Human Routine Behavior Modeling for In-situ Opportunistic Health Interventions using Mobile Robotic Platforms

A dissertation

by

Yunus Terzioğlu

to

the Khoury College of Computer Sciences

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Computer Science

Northeastern University
Boston, Massachusetts

September 2025

Contents

Abstract

Socially assistive robots are envisioned to be deployed in unstructured environments among humans for a variety of applications such as health-related behavior change interventions. A critical property that such robots should have is the ability to adapt to their users' behaviors and needs. To enable practically applicable user adaptation, an efficient approach to learn useful representations of human behaviors and routines is necessary. Building on a case study of opportunistic interventions in a health-related decision-making task, the work presented in this manuscript demonstrates how a mobile home robot can be useful in delivering *in-situ* health interventions and what some of the challenges entailing real-world deployment of such systems are. Following findings and observations from two user studies, this manuscript presents a solution based on human routine behavior modeling that is designed to be generalizable to a variety of different task domains by leveraging spatiotemporal context for opportunistic human-robot engagement. Using a probabilistic model based on Gaussian processes, the proposed work attempts to functionalize this solution in an automated system that actively tracks, models, and predicts routine movement patterns and object interactions of individuals in an unstructured environment with the goal of making it practically possible to make sure that an assistive robot can be in the right place at the right time in a novel user's environment using minimal amount of data.

1 Introduction

Machines are coming into our lives in a variety of forms, including voice agents [60, 58, 83], screen agents [6, 78], chatbots [84, 91], social robots [107, 125], collaborative industrial robots [61], and commercial robots [62, 59], fulfilling the roles of personal assistants, workmates, and even sports coaches [137, 34, 119]. Among possibly the most sought-after variety in this scope of machines are Socially Assistive Robots (SAR, [35]), which are envisioned as interactive companions rather than mere assistants. SARs have the potential to be deployed in dwelling environments to support humans in a variety of tasks, and this thesis investigates facilitation of SARs in health behavior change and habit formation.

Arguably the most prominent utility a SAR offers is the potential to proactively engage with humans at the time and place an intervention opportunity emerges. Facilitated by the robot's mobility and strengthened by proxemics and facial and deictic gestures, engaging with humans in their *social zone* has been shown to be more effective in persuading humans toward behavior change compared to stationary platforms [123, 124]. Such interactions–referred to as *opportunistic interventions* henceforth–are highly valuable as they can assist healthy habit formations, which requires making contextual connections between behaviors and external stimuli [37, 79]. However, cues such as preprogrammed reminders from a mobile phone application do not necessarily facilitate these connections since habitual context is highly dependent on *environmental cues* such as time of day, the physical environment the human is in, and leading and/or preceding activities [144, 113, 3]. In this sense, it can be argued that a proactive SAR with some ability to retain contextual connections between environmental cues and its user's behaviors, can help reform or refine the user's habitual behaviors. Moreover, by highlighting new or better contextual connections in the correct time and place a target behavior is about to be employed by the user, such a robot can help create new habits by serving the role of a mediator between the user and their environment. This requires a SAR to be *aware* of at least a subset of necessary cues within the *context space* which trigger the target behaviors for the user, and able to situate itself within this space for effective reformation of *stimulus control* [99].

One major obstacle in the way of achieving effective opportunistic interventions using SARs, and the initial inspiration and motivation behind the work presented in this manuscript, is these platforms' lack of sufficient agility in navigation and locomotion. As an example, it takes about one minute for a Pepper robot (Figure 1, [107]) to navigate autonomously into the social zone ([45]) of a human three meters away–a task which can be completed in a matter of a couple seconds by a human. This is due to practical issues such as the environment being unstructured and prone to changes, the mapping and localization being imperfect and prone to errors, and other limiting factors due to safety constraints. Such a delay means that by the time the SAR makes it to the human, the *opportunity window* of contextual coherence to deliver an in-place intervention regarding the human's activity is likely missed. Thus, in order for a SAR to be able to locate itself within this window, it either needs to inherit human-

level locomotion agility, which the state-of-the-art social robots do not offer yet; or inherit prediction capabilities so that it can position itself preemptively and wait for this window to form around itself. Such a predictive robot can, for example, predict a human is going to take an action that it wants to intervene in (e.g., *unhealthy snacking*) in a relatively distant future (e.g., *five minutes into starting watching TV*), and preemptively position itself in a contextually coherent position (e.g., *kitchen*) to cue for behavior change (e.g., *recommend apples instead of chips*) when the human eventually enters the opportunity window (e.g., *enter the kitchen to grab snacks five minutes after starting watching TV*). Notice here that the desiderata for such a robot quickly extends beyond the ability to compensate for its inferior agility, and starts entailing the ability to untangle underlying structures of human activity chains effectively.

Beyond the technical shortcomings due to locomotion and navigation, which are likely to become much less limiting in the near future, proactivity–the ability to take actions without receiving explicit commands–itself is an integral component of fluent human-robot interactions (HRI). Human domestic partners create and retain mental models for causal chains governing their partners' actions, habits, and routines. Using these models, humans can make predictions and assumptions regarding current or future actions of their partners and shape their own actions accordingly [130, 70]. This level of understanding significantly reduces the cognitive labor due to repeated communication of otherwise redundant information, and (usually) enables a highly harmonious interaction between partners and is foundational for a fluent relationship. Similar proactive abilities have been shown to be preferred over repeatedly giving explicit commands to robots in human-robot interactions as well [43, 44, 92, 96].

The state-of-the-art and future directions in robotics as well as the knowledge base in social sciences raise several interesting research questions and technical challenges for the future SARs. The work presented in this thesis aims to address some of these questions and challenges by focusing on just-in-time health interventions using SARs: Primarily the problem of predicting human activities with the goal of enabling *opportunistic interventions* through *proactive engagement* and providing *context-sensitive* advice at the time and place a human makes a health-related decision. To functionalize such a proactive robot, a model of causal human routine behaviors is necessary and likely sufficient. Considering that the context space for each human incorporates a high variability due to differences in the environments, behavioral routines, and intervention goals, this modeling problem is effectively a user-adaptation problem, which is yet another open challenge in the literature. Solution attempts to such modeling and adaptation problem leverage deep networks, which are extremely data hungry and unlikely to be useful using a single person's data in a feasible amount of training time. Some other approaches leverage underlying dynamics of the system (such as of a car) in predicting behavioral (driving) trajectories. However, the dynamics of a human is not solely governed by physical constraints but also driven by the *intentional stance* [30], which is non-trivial to formulate. Even if such models were available, state-of-the-art trajectory prediction algorithms typically can only predict up to several seconds, which is insufficient to compensate for the

agility mismatch between a robot and a human in an unstructured environment. Such short prediction horizons are also not sufficient for longitudinal prediction requirements of human-robot co-habitation.

Taking dietary interventions of a hypothetical home robot as an initial case study for deploying SARs in real-life amongst humans with the overarching goal of enabling effective human behavior modeling for a variety of tasks and task domains, the work proposed in this thesis aims to address the following questions:

- **Spatiotemporal Causality:** What kind of a (probabilistic) model can retain sufficient spatiotemporal information and provide continuous-time causal models for fluent human-robot interaction?
- **Proactivity:** How can such a model be useful in mitigating the shortcomings of a SAR in an unstructured environment by acting pre-emptively?
- **Adaptability:** What kind of model can be useful within a practically reasonable amount of time and data requirements for individuals with presumably significant differences between their behavioral routines and environments?
- **Generalizability:** Can such a model be practically and effectively generalized to different task domains?

To address these questions, a modular opportunistic intervention framework and a Gaussian process-based behavior modeling approach are proposed and discussed towards the end of this manuscript. Moreover, demonstrations of the utility of the proposed approach are provided through findings from a human subjects study along with a discussion of possible shortcomings of the approach and their possible solutions.

2 Related work

2.1 Behavior Modeling

Interactive *intelligent* machines are becoming widespread and have been taking on variety of tasks in different forms such as social robots, collaborative robots, virtual wellness coaches, chatbots, voice assistants, robotic vacuum cleaners, targeted advertisement systems, and content recommendation engines. Even though being inarguably distinct in nature, all of these machines are envisioned to be capable of tackling task variations on different levels. These include, but are not limited to, variations in individual user preferences as well as population preferences. Such variations necessitate *adaptive* machines through scalable and generalizable algorithms. In a human-centered computing perspective, this challenge is typically approached by finding meaningful and hopefully reusable representations of human behavior, and tying them together as causal chains of events to create models of human behavior, which can then be utilized to

classify, predict, or simulate human activities [136]. The most common of these approaches make use of parametric and non-parametric models, probabilistic methods, dynamics, machine learning, mathematical models of neurological processes and psychology; and span a wide range of applications such as depression detection from biosignal trajectories and robot learning from demonstrations.

