

My research focuses on developing **machine learning methods and data-driven interfaces for understanding and augmenting human behavior**, with the goal of advancing high-impact application areas like health, well-being, computer-mediated communication, and collaboration.

I analyze and model **complex datasets from multiple modalities**, such as smartphone and wearable data, web interaction logs, text and speech data, face and body video, and electronic health records by leveraging my diverse skills in machine learning, human-computer interaction, natural language processing, speech processing, and computer vision. To develop data-driven interfaces, I also leverage iterative prototyping and qualitative methods like user interviews and thematic analysis, thereby creating interfaces tailored to the needs of the user.

Over the last decade, technology and computing have become an inseparable part of our daily lives. People are constantly interacting with sensor-loaded devices such as smartphones and wearables, and Internet of things (IoT) devices containing cameras and microphones. These devices capture contextual data and digital traces that can provide valuable insights into human activity, such as exercise, sleep, social interactions, financial spending, and content consumption. Mining and analyzing this data can help us better understand behavior and inform data-driven interventions that can drive advancements in fields like health, finance, and collaborative technologies.

While earlier research in this area has been seminal, it has left several open questions, including how to best operationalize sensor features for different application domains and how to overcome the curse of dimensionality when analyzing thousands of features from multiple modalities. My research aims to **address these questions and bridge the gap between technology and human behavior** by developing data-driven solutions to real-world problems.

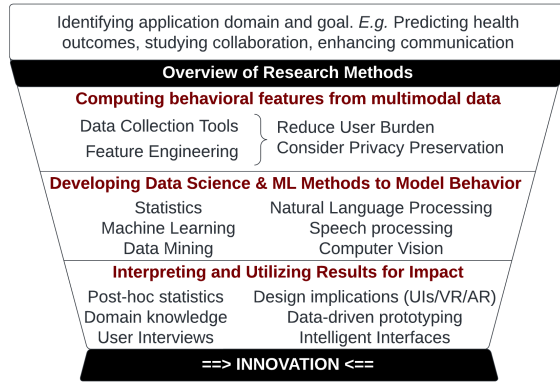


Fig. 1. I develop data science methods and systems to understand and augment human behavior for high-impact application areas.

In my work, I typically address the following **fundamental challenges** (see fig 1): (1) Raw sensor data does not measure behaviors like routines or exercise bouts or internal states like mood or fatigue. New computational techniques are required to operationalize these features and compute them using data from one or more sensors; (2) Scientific research in human behavior has primarily relied on subjective and qualitative methods. New data science and ML methods must be developed to study or model human behavior; (3) Interpreting and utilizing the results from these studies is not trivial. We need to conduct post-hoc statistical analysis, discuss design implications, prototype novel user interfaces, and conduct user and expert interviews.

My work has been published in top-tier computer science venues, including ACM CHI, IMWUT (UbiComp), CSCW, TOCHI, and JMIR, and been recognized through the Snap Research Fellowship from Snap Inc. and the Center for Machine Learning and Health Fellowship from CMU. Through my research, I strive to make meaningful **contributions to the field of data science and human-computer interaction** and positively impact people's lives.

## Modeling and Predicting Health Outcomes Using Multimodal Machine Learning

Depression is a common and serious mental health disorder, affecting 1 in 3 college students in any given year. Its lifetime prevalence is 1 in 2 among people with neurological disorders like Multiple Sclerosis, and 1 in 3 among all adults in the US. Though highly treatable, many people are unable to get treated due to lack of awareness or lack of access to diagnosis and treatment. To tackle this problem, I developed a machine learning approach that uses data from sensors embedded in smartphones and fitness trackers to detect depression. When trained and validated with 138 students, **my algorithm identified college students who experienced depressive symptoms at the end of the semester with an accuracy of 85.7%**, and students whose depressive symptoms worsened over the semester with an accuracy of 85.4%. My analysis also showed that it is possible to **predict the chance of post-semester depression with an acceptable accuracy (>80%), 11-15 weeks before the end of the semester** [7]. I was also able to use my proposed algorithm to detect post-semester loneliness in college students with an accuracy of 80.2%, thus validating its ability to generalize to other outcomes [10].

