbi M, Lin M, Yang X, lu H, Ali S, Doryab A, Berke E, Choudhury T, Campbell A. Bewell: a smartphone application to monitor, model and promote wellbeing. In: Proc Intl Conf Pervasive Comput Tech Health (ICST). 2011.

18. Wang R, Chen F, Chen Z, Li T, Harari G, Tignor S, Zhou X, Ben-Zeev D, Campbell AT. StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones. In: ACM Proc Intl Jt Conf Pervasive Ubiquitous Comput (UbiComp). 2014 Sep. Seattle, WA. 2014. p. 3–14.

19. Wang W, Nepal S, Huckins JF, Hernandez L, Vojdanovski V, Mack D, Plomp J, Pillai A, Obuchi M, Dasilva A. First-gen lens: Assessing mental health of first-generation students across their first year at college using mobile sensing. Proc ACM Interact Mobile Wearable Ubiquitous Tech. 2022;6(2):1–32.

20. Nepal S, Mirjafari S, Martinez GJ, Audia P, Striegel A, Campbell AT. Detecting job promotion in information workers using mobile sensing. Proc ACM Interact Mobile Wearable Ubiquitous Tech. 2020;4(3):1–28.

21. Mirjafari S et al. Differentiating higher and lower job performers in the workplace using mobile sensing. Proc ACM Interact Mobile Wearable Ubiquitous Tech. 2019;3(2):1–24.

22. Pillai A, Nepal S, Campbell A. Rare life event detection via mobile sensing using multi-task learning. arXiv preprint arXiv:2305.20056. 2023;.

23. Nepal S, Wang W, Vojdanovski V, Huckins JF, daSilva A, Meyer M, Campbell A. COVID student study: a year in the life of college students during the COVID-19 pandemic through the lens of mobile phone sensing. In: Proc SIGCHI Conf Hum Factors Comput Syst (CHI). 2022. p. 1–19.

24. Kyonka EG, Schutte NS. Probability discounting and gambling: a meta-analysis. Addiction. 2018;113(12):2173–2181.

25. Reynolds B, Richards JB, Horn K, Karraker K. Delay discounting and probability discounting as related to cigarette smoking status in adults. Behavioural Processes. 2004;65(1):35–42.

26. Hart KL, Brown HE, Roffman JL, Perlis RH. Risk tolerance measured by probability discounting among individuals with primary mood and psychotic disorders. Neuropsychology. 2019;33(3):417.

27. Acuff SF, Boness CL, McDowell Y, Murphy JG, Sher KJ. Contextual decision-making and alcohol use disorder criteria: delayed reward, delayed loss, and probabilistic reward discounting. Psychology of Addictive Behaviors. 2023;37(1):121.

28. Johnson MW, Strickland JC, Herrmann ES, Dolan SB, Cox DJ, Berry MS. Sexual discounting: a systematic review of discounting processes and sexual behavior. Experimental and Clinical Psychopharmacology. 2021;29(6):711.

29. Olson EA. Delay and probability discounting: a longitudinal study of neural, cognitive, and emotional processes contributing to adolescent development. [thesis]. 2010.

30. Wagenmakers EJ, Farrell S, Ratcliff R. Estimation and interpretation of 1/fα noise in human cognition. Psychonomic Bulletin and Review. 2004;11(4):579–615.

31. Van Orden GC, Holden JG, Turvey MT. Self-organization of cognitive performance. Journal of Experimental Psychology: General. 2003;132(3):331.

32. Likens AD, Fine JM, Amazeen EL, Amazeen PG. Experimental control of scaling behavior: what is not fractal? Experimental Brain Research. 2015;233:2813–2821.

33. Adamson T, Pillai A, Campbell A, Tapia L, Obregon LS, Carter EC. Long-term analysis of a human-ai collaboration study using a mobile game. DEVCOM Army Research Laboratory (US); 2023. Report No.: ARL-TR-9628.

34. Adamson T, Wang W, Hasan Y, Campbell A, Tapia L, Carter E. A mobile data collection system for studying human autonomy teaming in conjunction with passive context and psychophysiological sensing. DEVCOM Army Research Laboratory (US); 2021. Report No.: ARL-TR-9359.

35. Lee MD. Bayesian methods in cognitive modeling. In: Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience. John Wiley Sons, Ltd; c2018. 2018. p. 1–48.

36. Likens AD, Wiltshire TJ. fractalRegression: an R package for multiscale regression and fractal analyses. https://github.com/aaronlikens/fractalRegression [accessed 2023 Sep 27].

List of Symbols, Abbreviations, and Acronyms

AI	artificial intelligence
DFA	detrended fluctuation analysis
GPS	global positioning system
MRA	multi-timescale regression analysis
RMS	root-mean squared

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

1 DEVCOM ARL
(PDF) FCDD RLB CI
TECH LIB

1 DEVCOM ARL
(PDF) FCDD RLA FA
E CARTER