Zhang et al. [144] used computational behavioral models of habit forming and habit strength in a real-world user study focusing on behavior change support systems, investigating dental hygiene as a case study. Using the data collected from sensors attached to participants' toothbrushes, they have shown that theory-based mathematical models of habit formation can explain and predict habit strength and health behavior change better than self-reported measures collected from the participants. In another study, Klein et al. [75] have studied computational models of habitual behavior leveraging neurology literature and proposed a formalism based on Temporal Trace Language to model habitual behavior change mechanisms. They validated their model through simulations generated by their model. In [2], Baker et al. have shown how Bayesian models can be used to reproduce reasoning and prediction patterns from developmental psychology and *intentional stance* ([30]) perspectives. Through a human subjects experiment on a "sprite world" domain, they have shown that such models can successfully model observed human intentional reasoning behavior.

Earlier work by Magnusson [86] introduced a definition of temporal patterns– *T-patterns*–which represent behavioral patterns in time series data, and further, proposed an algorithm for these patterns' detection. Building on this framework, Brdiczka et al. [10] demonstrated how T-patterns can be leveraged to model routine behaviors of office workers. Wang et al. [133] proposed an extension to Topic Models, transforming them from a *bag-of-words* approach to *n-grams* to have contextually coherent clusters of topics in language modeling. Their work inspired further behavior modeling work such as [56]. In this study, Huynh et al. leveraged Topic models to automatically extract activity patterns that model routine behaviors in a generalizable way. Using data from wearable sensors, they modeled daily routine behaviors of elderly in home settings and office workers in their work environment. They have shown that such an approach can be used to represent routines as combinations of activity patterns, and can still be useful without annotated data.

There are several studies in the recent literature that leveraged a framework of Bayesian networks, Inverse Reinforcement Learning (IRL, [102]), and Maximum Causal Entropy algorithm ([147, 145]) to learn and model human behavior. Banovic et al. [3] used this framework to capture causal relationships that define humans' routine behavior. They used existing activity prediction algorithms and proposed an extension to these that offer a more generalizable framework for causal routine prediction. They tested their approach on modeling the differences and predicting the behavior of aggressive and non-aggressive drivers. Ziebart et al. [146] used IRL to model driver behavior on public roads, and have shown how such a model can be used to predict human behavior in a causal manner. In [104], Reddy et al. have shown how this framework can be used to learn and refine beliefs of users' understanding of their environment.

They used a navigation task in a grid world domain and the Lunar Lander game ([13]) to show that their approach can successfully model user beliefs, and provide user-adapted corrective signals to improve their performance in the game significantly. In a recent study, Hossain et al. [54] used a similar IRL framework to infer if a certain dataset is sufficiently large to model certain human behavior. They have demonstrated the functionality of their approach in a case study using data from people with Multiple Sclerosis.

Parametric models, i.e., models that attempt to represent and generalize observed data with a (small) set of parameters, have been in the literature for a long while. In their work [51], Helbing and Molnár have shown how pedestrian trajectories can be modeled and predicted with great accuracy using their Social Force Model. Further, building on the seminal work of Kalman and Bucy [68] to improve the capabilities of Hidden Markov Models (HMMs) in characterizing human behavior, Pentland and Liu [97] proposed a method to combine a (parametric) dynamics model with a (stochastic) Markov Chain to model and predict human behavior. They tested their approach in predicting driver trajectories a few seconds into the future from their preparatory actions. In a more recent work, Wiest et al. [134] leveraged Chebyshev polynomials to reduce state-space complexity and predict driver trajectories through probabilistic inference using Gaussian Mixture Models (GMMs). They used past trajectory snippets to predict the driver trajectory a few seconds into the future, and discussed how such models can be useful for improving driver assistance systems to improve traffic safety. More recently, Deng and Söffker [29] used layered HMMs to model and predict driver lane-change behavior, and used a genetic sorting algorithm to find optimal model parameters. They have shown that this approach outperformed other baseline machine learning methods such as artificial neural networks and support vector machines. Drawing on the observation that most context-sensitive information to model human behavior is in the form of time series data (e.g., trajectories, biosignals) and most probabilistic models work on discrete data, Bobu et al. [7] revisited discrete probabilistic models to make them better applicable to continuous domains. Building on the literature in Boltzmann noisily-rational decision models and Luce's choice axiom ([85]), they proposed a method to cluster time series trajectories and formulations to have these clusters inform probabilistic models. They show that their representation outperforms conventional models in modeling observed human behavior in two human subjects studies.

Emphasizing the shortcomings of parametric modeling such as their inflexibility to data variations and laborious parameter estimation, Sun et al. [118] proposed a non-parametric Bayesian method to discover routine human behaviors which is based on Dirichlet Process Gaussian Mixture Models and Hierarchical Dirichlet Processes. As case studies to show that their approach outperforms a more conventional parametric method, they made use of a daily human activity dataset that incorporates atomic and composite actions as well as wearable sensor readings (e.g., accelerometer), and a transportation mode dataset which incorporates longitudinal position trajectories which are annotated with the associated transportation methods the participants used. Ellis et al. [32]

proposed a non-parametric long-term trajectory prediction method based on Gaussian process regression. Using their clustering algorithm on surveillance camera footage, they successfully created models of pedestrian trajectories in public settings, and highlighted the scalability issues non-parametric methods are likely to suffer. There is a large body of literature on human trajectory prediction some of which is summarized by Rudenko et al. in their survey paper [109] covering both parametric and non-parametric approaches.

In their recent work, Xu et al. [142] drew attention to generalization difficulties in human behavior modeling research which stem from variations in data both within- and across-population. They propose a clustering and a sampling algorithm in an attempt to tackle the challenge of generalization across datasets. They verify their approach in comparison with several other methods from prior work using a 4-year-worth longitudinal passive dataset they collected using mobile devices and wearables to detect depression. They publish the algorithms they developed as well as the baseline algorithms on a benchmarking platform for other researchers to make use of. In their earlier work [141], Xu et al. have shown how human routine behaviors could be extracted from passive multi-modal sensor data such as nearby Bluetooth devices, location, charging state etc. and this data could be used for depression detection using Association Rule Mining method. In their work [98], Pierson et al. also made use of a clustering-based method and proposed Cyclic HMMs to featurize cycles, recognize feature variability in cycles, and cluster similar groups with the goal of tracking menstrual cycles. They have shown that their method is robust against missing data and their model can extract new (health-related) information, even when the specific information was not explicitly collected from the donors.

Other example use cases from human behavior modeling can be found in commerce. In their work, Yanchenko et al. [143] demonstrated how targeted discounts and modeling customer behavior can be leveraged to optimize profit in a scalable way. Using Bayesian Dynamic Mixture Models, they model the purchasing trends based on the spending characteristics of different household groups at the individual buyer and also individual item level to do forecasting. In another work to improve service quality, Kanda et al. [69] used a clustering algorithm to classify individuals around a shop, and predict their motions a few seconds into the future to decide whether they are potential customers or passerby. Deploying their algorithm on a mobile humanoid robot, they approached and invited the potential customers into their shop, and showed that with this anticipatory behavior, they were able to increase the services they have provided with the robot. Focusing on building automation systems in residences, Bruckner and Velik [15] used HMMs and data collected from proximity sensors placed in an office environment to model routine activities within the office. They have shown how unannotated data can still bear useful representations of routine data in such environments. In their work [63], Jahanmahin et al. surveyed the literature on human behavior modeling for human-robot interactive manufacturing.

Observing the difficulty of collecting real-life behavioral data, especially from dwelling environments, some studies in the literature focused on synthesizing be-

havioral data by using *generative* behavioral models. Puig et al. [101], collected a dataset of household activities from participants through Amazon Mechanical Turk. The participants composed the prompted activities using a user-friendly drag-and-drop programming interface with atomic actions. Using these atomic action sequences, video sequences, and language descriptions with a variety of machine learning techniques, they generate a "program" that encapsulates the relations between these modalities and final activities. They also developed a home simulation environment, animated atomic behaviors, and deployed this program on this simulation environment to successfully synthesize household behaviors from prompts. Idrees et al. [57] proposed a framework for daily human activity simulation to help generate data for research and development of product testing for home robots (e.g., Amazon Astro). They used behavioral data from three publicly available household activity datasets, and applied their data augmentation algorithm, which is constraint- and heuristic-based, to add variations and noise to the readily available activity schedules. They show that the generated schedules are similar to the schedules in the training set using a distance metric. Elbayoudi et al. [31] focused similarly on household data scarcity, and proposed an algorithm to synthesize daily routine behaviors of older adults. Their training data was composed of timestamps, occupancy, and movement activities within the house. Using HMMs and Direct Simulation Monte Carlo methods, they have shown that their simulator can generate behavior similar to the data in the training set.

This body of literature draws attention to the significance of human behavior modeling in a vast pool of application areas such as behavior change, elderly care, driver assistance systems, and commercial market. The studies show probabilistic models, machine learning, system dynamics, and formalisms on routine and/or habitual behavior can be effective tools for approaching this modeling task. The primary open challenges in behavior modeling literature entail the difficulties around generalization of the proposed models and personalization of the modeled behavior due to data scarcity. Majority of the studies in the field rely on domain-specific information, datasets, and heuristics, which restricts the applicability of the approaches and, thus, their generalizability. The trajectory prediction methods typically can only predict up to a few seconds into the future, or they rely on clustering tools which require further domain-specific information. The work proposed in this manuscript primarily aims to address these shortcomings of the models in the literature regarding generalization and personalization, and offer a method that does not rely on domain-specific heuristics to achieve long-horizon behavior prediction.