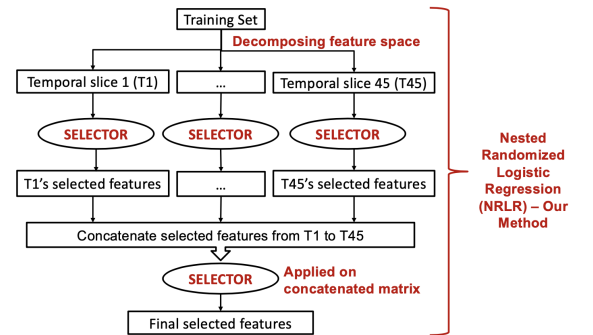


Fig. 2. NRLR – a novel feature extraction approach for multimodal sensor data.

To develop this approach, I extracted features from different time periods during the semester by **operationalizing behavioral phenomenon** such as regularity in movement patterns due to routines, and bouts of activity or inactivity of different lengths due to exercise and sedentary behavior. I also overcame the curse of dimensionality in the feature space by **devising a novel feature extraction approach called Nested Randomized Logistic Regression (NRLR)** that selects the best features by decomposing the feature space and performing feature selection hierarchically (see fig 2). My work has several implications for interventions. For example, detecting post-semester depression allows us to identify students who may have a depressive disorder at the end of the semester, while predicting the chance of post-semester depression weekly, enables us to reach out to students who may be at-risk for depression before the situation worsens.

I successfully adapted the above approach to **predict frequently co-morbid health outcomes for patients with Multiple Sclerosis (pwMS)** during a state-mandated "stay-at-home" period that occurred due to a global pandemic. For this purpose, I collected passive sensor data from smartphones and fitness trackers of 56 multiple sclerosis (MS) patients, and computed behavioral change features that capture the changes that occurred in their behavior due to the "stay-at-home" order. Using these features, my NRLR ML pipeline was able to detect depression with an accuracy of 82.5%, high global MS symptom burden with an accuracy of 90% , severe fatigue with an accuracy of 75.5%, and poor sleep quality with an accuracy of 84%. This work could help clinicians better triage patients with MS and potentially other chronic neurological disorders for interventions and aid patient self-monitoring in their own environment, particularly during extraordinarily stressful circumstances such as pandemics that would cause drastic behavioral changes [5].

### Revealing the Deep Structures of Collaboration Using Multimodal Data Science

Today, distributed teams worldwide rely on computer-mediated communication (CMC) to communicate and collaborate on projects. The quality of communication in these teams can effect their performance, and their wellbeing related outcomes such as satisfaction. Yet, little is known about the behavioral and physiological underpinnings of what makes a team smart and content. To better understand the behavioral and physiological processes underlying collaborative outcomes, I conducted a study in which I **operationalized physiological and behavioral synchrony using Dynamic Time Warping** as an indicator of coordination (see fig 3), and examined its relation with Collective Intelligence (which is a indicator of performance) and group satisfaction. Specifically, we recorded electrodermal activity (EDA), heart rate, facial expressions, and speech of dyads collaborating via video-conferencing on a test that measured Collective Intelligence (CI). The findings showed that synchrony in facial expressions positively correlated with CI, while synchrony in EDA positively correlated with group satisfaction. Further analysis revealed that synchrony in facial expressions is partially responsible for the relationship between social perceptiveness and CI. This creates the potential for technological or behavioral interventions that increase CI by augmenting facial expression synchrony. Examples include, video conferencing systems that make expressions more visible and salient, or training people in effective mimicry [11].

