# Method for Understanding Complex Human Routine Behaviors from Large Behavior Logs

**Nikola Banovic**
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15213, USA
nbanovic@cs.cmu.edu

## Author Keywords
Human Routine Behavior; Inverse Reinforcement Learning; Data Mining; Visual Analytics; Sensemaking.

## ACM Classification Keywords
H.5.m. Information interfaces and presentation (*e.g.*, HCI): Miscellaneous.


CHI'17 Extended Abstracts, May 06-11, 2017, Denver, CO, USA
ACM 978-1-4503-4656-6/17/05.
http://dx.doi.org/10.1145/3027063.3027135

## Abstract
The increasing ability to collect large amounts of human behavior data can inform technology that has the potential to help people improve their behaviors and thus improve the quality of their lives. To design and implement such technology requires understanding of those very behaviors that the technology is trying to diagnose and improve. However, existing methods to explore and make sense of human behaviors are not well suited to address the increasingly large amount of data collected in behavior logs. My research focuses on the domain of human routines where I model behaviors as sequences of actions people perform in specific situations. I leverage those computational models of routines together with different visualization tools to aid researchers and domain experts in exploring, making sense of, and generating new insights about human behavior in a principled way. My research informs the design of technology that helps people be productive, healthy, and safe.

## Context and Motivation
Good routines help people be productive, healthy, and safe in a variety of contexts, ranging from people's daily schedules to their driving routines. When people lack good routines or develop poor routines over time,

their wellbeing is negatively impacted. For example, poor sleeping or exercising routines can negatively impact people's mental and physical health. Routines are made up of behavior instances—frequent and repetitive series of actions that people perform in different situations that are the cause of those actions [10]. Having an understanding of those causal relationships could inform technologies that detect the causes of poor routines and help improve people's lives.

To develop interventions that help people improve their routines requires careful understanding of how different *routine variations*, or different behavior instances characteristic of a routine, affect them and their goals. Such variations are different from *deviations* and other uncharacteristic behaviors that are not part of their routines. For example, to help students balance the demands of their academic life and their wellbeing requires an understanding of how their schedules impact their mental health. To coach them on how they can improve requires understanding how students who are successful and healthy manage their schedules.

To understand human routines requires collecting and studying data about how people behave in their frequent, day-to-day tasks. The recent proliferation of sensors in the environment and people's personal and wearable devices has significantly increased our ability to collect massive behavior log data compared to traditional journal-based data collection. However, unlike behavior data collected in journal-based studies, automatically collected behavior logs do not contain ground truth about people's behavior. Researchers often have to perform tedious manual exploration of such large behavior logs to make sense of the data.

As a result researchers are forced to manually pick only a small subset of available data for a specific statistical analysis to tests a particular preconceived hypothesis. Such approaches may work for *hypothesis testing*, but are difficult to apply to *hypothesis generation* for domain experts that may not be knowledgeable about data analytics approaches. This leaves a massive amount of data from behavior logs underutilized when researchers do not have a hypothesis they would like to test beforehand. To be able to explore and make sense of large behavior logs that we are able to collect nowadays requires new data exploration techniques that help researchers to first understand their data, and then generate and finally test their hypotheses. Thus, the goal of my research can be summarized in the following thesis statement:

*A rich computational model of human routine behavior, that captures causal relationships between situations and actions, can be used to describe, reason about, and act in response to human behavior, stored as event-driven data in large behavior logs. I hypothesize that such a model can aid domain experts and end users in sensemaking about behavior of individuals and populations. I also hypothesize that such a model can be used to automatically detect classes of behavior, such as poor routines, and prescribe changes, such as simulating a better routine.*

## Background and Related Work

The need for methods that help researchers understand human routine behaviors span different stakeholders and domains. For example, designers may build on such knowledge to design technologies that help people improve their routines [6]; individuals may wish to reflect about their own routines for supporting behavior change [9]; clinicians may use such knowledge to explore the effects of treatments on their patients [12].