2.2 Persuasive Robotics

Persuasion, in relation to persuasive computing, is defined by Fogg [36] as "an attempt to shape, reinforce, or change behaviors, feelings, or thoughts about an issue, object, or action." Persuasiveness has been shown to be an important social process that facilitates social influence, cooperation, and attitude changes [103, 148]. Persuasive robotics requires an understanding of how persuasion in

human-human interactions translates to human-robot interactions [11]. There is now a large body of literature on persuasive computing [36] in general and persuasive robotics [115] in particular. Several studies in the literature have investigated verbal and nonverbal strategies that robots can use to persuade users.

Leveraging the Elaboration Likelihood Model (ELM) [19, 94], an established model of persuasion in human-human interaction, Winkle et al. [138] evaluated the strategies of goodwill, similarity, and credibility on health behavior compliance, in having participants perform a simple wrist rotation exercise. This study also explored the use of the expert strategy from the credibility peripheral cue of the ELM, which establishes the persuader (in this case the robot) as an authority on the subject matter for which persuasion is required.

Ham et al. [47, 46] studied robot persuasion to improve attitudes toward energy conservation. They found that social feedback was better than factual feedback and negative feedback was better than positive feedback to elicit behavior change. By using negative feedback and robot proxemics, they were able to create a sense of urgency and persuade users to take remedial action.

Rincon et al. [106] designed a social robot, *EmIR*, for assisting older adults in their daily activities. The authors used three argumentation strategies, namely analogy, popular practice, and expert opinion, to persuade users to do activities using a framework proposed by Costa et al. [27]. This framework leveraged users' interests and preferences to generate tailored persuasive strategies, which enabled the EmIR robot to use proxemics to persuade older adults to engage in activities that they might not have considered otherwise.

Ghazali et al. [39] demonstrated that apparent gender congruence between user and robot may lead to persuasion, but not necessarily improve the trust between them. On the other hand, Siegel and Breazeal [115] showed that male users were more likely to donate money to a female robot compared to a male robot, while female users demonstrated no preference. They have also shown that participants found robots of the opposite sex more credible, trustworthy, and engaging.

Hashemian et al. [49] explored social power and persuasion in HRI. Two humanoid robots each of whom employed a different social power strategy to encourage participants to pick one of three different coffee packs. One robot used a reward strategy and the other established itself as a coffee expert to do so. Both approaches were found to be equally persuasive. They also demonstrated that the relationship between social power and persuasion is not linear and that repeated persuasion attempts do not decay the perceived value of the rewards when rewards are used as the social power strategy [48].

Herse et al. [53] explored how different embodiment types influence persuasion in recommendation systems. Through a vignette study, they compared persuasion across two recommendation statements: one related to the atmosphere and one about the staff; and three hypothetical embodiment types: human, robot, and information kiosk, which aimed to persuade the participants to choose one of two restaurant options. They found that using a human embodiment for recommendation is better than using a robot or information kiosk.

However, this finding was only significant in the case of one statement which spoke about the good atmosphere at the restaurant. They also used cartoon icons to depict each embodiment which may have affected the participants' perception of the recommender.

One study by Ju and Sirkin [67] explored the use of a waving robotic hand as a mechanism for attracting the attention of people passing by a public kiosk. The study found that the physicality of the robot was a powerful tool for persuasion, as the waving hand was more effective in engaging participants with the kiosk than an animated hand on a display. This demonstrates the potential of proxemics as a powerful mode of persuasion in public spaces, where robots can use physical gestures to draw people's attention and engage them in interaction.

Another study by Chidambaram et al. [23] explored how manipulations in a robot's nonverbal cues affected participants' perceptions of its persuasiveness and compliance with suggestions. The study found that nonverbal cues, including deictic and metaphorical gestures, were effective at persuading participants when combined with speech cues.

These studies demonstrate that behavioral theories on human-human persuasion can be transferred to human-robot interactions. Moreover, they highlight the importance of robot embodiment, form factor, and role-setting on the effectiveness of persuasion. Additionally, they show that proxemics can be a powerful mode of persuasion, and by combining proxemics with physical gestures, spatial cues, and other nonverbal cues, robots can create persuasive messages that engage users and encourage behavior change.

2.3 Proxemics and Nonverbal Behavior

There are several kinds of communication behavior categorized under nonverbal behavior that are studied extensively in social human-robot interaction [110]. Proxemics is one of the branches of nonverbal communication that examines the use and interpretation of physical space in interpersonal interactions. Research on proxemics has extended to human-robot interactions, investigating how proxemic behavior affects humans' perceptions of and experiences with robots. Kim and Mutlu [74] explored the effects of social distance on user experience by manipulating the physical proximity, organizational status, and task structure of a humanoid robot interacting with participants. They found that participants had a better experience performing a task when the robot was established as a supervisor and it was physically closer to the participant. According to further studies by Walters et al. [132], Koay et al. [76], and Mumm and Mutlu [90], factors such as distance and likability can affect people's attitude and behavior toward robots.

Other studies investigated how user gender and cultural background, and robot appearance and voice can affect proxemic behavior. Obaid et al. [93] did not find a significant effect of gender on human-robot proxemics (HRP), and Eresha et al. [33] found a significant effect of culture on HRP. In terms of robot appearance, the study by Samarakoon et al. [111] found that participants' preferred interaction distance with a manipulator robot was significantly less than

it was with other robots. Terzioğlu et al. [122] have shown that gaze cues can be leveraged to increase collaborator rapport with an industrial manipulator. The degree of anthropomorphic attribution is linked to higher expectations from the robot in terms of proxemics norms, according to Syrdal et al. [121]. The study by Walters et al. [131] suggests that a synthesized voice leads participants to stand significantly closer to the robot compared to a robot with a natural voice or no voice.

Mead and Mataric [87, 88] investigated how robot task performance affects humans' perception of robots' competence, anthropomorphism, engagement, and likability. The study found that minimum and average robot performance levels were correlated with how humans perceived the robot's competence, anthropomorphism, and engagement, while the maximum performance level was not. However, likability was significantly correlated with all three levels of performance. Papadopoulos et al. [95] explored the effects of proxemic positioning on engagement and collaboration in a memory task. They found that participants were more engaged with and had higher preferences for a NAO robot in a frontal position compared to a lateral position. Additionally, a robot with helpful speech behavior was found more engaging and preferable compared to a robot with neutral behavior.

Siegel [114] examined the effects of social distance on donation behavior towards a Mobile Dexterous Social humanoid-torso robot. The study found that male participants donated more in a 'close' condition (75 cm distance) than in a 'normal' condition (1.5 m distance), while female participants donated more in the 'normal' condition compared to the 'close' condition. Syrdal [120] explored the role of spatial behaviors in building human-robot relationships. Overall, the findings suggest that physical interactions play an important role in building a relationship between a domestic robot and a user, even when comparing robots with similar interaction capabilities. Participants reported feeling closer to the mobile robot and rated it as more likable than the stationary robot, suggesting the importance of negotiating shared physical space in real-time.

These studies indicate how robot proxemics and nonverbal behavior in combination with robot features such as form factor and voice can be manipulated to positively affect the attitude towards and also help shape expectations from the robot.

2.4 Robot Persuasion through Nonverbal Behavior

Persuasion has also been shown to be affected by an individual's use of nonverbal cues such as gaze and proximity [105]. Nonverbal behaviors such as gaze, proximity, gestures, posture, and facial expressions; and verbal behaviors such as vocal tone, and expressions can also shape nonverbal immediacy, the degree to which someone feels connected to another, which is closely related to persuasive ability [89]. Joint activity, for example, which requires coordinated efforts from two or more partners, shows that nonverbal signaling is especially important as means of directing attention to particular objects or regions. Gaze has also been shown to have some influence on joint activity tasks such as deciding

between two choices [20]. Ju and Sirkin [67] explored the use of different mechanisms to attract the attention of people passing by a public kiosk. They found that the use of a waving robotic hand was better at persuading participants to engage with the kiosk compared to an animated hand on a display, demonstrating the enhanced persuasive ability of physicality in robots. Ham et al. [47] demonstrated that persuasion was increased when a robot used a combination of both gazing and gesturing in a storytelling task. Nonverbal behaviors have also been shown to be effective in improving the retention rates of participants in a storytelling-based task where robots used different kinds of iconic gestures (or not) [129].

Ghazali et al. [41] investigated the use of social cues as persuasion strategies. In a series of studies, they looked at how cues such as mimicry, praise, as well as emotional intonation, head movements, and facial expressions, can be used to persuade people to change their choices in trivial tasks. However, these studies relied on robot speech as the primary modality for persuasion, and did not investigate the persuasive effect of nonverbal behavior in isolation [41, 40].