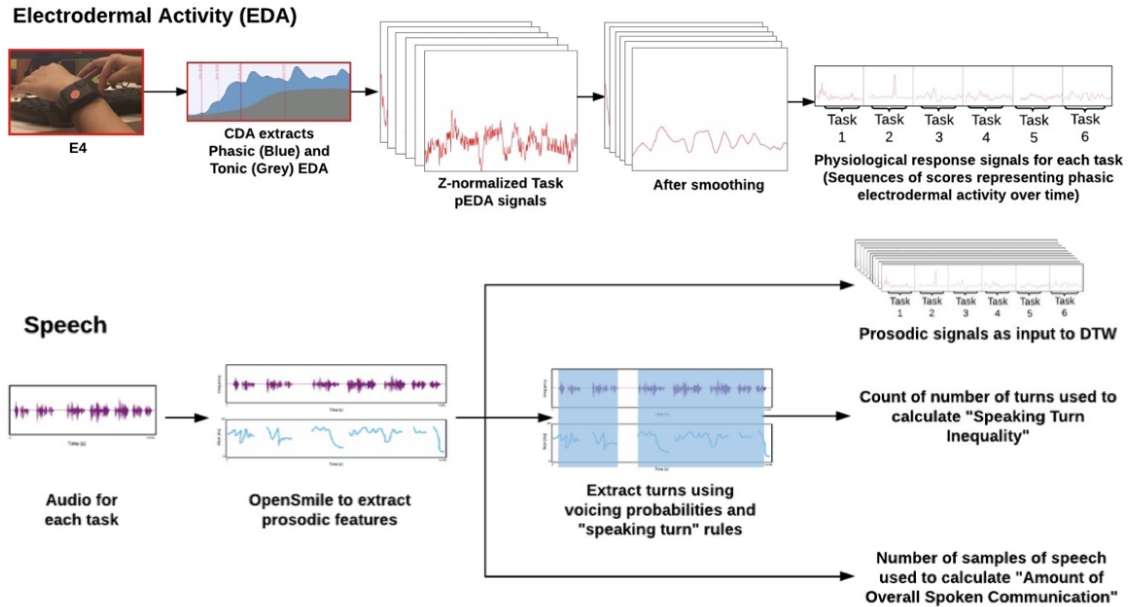


Fig. 3. Computing synchrony electrodermal activity and prosody of speech, and equality in turn-taking in a dyad.

Later, I ran another study that collected the above sensor data from dyads collaborating via audio-only-conferencing instead of video conferencing. In addition to behavioral and physiological synchrony, in this work, I also **computed synchrony in vocal prosody** (e.g. pitch) and **equality in turn-taking** between the dyad (see fig 3). Here, we found that in the absence of facial cues, synchrony in facial expressions does not correlate with CI. Instead, vocal synchrony and

turn-taking positively correlated with CI. Further, when comparing the data from the two studies, we found that in the absence of facial cues, teams have higher vocal synchrony and engage in more equal turn-taking. This dispels the popular myth that video conferencing is superior to audio-only conferencing. Instead, we contend that **video conferencing and audio-only conferencing are both effective** means of collaboration but the physiological and behavioral processes underlying them are different. Hence, this work calls into question the necessity of video support [6].

### Understanding Mental Health and Support Interventions Using Privacy-Preserving NLP

Online mental health interventions are increasingly important in providing access to, and supporting the effectiveness of, mental health treatment. While these technologies are effective, user attrition and early disengagement are key challenges. Evidence suggests that integrating a human supporter into such services mitigates these challenges, however, it remains under-studied how supporter involvement benefits client outcomes, and how to maximize such effects. Hence, during my internship at Microsoft Research, I conducted an analysis of over 200K supporter messages to discover how different support strategies correlate with clinical outcomes. For this purpose, I devised a three fold approach: (i) I **clustered supporters** based on multiple client outcomes in order to identify more and less successful supporters; (ii) I extracted and analyzed linguistic features from supporter messages using a **privacy-preserving feature extraction approach**; (iii) I adapted an **associative rule mining approach** to identify context-specific patterns of support that **reveal best support strategies in very specific multidimensional client contexts**. Our findings indicate that concrete, positive and supportive feedback from supporters that reference social behaviors are strongly associated with better outcomes (see fig 4); and show how their importance varies dependent on different client situations. This work impacted supporter training for human supporters on this platform, and has design implications for personalized supporter interfaces [8, 9].

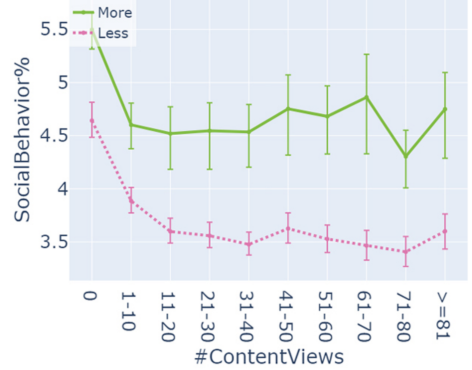


Fig. 4. More successful supporters reference social behaviors more than less successful supporters in all client contexts.