However, it is challenging to apply existing data exploration methods to understand data from large behavior logs. Using a visualization to manually explore data from such logs (e.g., [13]) does not guarantee that the user will be able to find patterns of behaviors that form routines as the number of data features grows. The lack of ground truth labels in large behavior logs makes it challenging to use existing supervised machine learning algorithms (*e.g.*, [6]) to find behavior instances that are characteristic of a particular routine. This places burden on researchers to manually label enough data (even with semi-supervised methods [5]) to train such algorithms. Unsupervised learning methods (*e.g.*, [8]) cluster behaviors without prior knowledge of labels, but offer no guarantees that the resulting clusters represent routines.

Even when using existing routine models [1, 3, 4, 7, 11] to automatically extract patterns of routine behaviors from behavior logs, manually exploring those individual patterns remains tedious. Furthermore, researchers and domain experts who lack intimate knowledge about such modeling methods face obstacles when trying to combine them together with other approaches to create their own data analysis pipelines. Thus, currently, without an ability to explore routine behaviors, we face significant obstacles in developing technologies that help people improve their routines.

## Research Objectives

My goal is to design and implement a domain-agnostic, mixed-initiative method for exploring and making sense of complex human routine behaviors from large behavior logs. The method consists of a carefully curated pipeline of machine learning and data mining algorithms coupled with visualization tools that guides researchers through their exploration of behavior data to generate and test their hypotheses about behaviors in a principled way. At the highest level, the method helps researchers answer the following questions about routines of an individual or a population to generate their hypotheses:

*Q1: What are the variations that are characteristic of a particular routine?*

*Q2: What are the situations and/or actions that are the cause of those variations?*

*Q3: What are the differences between routine variations and deviations between different individuals and different populations?*

To allow researchers to answer those questions, the interactive pipeline will guide them from pre-processing their log data into behavior instances and automatically extracting routine variations and deviations from the logs for both individuals and populations, to providing them with an exploration-time feature selection method to find which features best describe the behaviors in the logs, and automatically detecting routine variations and deviations to allow them to explore people’s behaviors in different “what-if” scenarios. Once they form their hypotheses about human behaviors through model exploration, the pipeline will provide them with appropriate methods to test those hypotheses.

## Research Approach and Methods

To study how my proposed method aids researchers and domain experts in exploration of their data, I plan to develop a system that implements the pipeline and test it in a combination of lab and field studies in multiple domains. In the lab studies, we intend to evaluate the usability of the system to ensure that researchers across different domains that do not have extensive knowledge of machine learning and data mining can use the system. Field studies will allow me to gather qualitative data about how researchers make use of the system in their everyday data exploration.

## Dissertation Status and Current Results

I have devised an algorithm for modeling routines [1] based on an existing Inverse Reinforcement Learning algorithm [14]. This model is at the core of the proposed pipeline and enables systems that can automatically reason about routines. Results from a lab study showed that the algorithm extracts meaningful patterns of routine behaviors in the domains of daily commutes and driving [1].

I have also shown that my routine models can automatically detect and generate variations that are characteristic of a particular routine [2], which helps explore behavior data when ground truth labels about behaviors are not present in the logs. The model also enables algorithms that can automatically select important behavioral features that can guide behavior exploration. The next step in developing the system is to devise an algorithm that captures the procedural knowledge required to model routines using my approach. I will then implement a visual analytics tool, which incorporates these algorithms, and deploy it in the field to study its effects on researchers' ability to make richer insights about their data when using my method compared to existing data analysis methods.

## Expected Contributions

Aside from generating knowledge about routine behaviors, my method will produce models of routine behaviors that can directly power technologies that help people improve their routines. For example, my work that is currently in submission to a conference shows how aggressive and non-aggressive driving routine models can be used to automatically detect and raise drivers' awareness of their aggressive behavior instances and coach them on how non-aggressive drivers would behave in the same situations. The method is grounded in statistical principles, which ensures it generalizes to multiple domains (*e.g.*, health, fitness, productivity). It is uniquely positioned to increase the intelligibility and accountability of such systems because the algorithms that are at the core of the method help in understanding the behaviors they capture. My routine exploration approach will inform technology that benefits society by helping people improve their wellbeing.