Chidambaram et al. [23] explored how manipulations in the use of a robot's nonverbal cues (with and without the addition of speech cues) affected the participants' perceptions of the robot's persuasiveness and compliance with the robot's suggestions. Participants performed the *Desert Survival Task* [80] on a computer with a robot providing suggestions using a combination of verbal and nonverbal cues. The gestures used by the robot included pointing (deictic) gestures to reference itself and the participant, metaphorical gestures, to show a space containing an idea, and other kinds of gestures where appropriate. The study demonstrated that nonverbal signaling outperformed having no signals; and that nonverbal cues were effective only when combined with speech cues.

These studies suggest that nonverbal behavior can be an effective tool for persuasion in human-robot interactions and encouraging behavior change. Overall, the body of research on proxemics, nonverbal behavior, and robot persuasion supports our claims and findings (Section 3) regarding the use of robot mobility as a facilitator of proxemics and robotic gestures to reinforce verbal cues during human-robot interactions.

2.5 Further Topics

There are several other fields of research that are relevant to the work proposed in this manuscript regarding behavioral medicine [4, 16]; social influence and persuasion [24, 117]; robotic persuasion [12, 139] and coaching [65, 8, 137, 119, 34]; human intent recognition [55, 64] and activity recognition [66, 5, 22].

3 Foundational Studies

There is a large body of HRI literature investigating robot capabilities provided by their physical form factor in a wide range of application domains. However, it is still unclear how a physical mobile agent stands out compared

to other technologies–such as voice assistants–in health-related decision-making interventions. Understanding such difference, if any, and how each competency provided by the form factor contributes to this difference is crucial before making any comprehensive attempt to utilize a mobile robot in a health intervention application. The first reason why it is crucial to develop this understanding is the *cost* of working with such robot platforms, which comes in two forms: The monetary expense of purchasing and maintaining such platforms; and the technical labor required to functionalize a robot to move around unstructured environments and interact with humans. The second reason is the many confounding factors that influence the interaction–especially when working with humanoid robots–and the difficulty of isolating contributions from individual competences (e.g., speech, movement) and design elements (e.g., face, human form) on the investigated outcomes.

The discussion of such confound effects might have indirect roots in theories such as Attribution Theory [50] or more direct roots related to human physiology, such as our nervous system being *hard-wired* to perceive faces [1]. Although an investigation of such effects individually overreaches the scope of this work, a series of controlled trials were conducted to validate the effects of factors with the most relevance–such as mobility, gestures, and role-setting–which distinguish such a physical robot platform from other possible technologies. For this purpose, a testbed comprised of a fake food pantry was created (Sect. 3.1). As the intervention platform, a semi-humanoid robot (Sect. 3.2) was used in a Wizard-of-Oz [28] human subjects experiments (Sect. 3.3) [123, 124]. Using this testbed, we were able to demonstrate significant effects of robot verbal and nonverbal communication and proxemics–which were primarily enabled by robot mobility–on persuading participants to make healthier decisions.

3.1 Nutritional Testbed and Fake Food Pantry

An experimental testbed to facilitate the evaluation of persuasive strategies that a mobile social robot could use to influence food choice was developed. This testbed was inspired by prior work by Bucher et al. [17, 18] and Ung et al. [126] on the use of a *fake food buffet* and a *virtual cafeteria*, which are simulations of a cafeteria, stocked with fake food of varying nutritional quality, in which study participants can make food selections in order to evaluate the impact of different interventions on food choice. A repeated-trials experiment design was used to test the persuasive interventions several times for each participant. In these experiments, the participants were asked to assemble 20 meal boxes during each study session.

Our food pantry (Figure 1) contained 45 food items that could be included in a lunch. These items were intentionally selected to be the kinds that did not require cooking (e.g., salad, sandwich) or that could be consumed by simply heating up or adding hot water (e.g., soup, ramen). We used real food for long-shelf-life items and fake food for perishables. The food items on the shelves had highly varied healthiness levels. Building primarily on the nutritional indices provided by the Food and Nutrient Database for Dietary Studies (FNDDS) [127]

and Food Patterns Equivalents Database (FPED) [128], and the food quality assessment scheme provided by the Healthy Eating Index (HEI) [77], a meal quality evaluation tool was implemented to score individual food items in our pantry as well as combinations of meals that can be assembled using these items. This tool computed a score between 0 (least healthy) and 100 (most healthy) using the FNDDS and FPED nutritional indices for our inventory of food items and the scoring formulae provided by HEI.

Figure 1: Our fake food pantry which is used to evaluate how robot proxemics, through mobility and multimodal cuing, can influence health-related decision-making in a collaborative meal assembly task.

3.2 Intervention Platform

As our intervention platform, a 120-cm-tall semi-humanoid robot was used (Pepper [107]; Figure 1). This 19-degree-of-freedom (DoF) robot can execute sophisticated gestures using its head, arms, and body; can move around using its wheels; and communicate verbally. The robot was controlled with a handheld game controller in a Wizard-of-Oz configuration [28] by the experimenter who was hidden in the adjacent room during the experiment sessions.

During our experiments, participants engaged in a counseling session led by the robot at the beginning of their sessions. The counseling conversation with the robot was intended to 1) build rapport, trust, and therapeutic alliance with the robot, 2) establish the nutrition expertise of the robot, and 3) establish the role of the robot during meal assembly task. The conversation consisted of a brief greeting and a few turns of social chat, a review of the current US guidelines for nutrition, a brief counseling dialogue in which the robot assessed

the participant's own dietary behavior and motivation for change, a discussion of what the robot will be doing during the meal assembly ("I will help you by occasionally making healthy suggestions."), and a statement of collaboration ("I look forward to working with you. Let's get to work!"). During the dialogue, the robot used appropriate conversational nonverbal behavior, including eye gaze and hand gestures. This dialogue was pre-scripted and was also controlled by the experimenter who observed the interaction from a different room.

3.3 User Studies

Two user studies on the effects of health interventions using the mobile robot in our fake food pantry were conducted.

The first study [124] explored the use of mobility and deictic gestures (without speech) by a humanoid robot to indicate healthy food choices at the time and place a participant is making a decision. Providing such in-context reminders–referred to as 'stimulus control' in the transtheoretical model of health behavior change [100]–is among the most powerful techniques for changing behavior. In order to motivate the robot's behavior and establish trust and therapeutic alliance with individuals prior to health decision events, the use of trust-building and health counseling dialogue by a robot was also explored in conjunction with cuing behavior. In this study, the robot's role was manipulated between participant groups: A randomly assigned subgroup of participants went through the counseling dialogue with the robot before beginning the experimental task. This dialogue was intended to set the robot's role as a counselor under `SOCIAL` condition. The remainder of the participants directly started the meal assembly task without the initial dialogue with the robot, setting the robot's role as a `BYSTANDER`. During the experiment sessions in both `SOCIAL` and `BYSTANDER` conditions, for a randomly selected half of the lunch boxes that the participants assembled, the robot approached the pantry and the participant who was assembling a lunch box and *cued* a randomly selected healthy food option from the pantry shelves. This cuing behavior consisted solely of physical gestures such as pointing and gazing at the food item as well as gazing at the participant's face without any further verbal interaction.

A second study tested two factors at each decision point: The effects of multimodal advice by the robot compared to no advice, and the effects of robot mobility and proxemics on food choice, comparing whether the robot moved into the participant's workspace to deliver advice (`MOBILE`) or remained stationary, giving advice at a distance (`STATIONARY`) [123]. During the intervention trials under `MOBILE` condition, the robot approached the participants to deliver advice, using the same nonverbal behavior as in the first study, but accompanied these gestures with verbal cues. These verbal cues were composed of a propositional statement selected randomly from a pool of templates (e.g., "*How about <SOME BROCCOLI>?*") followed by a supporting argument which was selected randomly from a pool of statements for each healthy food item (e.g., "*Broccoli is full of vitamins, minerals, and antioxidants.*"). During the intervention trials under the `STATIONARY` condition, the robot stayed stationary in

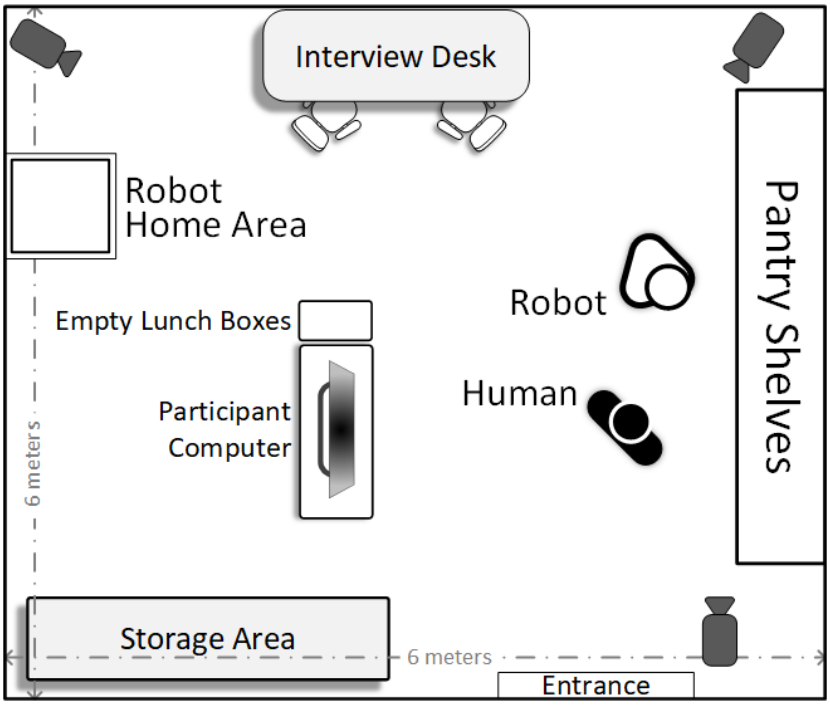

Figure 2: The lab layout used in the food pantry experiments.

its *Home* Position (Figure 2) and delivered the same multimodal advice from a distance. During the control trials, the robot stayed in its home position running idle animations and did not interact with the participants.