### Designing Novel Visual-First Communication through Word Embeddings and Personalization Rules

Emoji are known to have a strong impact on emotions and positively affect relationships. People in close relationships use emoji in very personalized and secretive ways, thereby forming a shared vocabulary. During my internship at Snap Research, I developed EmoChat, **an emoji-first messaging app** for pairs of users in close relationships that leverages the pair’s closeness and shared experiences to create a **highly personalized "communication zone"** with the goal of enhancing communication. In EmoChat, each message is fully translated to emoji. The recipient always sees the emoji first, but can view the text on demand with a tap (see fig 5) [1, 2, 3].

To achieve the default word to emoji translations, I trained a word2vec model on millions of tweets containing emojis as well as emoji dictionaries containing emoji definitions. These default emoji mappings were then modified by the users’ custom mappings based on their shared vocabulary and relationship.

Through a field study, I investigated the affordances of personalized emoji-first messaging for close relationships. I conducted **user interviews and thematic analysis** of the interview transcripts. I found that emoji-first messaging contributes the following values to text messaging: it adds to humor; it builds secrecy for the exchange of private messages; it promotes creativity, gamification, fun, and play; and it promotes companionship through affection and the sharing of perspectives. To the best of our knowledge, EmoChat is the first ever emoji-first messaging app. This work created a new paradigm for messaging that enables the communication of lightweight symbols before text and has implications for the future ecosystem of communication apps [4].

### Research Agenda

My goal is to build data-driven models and interfaces that can measure, predict, and enhance human behavior and outcomes that impact our lives. I’ve made progress towards this goal during my PhD, but much work remains to be done.

- **Deploying predictive models and studying their acceptability:** Through preliminary interviews with clinical experts and users, I’ve realized that for these models to be widely accepted, explainability and active learning are

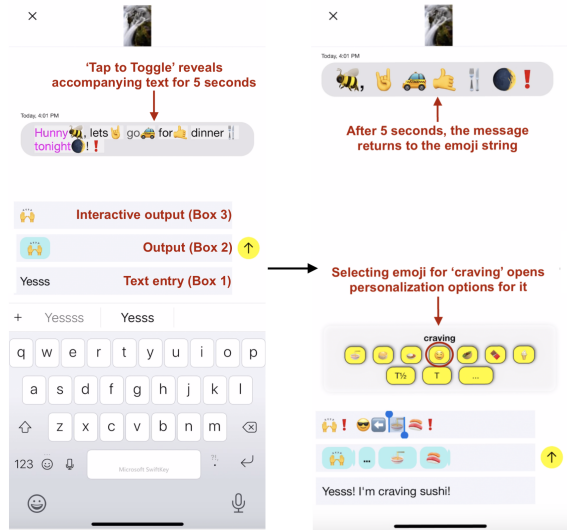


Fig. 5. EmoChat, an “emoji-first” messaging app.

the most important features to have beyond accuracy. That is, users should be able to understand what led to the prediction and the model should learn from the user’s feedback over time. We could, for example, accomplish this by designing interactive visualizations of the feature importances learned by the model. Users could provide feedback by disagreeing with the model’s prediction or by rejecting the model’s inference about a certain feature’s importance. The model can then be retrained using the user’s feedback. While deploying the models, I would also like to conduct user studies and interviews to better understand user preferences for frequency of predictions, level of detail for model explanations, their feelings about false positives, etc.