## Acknowledgements

I would like to thank my advisors Jennifer Mankoff and Anind Dey. My research was supported by NSERC (PGSD3-438429-2013).

## References

1. Nikola Banovic, Tofi Buzali, Fanny Chevalier, Jennifer Mankoff, and Anind K. Dey. 2016. Modeling and Understanding Human Routine Behavior. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16). ACM, New York, NY, USA, 248-260. http://dx.doi.org/10.1145/2858036.2858557

2. Nikola Banovic, Anqi Wang, Yanfeng Jin, Christie Chang, Julian Ramos, Anind K. Dey, and Jennifer Mankoff. 2017. Leveraging Human Routine Models to Detect and Generate Human Behaviors. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (CHI '17). ACM, New York, NY, USA. http://dx.doi.org/10.1145/3025453.3025571
3. Oliver Brdiczka, Norman Makoto Su, and James Bo Begole. 2010. Temporal task footprinting: identifying routine tasks by their temporal patterns. In *Proceedings of the 15th international conference on Intelligent user interfaces* (IUI '10). ACM, New York, NY, USA, 281-284. http://doi.acm.org/10.1145/1719970.1720011
4. Paolo Buono, Aleks Aris, Catherine Plaisant, Amir Khella, and Ben Shneiderman. 2005. Interactive pattern search in time series. In *Electronic Imaging 2005*, 175-186. http://dx.doi.org/10.1117/12.587537
5. Olivier Chapelle, Bernhard Schlkopf, and Alexander Zien. 2010. *Semi-Supervised Learning* (1st ed.). The MIT Press.
6. Scott Davidoff, Brian D. Ziebart, John Zimmerman, and Anind K. Dey. 2011. Learning patterns of pick-ups and drop-offs to support busy family coordination. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '11). ACM, New York, NY, USA, 1175-1184. http://doi.acm.org/10.1145/1978942.1979119
7. Nathan Eagle and Alex S. Pentland. 2009. Eigenbehaviors: identifying structure in routine. *Behav. Ecol. Sociobiol.*, vol. 63, no. 7, 1057-1066. http://dx.doi.org/10.1007/s00265-009-0739-0
8. Katayoun Farrahi and Daniel Gatica-Perez. 2012. Extracting Mobile Behavioral Patterns with the Distant N-Gram Topic Model. In *Proceedings of the 2012 16th Annual International Symposium on Wearable Computers (ISWC)* (ISWC '12). IEEE Computer Society, Washington, DC, USA, 1-8. http://dx.doi.org/10.1109/ISWC.2012.20
9. Ian Li, Anind Dey, and Jodi Forlizzi. 2010. A stage-based model of personal informatics systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '10). ACM, New York, NY, USA, 557-566. http://doi.acm.org/10.1145/1753326.1753409
10. Geoffrey M. Hodgson. 1997. The ubiquity of habits and rules. *Cambridge Journal of Economics*, *21*(6), 663-684.
11. Magnus S. Magnusson. 2000. Discovering hidden time patterns in behavior: T-patterns and their detection. *Behavior Research Methods, Instruments, & Computers*, *32*(1), 93-110. http://dx.doi.org/10.3758/BF03200792
12. Adam Perer and David Gotz. 2013. Data-driven exploration of care plans for patients. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '13). ACM, New York, NY, USA, 439-444. http://dx.doi.org/10.1145/2468356.2468434
13. Catherine Plaisant, Brett Milash, Anne Rose, Seth Widoff, and Ben Shneiderman. 1996. LifeLines: visualizing personal histories. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '96), Michael J. Tauber (Ed.). ACM, New York, NY, USA, 221-227. http://doi.acm.org/10.1145/238386.238493
14. Brian D. Ziebart, J. Andrew Bagnell, and Anind K. Dey. 2010. Modeling interaction via the principle of maximum causal entropy. In *Proceedings of the 27th international conference on Machine learning (ICML '10)*. ACM, New York, NY, USA., 1247-1254.