3.4 Effects of Role-setting

The findings from the first study revealed that the `SOCIAL` robot was rated significantly higher than the `BYSTANDER` robot in satisfaction ($U = 94.0, p = .025, \eta^2 = 0.228$) and perceived knowledgeability ($U = 104.5, p = .002, \eta^2 = 0.383$). Moreover, the participants felt significantly safer around the `SOCIAL` ($t(19) = 2.045, p = .028, d = 0.876$) robot and trusted it more than the `BYSTANDER` robot ($t(20) = 2.334, p = .015, d = 0.999$). Additionally, there were trending positive main effects on ease of interaction ($U = 89.0, p = .051, \eta^2 = 0.166$) and the naturalness of interaction ($U = 88.5, p = .059, \eta^2 = 0.162$) in favor of the `SOCIAL` robot. We found a trending positive main effect of the `SOCIAL` condition on HEI scores ($F(1, 20) = 3.019, p = .098, \eta^2 = 0.131$, Figure 3) compared to the `BYSTANDER` condition. Also, the participants who interacted with the `SOCIAL` robot scored significantly higher in the post-experiment meal quality assessment test ($t(20) = 2.844, p = .010, d = 1.218$).

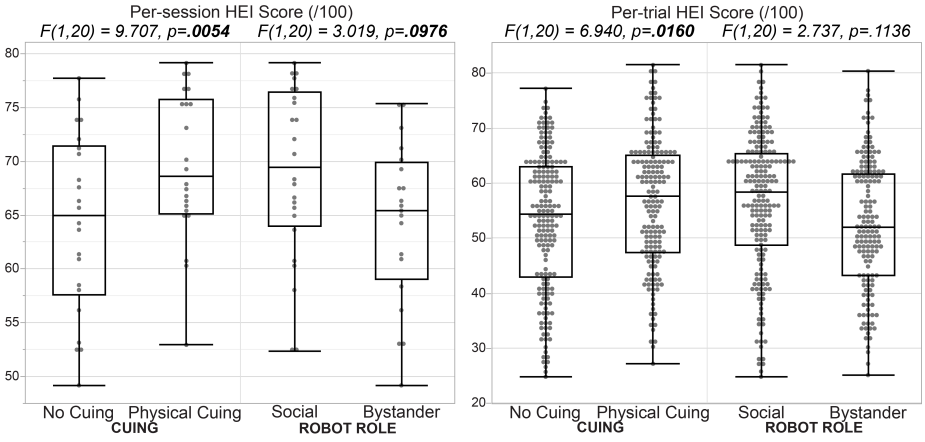

Figure 3: ANOVA results for the primary meal quality outcomes along with the data distributions from the first study.

3.5 Effects of Non-verbal Cues

The analysis on the first study's primary outcomes revealed significant positive main effects of physical cuing on the assembled meal quality. Per-trial scores–the meal scores of each basket the participants assembled–were significantly higher in the physical cuing trials compared to the no cuing trials ($F(1, 20) = 6.940, p = .016$, Figure 3). Additionally, the aggregated per-session scores–the total meal score of all baskets for each participant under both within-participant conditions–were also significantly higher in the physical cuing trials compared to the no cuing trials ($F(1, 20) = 9.707, p = .005, \eta^2 = 0.327$).

3.6 Mobility, Proxemics, and Multimodal Cues in Compliance, Engagement, and Alliance

The findings from the second study showed a significant positive correlation between `MOBILE` condition and participants' compliance with the robot's recommendations ($\chi^2(1, 218) = 12.376, p = .0004$). Moreover, the participants under the `MOBILE` condition assembled significantly healthier meals compared to the participants in the `STATIONARY` condition ($F(1, 20) = 4.556, p = .045$). Similar to the findings from the first study, the multimodal cues used in the second study–speech as well as deictic and gaze cues–had a significant positive main effect on the assembled meal quality compared to the trials under no cuing condition ($F(2, 40) = 7.833, p = .001$).

The participants' vocal responses–verbal or non-verbal (*interjections*)–to robot recommendations were coded as a metric of engagement for the second study. An analysis of covariance revealed a significant positive main effect of robot mobility on participant engagement ($F(1, 19.99) = 26.314, p < .0001$)

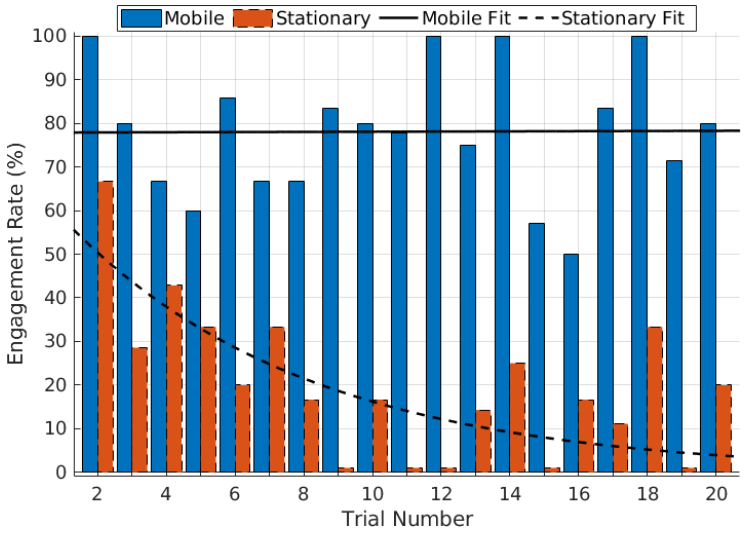

Figure 4: Aggregated engagement trends of the participants throughout 20 experimental trials for `STATIONARY` and `MOBILE` robot conditions along with the associated exponential curve fits. The decay rate for the `STATIONARY` condition was found as 0.15 ($\chi^2 = 10.86, p = .001$) while there was no significant decay observed in the engagement rates for the `MOBILE` condition.

as well as a significant effect of trial number on engagement ($F(1, 16.47) = 8.267, p = .011$). Further logistic regression analyses for both *STATIONARY* and *MOBILE* experiment conditions revealed that the trial number was a significant predictor of participant engagement for the *STATIONARY* condition ($\chi^2(1) = 6.462, p = .011$), whereas there was no such significance for the *MOBILE* condition ($\chi^2(1) = 0.008, p = .926$). As demonstrated in Figure 4, the participants' engagement rate was relatively consistent throughout the trials for the *MOBILE* condition, whereas the engagement decayed exponentially through the trials under *STATIONARY* condition. Additionally, a significant correlation between participant engagement and compliance ($\chi^2(1, 217) = 15.963, p < .0001$) was also found, indicating the significance of engagement on collaborative alliance.

To analyze the effect of robot mobility on therapeutic alliance between the participants and the robot, we ran independent samples t-test on Working Alliance Inventory-Bond scores. Finally, the participants under `MOBILE` condition rated a significantly higher working alliance with the robot ($t(18.34) = 1.920, p = .035$) as well as a significantly higher intention to continue working with the robot ($U = 31.0, n_1 = n_2 = 11, p = .029$) compared to the `STATIONARY` condition.

3.7 Conclusions

The findings from these foundational human subjects experiments provided sufficient validation and justification regarding the utility of a humanoid robot form factor in opportunistic health interventions. Based on these findings, *nudging* humans at the time and place a health-related decision is being made was found to be highly valuable for improving the outcome of the decision, even if the nudge consisted solely of physical gestures.

The studies have shown that the robot can be framed as an expert on health topics simply through social dialogue. This role-setting dialogue has proven valuable for increasing trust in the robot, potentially leading to higher compliance with the robot's recommendations.

The gestures alone were shown to be legible within this context to significantly improve the participants' decision-making, even when the majority of the participants complained that it was difficult to understand what the robot was trying to communicate with its gestures. Addressing this shortcoming and using multimodal cues in the second study, which combined the gestures with speech, had even stronger effects on improving decision-making and resulted in significantly higher quality meals.