- **ML-driven personalized interventions and user-driven self-experimentation:** I would like to build on the early work on personalizing interventions using reinforcement learning methods like multi-armed bandits to improve user outcomes in high-impact application areas. However, beyond personalized interventions, I’m very excited about enabling self-experimentation. Self-experimentation is an emerging field that involves giving the users access to tools that allow them to run self-experiments for better decision-making. We could use data-driven insights based on the user’s historical data, and incorporate them in a interactive interface to help them make decisions about what their next actions should be. For example, in the online mental health intervention I worked on, supporters can be shown statistics about how their potential actions can impact the client given their situation – "This client is not engaged. After receiving support messages of over 100 words, 40% of similar clients were able to re-engage. Whereas, after receiving short support messages, only 10% of similar clients were able to re-engage." We can also use a similar approach to help users plan their day. For example, each day the user has a set of actions to choose from such as spend time with family, exercise, and practice mindfulness. If the user has limited time, the UI they’re using to plan their day can help them choose between tasks based on the outcome they care about the most on that day – "So you want to focus on mood today. Your data shows that the days you do mindfulness, your mood improves by 50%. Whereas the days you exercise, your mood improved by 20%. The days you don’t spend time with your family, your mood declines by 50%. What would you like to do next?"
- **Exploring understudied high-impact application areas:** I would like to extend my work to high-impact areas that are largely understudied. For example, I’d like to explore how we can use multimodal behavioral data and data from financial institutions to understand a user’s spending habits and help them make better financial decisions, thereby improving customer experience in finance. We could also create new business opportunities in finance by, for example, enabling "smart loans" that use a user’s behavioral data to determine their risk. We can also use IoT devices and the customers’ personal data enable automatic restocking and dynamic pricing in retail for products. I’d also like to explore if we can use preferences learned from a visually impaired person’s contextual history to better navigate a store by physically directing them to items they may be interested in.

## References

- [1] Bayer, K., **Chikersal, P.**, Nayar, S. K., & Smith, B. A. Client device processing received emoji-first messages. U.S. Patent 11593548, filed Apr 20, 2021, and issued on Feb 28, 2023.
- [2] Bayer, K., **Chikersal, P.**, Nayar, S. K., & Smith, B. A. Personalized emoji dictionary. U.S. Patent 11531406, filed Apr 20, 2021, and issued on Dec 20, 2022.
- [3] Bayer, K., **Chikersal, P.**, Nayar, S. K., & Smith, B. A. Emoji-First Messaging. U.S. Patent 20220337540, filed in 2022. Patent Pending.
- [4] **Chikersal, P.**, Bayer, K., Smith, B. A., & Nayar, S. K. When Emoji Speak Louder than Words: Personalized Emoji-first Messaging for Enhanced Communication between Partners and Close Friends. In Preparation for TBD.
- [5] **Chikersal, P.**, Venkatesh, S., Masown, K., Walker, E., Quraishi, D., Dey, A., Goel, M., & Xia, Z. Predicting Multiple Sclerosis Outcomes during the COVID-19 Stay-at-Home Period: Observational Study Using Passively Sensed Behaviors and Digital Phenotyping. In JMIR mental health (2022).
- [6] Tomprou, M., Kim, Y. J., **Chikersal, P.**, Woolley, A. W., & Dabbish, L. (2021). Speaking out of turn: How video conferencing reduces vocal synchrony and collective intelligence. PLOS ONE.
- [7] **Chikersal, P.**, Doryab, A., Tumminia, M., Villalba, D., Dutcher, J., Liu, X., Cohen, S., Creswell, K., Mankoff, J., Creswell, D., Goel, M., & Dey, A. (2020). Detecting Depression and Predicting its Onset Using Longitudinal Symptoms Captured by Passive Sensing: A Machine Learning Approach With Robust Feature Selection. ACM Transactions on Computer-Human Interaction (TOCHI 2020).
- [8] **Chikersal, P.**, Belgrave, D., Doherty, G., Enrique, A., Palacios, J., Richards, D., & Thieme, A. (2020). Understanding Client Support Strategies to Improve Clinical Outcomes in an Online Mental Health Intervention. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI 2020).
- [9] **Chikersal, P.**, Doherty, G., & Thieme, A. (2020). Towards Using AI to Augment Human Support in Digital Mental Healthcare. In Proceedings of the 2020 CHI Workshop on Technology Ecosystems: Rethinking Resources for Mental Health.
- [10] Doryab, A., Villalba, D. K., **Chikersal, P.**, Dutcher, J. M., Tumminia, M., Liu, X., Cohen, S., Creswell, K., Mankoff, J., Creswell, D., & Dey, A. K. (2019). Identifying Behavioral Phenotypes of Loneliness and Social Isolation with Passive Sensing: A Three-fold Analysis. In Journal of medical Internet research (JMIR 2019).
- [11] **Chikersal, P.**, Tomprou, M., Kim, Y. J., Woolley, A. W., & Dabbish, L. (2017). Deep Structures of Collaboration: Physiological Correlates of Collective Intelligence and Group Satisfaction. In Proceedings of the 20th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW 2017).