Probably the most striking findings from these studies entailed the reflections of proxemics theories in health-related robotic interventions. The ability to move into a human's social zone ([45]) while giving advice was found to be highly valuable in achieving a high rate of compliance as well as retaining engagement. This proxemics behavior was found to be the most significant factor in increasing the decision-making quality throughout our experiments; and its absence was found to be the primary factor behind losing engagement and, consequently, compliance over time. This ability, which is potentially the *costliest* component of such platforms, is also potentially the most valuable in making them stand out compared to other stationary platforms such as voice assistants.

In addition to providing sufficient fundamental understanding of how physical form factor can contribute to improving health-related decision-making using opportunistic cuing, these studies also provided two other important lessons: (1) Humans are fast! (2) Decision-making is perceivably momentary. The initial intention prior to starting these experiments was to have the robot drive autonomously within the lab. However, during the pilot studies, it became clear how naive this ambition was after observing how comically slow the Pepper robot was compared to a human. Even if the robot had access to the full map of the laboratory and was pre-loaded with the free configuration space, it still took 20-30 seconds for it to pace the 5 meters between its home position and the pantry shelves. Note that this duration was measured under the "best" conditions where the robot only drove over a completely straight line and there was no moving entities around it in the room to trigger collision avoidance behavior. Consequently, by the time the robot made it into the shelf after detecting the participant was picking up food items, the participant had already completed their lunch box, made it back to the logging computer, and was about to finish logging their current meal and starting the next trial. Even when the

autonomous navigation was replaced with a teleoperation system, where the experimenter drove the robot using a remote controller from the next room, it was still not practically possible to keep up with the participants. We were able to catch up with them by the pantry, more often than not, but the robot's recommendations were then usually followed by responses such as "Oh. I've already picked up <food item>. I'll get <the recommended item> later," indicating we have missed the decision-making point and it was *already* too late. Following these test runs, an intentionally bad meal logging user interface was designed and implemented for the participants. This bought us some crucial seconds as we preemptively drove the robot to the pantry to catch the participant at the next decision-making moment and, hopefully, deliver a successful intervention. This struggle was another reminder why proactivity is important and was the initial motivation behind the model presented in the following chapters of this thesis: A model that can anticipate user behaviors in unstructured environments so that we can use robots to provide just-in-time advice, without being comically late, and achieve opportunistic cuing for health behavior change.

4 Approach

The approach taken to tackle the problem of modeling and predicting the human behavior in this thesis is rooted in the assumption that the majority of human activity is repetitive, as defined by "routine behavior" [3, 108]. This assumption is further supported by earlier studies which have shown the human mobility to be 93% predictable [116] and follow consistent spatial and temporal patterns [42]. Furthermore, the significance of routine behavior becomes especially relevant and prevalent within the context of health behavior considering that the majority of such behavior is in fact habitual [14, 37, 81], meaning that they are repeated on a regular basis triggered by environmental cues [37, 79], which form the *context space*. This connection between health behavior and routines makes the problem of modeling behavior patterns and predicting future human activities quite tractable: If a behavior is repeated based on a human's past and current context states, then the state history should provide sufficient information for the their future activities during further repetitions of the routine behavior. In this sense, if one can define a context space to contain information that is descriptive enough to represent the connection between the human's state history, future states, and activities, then it becomes almost trivial to leverage this space to make predictions of future states based on observation histories. Then these state predictions can be used to compute the most probable future activity since the connections between the states and the activities have already been established by past observations from the context space.

4.1 A Generalizable and Causal Context Space

One of the core principles sought after through the work in this thesis is the generalizability of the approach, as discussed in Section 1. This means that the

context space, which is the most fundamental building block of a learning-based method to enable proactivity, should be generalizable across different task domains. This necessitates associating activities with a set of observations that are "universal" across domains, so that the underlying representation learning approach can be applied to different domains with minimal changes. As this universal space of observations, *time* and *space* are found to be the most fitting as these observations do not require prior information about the task domain and are effectively universal across any domain. Further, considering the aforementioned connections between the environmental cues, habitual behavior, and routines, the spatiotemporal cues can be argued to be strongly correlated with routine health-related behavior. Behaviors such as "brushing in the bathroom before bedtime," "grabbing a snack from the kitchen during TV time," or "going to the gym after work every other day," can all be examples of how these correlations can take shape. Notice how each statement can be defined by a set of $(activity, place, time)$. Drawing on this observation, positional time-series data–referred to as *trajectories* henceforth–and their association with activities are chosen to be used as the primary form of activity representations to be used within a spatiotemporal context space with the primary axes of $Time$ and $Position$. The regions within this space can then be *labeled* as necessary with associated activities to form the connections between the context and the activities. This way, the problem of enabling proactive engagement can be reduced to a trajectory prediction problem in its most fundamental and the most generalizable form. Moreover, using trajectories as the backbone of the routine representation learning, the solution then inherently becomes capable of making causal inference: Where an agent is right now has a strong influence on where it will be (and what it will be doing) in a moment.

As an example to illustrate causal connections in such a context space and the connections between the states and activities, a demonstration is provided in Figure 5. In this figure, a synthesized 3-day movement patterns of a hypothetical point human, who lives in a one-dimensional world and engages in daily routine behaviors, is shown. One may imagine that Days 1 and 2 are weekdays where the point human wakes up in their home (*Activity 1*), commutes to work (*Activity 2*), works (*Activity 3*), goes out for lunch (*Activity 4*), has lunch (*Activity 5*), and so on. The timings and the positions for each activity and activity transitions exhibit *minor variations* as one might expect to happen in real life (*Day 1* vs. *Day2*) as well as *major variations* (*Day 1* vs. *Day3*). On a weekend (*Day 3*), the agent exhibits a highly different behavior–a different routine–and goes out to the park (*Activity 9*) to have a picnic (*Activity 10*), and so on. Notice that the exact time and position scale of these activities and the axis labels are relatively irrelevant: Instead of daily commute behavior, these trajectories might as well be demonstrating in-home activities of the point agent or the movements of a linear actuator operating in a factory, addressing the generalizability principle of the proposed context space.

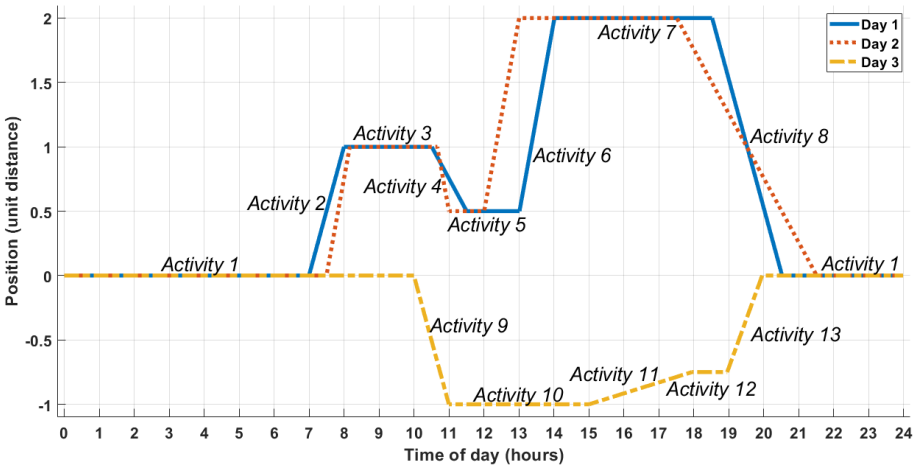

Figure 5: A toy example demonstrating how time and space can be associated with activities of a 1-dimensional point agent through a day for multiple days. The routine behavior has minor (Day 1 vs. Day 2) and major (Day 1 vs. Day 3) variations between each day.

4.2 Discussing Adaptability

Another core principle taken into account in this thesis is the adaptability of a proactive engagement system. Adaptability in this scope is considered twofold: User adaptability and task adaptability. Adapting to a specific user entails a model's ability to make inferences on a wide range of behavior routine patterns; across a wide range of environments where these behaviors are taking place. Whereas the task adaptability entails the models capacity to be useful in a wide range of intervention tasks.

The outstanding problem with adaptability, especially using deep learning-based parametric models, is that even the most "adaptable" model needs to have some prior knowledge of the underlying distribution that represents the variations in the observations. Even the approaches that propose the highest level of task adaptation based on models such as transfer learning or meta-learning require some sort of exposure to the full spectrum of variations in the underlying environment observations during the training time, and these models typically cannot be fine tuned if the observations during test time are significantly different from the training set. However, considering the level of variation individual's behaviors can have compared to one another, or how the environments that they are operating in might be different from one another, it is not practically feasible to assume access to a dataset that represents the whole distribution sufficiently during the training time before deployment. This effectively means such models will likely require a significant amount of re-training upon deployment, and it is questionable for how long an individual

should keep feeding their home robot with training data before it finally can understand to play along with their daily routines.

The proposed solution approach to enabling adaptability in proactive engagement is also twofold. First, instead of using a large parametric model, leveraging a probabilistic model for statistical inference is preferred. This type of models are typically much less data-hungry than any sort of deep learning approach. Using such a model, depending on how it is structured, it is possible to predict where a human is likely going to be in a relatively distant future compared to parametric model-based trajectory prediction methods, which typically offer up to several seconds of prediction horizon. Even though such horizon range may be more than enough and the models can successfully generalize over multiple scenarios in certain applications (e.g., autonomous driving), it is not sufficient for a human-robot co-habitation scenario where the system may benefit from much longer prediction horizons spanning minutes or hours.

The second aspect of adaptability–the task adaptability–suffers from a similar issue. But this time, instead of trying to tackle the possible variations in the user's behavior, the system needs to be able to handle the variations in the tasks assigned to the robot and the system's competences in handling them. In Figure 6, a stereotypical end-to-end learning model to generate proactive robot behavior is shown. Such models are trained with predefined set of competences (e.g., Activity Recognition, Object Recognition) and robot goals. What if one day, the user asks the robot to assist them with an activity that the robot was not trained to recognize? Or what if one day, the user buys a completely new coffee machine with a completely new form factor and the robot attempts to make some coffee before the user wakes up, but fails to locate the machine? Unless such a model was trained with every possible variation in the task domain prior to deployment–which is practically not possible–then it will likely require a full new training cycle to account for the users' previously unanticipated requests or futuristic coffee machines that do not have any buttons or knobs but fully operated using voice commands. Through leveraging a more modular approach to generating proactive behavior instead of using an end-to-end model, this level of adaptability can be handled in the architecture level rather than the model level. Such a modular architecture is proposed and further discussed in Section 4.4.

In the following subsections, the complete approach to the problem of opportunistic proactive engagement by leveraging modeling and prediction of routines

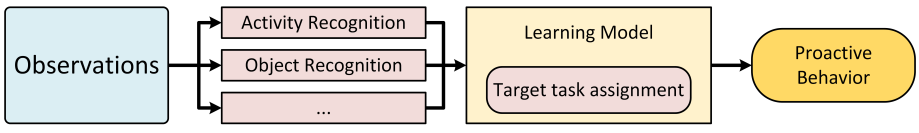

Figure 6: Conventional end-to-end approach to generating proactive agent behavior.

is described. Starting with a first look at the problem through probabilistic graphical models and discussing the shortcomings of such methods followed by a breakdown of the problem, a Gaussian process-based routine modeling method as well as a system-level architecture is proposed by the end of this section. Following the motivations and goals as outlined in Section 1, the proposed framework addresses the **Proactivity** problem by building on **Spatiotemporal Causality**, and offers an **Adaptable** machine learning method that is **Generalizable** across different task domains at different scales to learn routine behavior patterns .

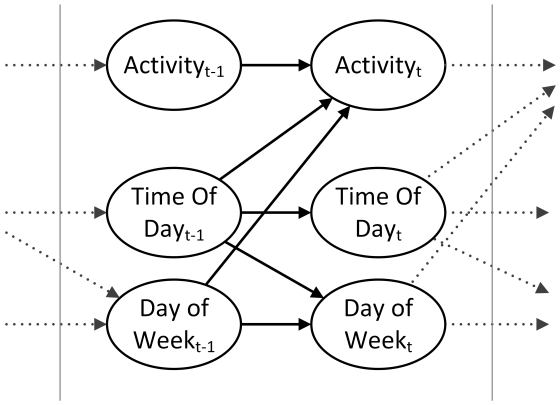

Figure 7: A two-timeslice Bayesian network for demonstrating a possible approach to model routine human behavior using probabilistic graphical models.

4.3 An Initial Solution Attempt with Probabilistic Graphical Models

The behavior patterns in Figure 5 can be represented using a probabilistic graphical model such as an N-timeslice Dynamic Bayesian Network (DBN) at the *activity level* . For example, a two-timeslice representation of the point person's behavior may take form as demonstrated in Figure 7. Notice that for a 7-day week, if the *Time of Day* was factored down into K slices (e.g., early morning, morning, noon, early afternoon, ...), then the transition probabilities between $Activity_{t-1}$ to $Activity_t$ can be represented by a table of size $K*7*13$-by-13 for the person in Figure 5. These probabilities can be computed by observing the human behavior *enough-many* times and then used to predict the likelihood of future agent activities. Notice that such model can be further simplified and/or it can be functionalized by actively reshaping the N-timeslice model with the observations of new transitions and activities.

Such a model offers a relatively straightforward way to keep track of activities and transitions that are causally informed by the observed states–or the *context* (e.g., [*Time of Day*, *Day of Week*, $Activity_{t-1}$] for this example). Moreover,

since it operates in a discretized state-space ($Activity \in \mathcal{A}$), it would not require any sort of prior knowledge of the system's internal dynamics nor necessitate an approximation of it. However, this approach inherently assumes access to an activity recognition tool[1]. This assumption, combined with the factorization of the observation space, limits such models' usefulness in two main ways: First, without a proper activity recognition tool, this approach cannot be used to learn any *useful* representations. Second, the approach cannot easily be transferred to other domains as this would require re-factorization of the observation space and access to another activity recognition tool. For example if the behavior demonstrated in Figure 5 represented the person's home routines instead of their daily commute routines, then we probably would have to have a finer-grained time factorization, and another tool that can recognize a completely different set of activities (e.g., cooking, watching TV, snacking, etc.). Although approaching the problem with graphical models provides a more formal insight on possible solutions and shows that a scalable, learning-based probabilistic inference model may be implemented using, for example, DBNs. However, such a model requires prior knowledge for factorization and assumes access to an activity recognition tool prior to behavior observations, which contradict with the adaptability and generalizability principles pursued in this work.

4.4 Structuring the Problem

The observation of dependence on activity recognition tools and the time fragmentation for human routine behavior modeling pose interesting questions: Which class of activities should such a model be aware of? Should this model be able to classify atomic actions (e.g., grasping an object) or more composite actions (e.g., making coffee) to be of the most use? How fine should the time discretization be (days, hours, seconds, years, ...)? Without being able to fully specify a target task for the intervention agent, one might argue that these questions do not have any good answers since the necessary tools and methods will differ greatly from task domain to task domain. Drawing on this observation, the generalizable approach sought after in this work aims to be as *task-agnostic* as possible in order to achieve the most reusability and and reprogrammability–in whichever domain it may have been deployed.

For example, in the context of routine behaviors, one might want to be able to model daily household routines as well as day-to-day commute behavior of a human, or the manipulation steps that a factory worker takes within workflows using certain machinery. Or, within the current scope of opportunistic health interventions case study, a robot might be tasked with detecting and intervening in unhealthy snacking behavior as well as engaging with a person who is *about to* skip an exercise session. Even though each of these tasks likely requires different data collection methods and activity recognition tools, one can contemplate the idea of modeling routine behaviors using a *representation* that is generalizable

[1] Note that activity recognition is a different branch of research and is out of scope of this dissertation.

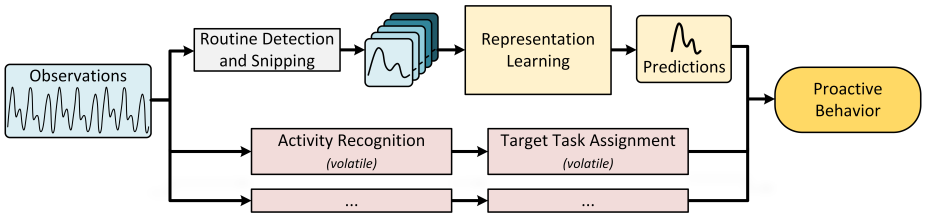

Figure 8: Proposed structure to approach the opportunistic interventions problem. The *volatile* blocks represent the functionality in such systems that are likely to change based on the intervention goals and/or the users' environment. The work proposed in this manuscript aims to focus on learning representations that are useful for a variety of intervention goals and environments.

across all such task varieties. Such an approach would provide a model that has a reusable behavior modeling *core*–or a *representation learning* model–instead of a task-specific end-to-end learning and modeling system. In Figure 8, the proposed modular system diagram for a proactive agent is demonstrated. Using such an architecture, the *volatile* blocks–the blocks that are domain- and task-dependent–can be detached from the behavior learning model. This way, it is possible for the learned representations to be useful for a variety of tasks with a possibility of applications in a variety of domains without needing to re-train an entire neural network in order to handle a task variation that was not available during the original training time.

The diagram in Figure 8 draws attention not only to the variability at the high-level tasks (home, commute, factory routines) as well as the nuances within the tasks. Thinking about the nature of the engagement (e.g., recommending a healthy food, suggesting setting a goal destination on a map automatically, fetching a tool in preparation of the next step in an assembly sequence), the target (intervention) tasks can be considered to be volatile and should have the flexibility of changing often or on demand. For example the agent might be required to recommend going out for a walk or suggest a different route to work today, or start boiling water for tea time. The activity recognition tools are also volatile and likely to change depending on the nature of the engagement, the environment, and the target task as well as technological advancements. Thus, instead of working for an end-to-end system, I believe it makes much more sense to work in a modular behavior modeling (or rather *representation learning*) backbone to maximize the utility of the approach for the widest variety of tasks and longest period of time possible.

For the representation learning part of this proposed architecture, it is intuitive to focus on time-series positional data as discussed earlier in Section 4.1. Such data, whether it is the location of a human in their house or out commuting; or the hand positions of a factory worker during an assembly step, carries

highly useful contextual information for the current and the future activities, referred to as *spatiotemporal causal context* henceforth. Unlike the output of any activity or object classification tool, the time and space are universal and are not subject to change depending on the task domain. This makes leveraging such information for a reusable and reprogrammable behavior modeling and intervention tool highly appealing, since making use of spatiotemporal context makes it possible to learn *useful* representations that one can build on in a fashion as demonstrated in Figure 8, without having the prerequisite of having access to an activity recognition tool, which is restrictive.

5 Representation Learning with Gaussian Processes

Considering how the modeling approach discussed in Section 4 relies heavily on the positional context (or simply the location of the human) and time, a design decision not to rely on discretization and to work in a continuous state-space is almost trivial to be made. To achieve this, the behavior of the human can be modeled as a Gaussian Process (GP), and the observations of the human behavior can be fed to a kernel learning algorithm to generate probabilistic models of transitions using Gaussian Process Regression (GPR). Notice that it does not make sense to use the trajectories in Figure 5 directly as is in such a model since the trajectory on Day 3 can be considered out of distribution of what we want to model for Days 1 and 2. This is because such a model, with a one-dimensional (1D) input (feature) space, only considers the time as the observed context variable and bases the inference on it, which is not *informative enough* to differentiate trajectories with major variations. However, if one chooses to consider the position itself as another context variable, thus, as another feature axis, then it is possible to make this distinction between the distributions governing Days 1, 2 and Day 3. In Figure 9, a demonstration of how the feature space can be reshaped for this purpose is given.

In Figure 9, the feature space is spanned by two axes, (*Time of Day*, *Current Position*), instead of time alone. And the output space represents the agent's position in 15 minutes into the future–measured from the current time and position. Note that the input points in this example are exactly the points given in Figure 5, and the output points are the identical position points as in this figure shifted with a constant time difference to represent a *prediction horizon*. In this sense, looking at such a three-dimensional (3D) plot from a top-down view towards the feature plane, Figure 9-bottom-right, yields the exact same plot as in Figure 5, and looking at it from other perspectives shows either the agent's position at the prediction horizon depending on the current time, Figure 9-bottom-left; or the relation between the future and the current positions, Figure 9-top-right.

This effectively forms a 3D inference space: with a two-dimensional (2D) feature space (the *context space*); and a 1D output space. In this space, the current

and the past context pushes the subsequent future contexts into different regions of the inference space. For example, to reach the region of *Activity 5* in the inference space, defined as the region around $(TimeOfday, Position) = (12, 0.5)$, the agent must have followed the trajectory through *Activity 1* and *Activity 2*; and taking this path moves the agent contextually far from *Activity 10*, which is visually apparent in Figure 9-top-left. In more clear terms, it is possible deduce whether the agent is more likely to be engaged with *Activity 5* or *Activity 10* at $t = t_i + 15minutes$ by checking the intersection of a ray perpendicular to the feature plane at $(t, x) = (t_i, x_{t_i})$ with the previously observed trajectory points. This means that it is possible to fit a GP model onto this reformed 3D space as the two feature axes can now be thought as two random variables that the data provides samples of and that the evidence is conditional to; and their distributions can be approximated by using GPR leveraging readily available kernel learning tools. The reasoning behind this approach is further discussed in detail in Section 5.1.

The primary benefit of using such a model is that it does not rely on access to an activity recognition tool or any other prior information about the user or the environment to be useful. It already inherits a representation of the human's behavior in terms of trajectories. These trajectories can possibly be segmented (e.g., by using frequency domain representation or simply by using 2D windows of timing and spatial information) and broken down into clusters of activity *regions*. These clusters of trajectory segments can then be used to inform a model such as a DBN. Considering our intervention goals such as recommending apples instead of potato chips during snack time, we are most likely interested in a very small subset of all possible agent activities. Cluster labeling can be custom tailored for each human using the system, depending on the users' goals. This can be done by using an activity recognition tool tailored for the task in conjunction with the behavior representation learning following the scheme shown in Figure 8, or maybe even simply through user input. For example, the agent might ask the user to let it know when they go to the kitchen to grab some snacks. Once the user does so, then the system can label the feature space region around the timing of this input as *Snacking Activity*, and assuming the representation learning model already had some trajectory observations from before, then it can immediately start predicting the intersections with the recently labeled feature space region during the future repetitions of the routines that include snacking behavior. Alternatively, this cluster labeling may be achieved using an additional recognition tool that can be *downloaded* or *updated* in the robot's firmware upon request, whenever needed, without requiring the robot to re-train a behavior representation model from scratch for each new request. This recognition tool might have as generic capabilities as recognizing a kitchen, or a pantry, or a cluster of snacks and outputting the locations of these entities. This location information can then be used to label the regions in the feature space with the associated activity using the readily existing trajectory representation data of the user. Such an approach should reduce the dependency on activity recognition and/or labeling by orders of magnitude compared to an end-to-end system since it only necessitates the labeling of the target in-

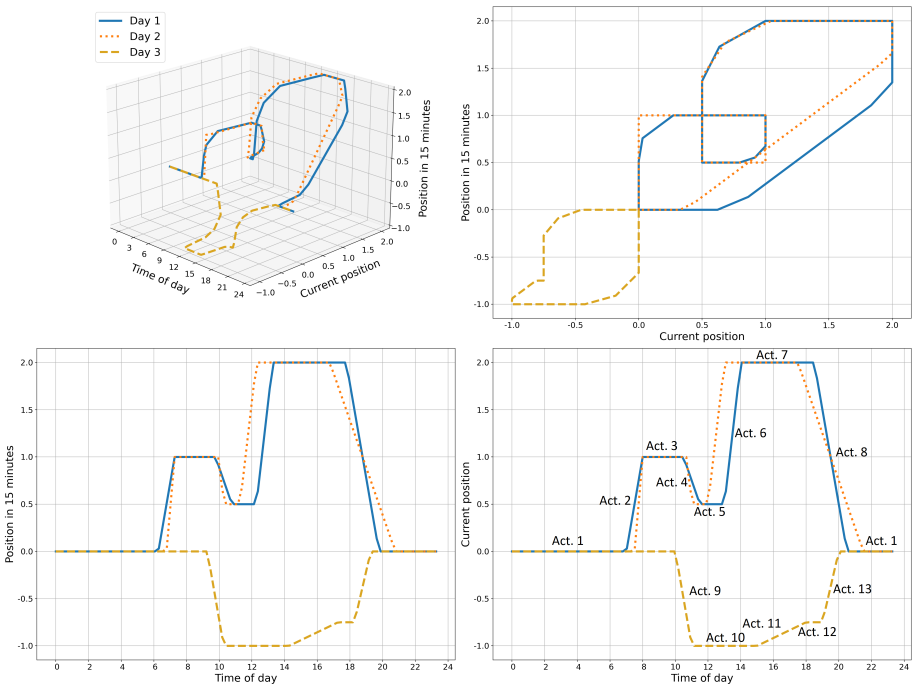

Figure 9: Reforming the example data points shown in Figure 5 into a new state-space representation. In this form, the *Time of Day* and the *Current Position* are treated as the feature axes, while *Position in 15 Minutes* is taken as the output axis. Doing so creates an inference surface that contains information about the future position of the point agent that is temporally causal which is useful to make trajectory predictions when there are major variations in the observed trajectories.

tervention activities instead of the whole spectrum of possible activities. Thus, it offers a higher flexibility in terms of the level of routine behavior abstraction it provides. A variety of continuous-time *routine behavior* data at different scales, such as day-to-day commute, in-home movements, hand positions during dish-washer operation, can be used effectively by the same framework after learning the associated kernel parameters.

Another significant advantage of GPs is that they are non-parametric and data-driven models. This makes it possible to represent the underlying data with very few training samples and make very few assumptions about the dynamics governing the modeled system. Effectively, they do not require training for thousands or millions of parameters to be able to generate useful representations. This makes it possible to implement a much more flexible system highly capable of adapting to variations in user behaviors and as well as environments and task

domains.

5.1 Gaussian Processes Regression for Modeling and Predicting (Long-term) Human Trajectories

A typical linear regression task investigates a noisy linear system

$$y = \mathbf{w}^T \mathbf{x} + \varepsilon \tag{1}$$

where $\mathbf{x}$ is the input to the system and y is the output such that the dataset $\mathcal{D}$ has the form

$$\mathcal{D} = \big\{(\mathbf{x}_1, y_1), ..., (\mathbf{x}_n, y_n)\big\} \tag{2}$$

and ε is the observation noise, typically modeled to have a Gaussian distribution with zero mean

$$\varepsilon \sim \mathcal{N}\big(0, \sigma_n^2\big) \tag{3}$$

The goal to such parametric regression task is to to find the parameters, $\mathbf{w}$, that describe the system the best using inference methods such as Maximum Likelihood Estimation (MLE) and maximizing the likelihood $P(\mathcal{D}|\mathbf{w})$; or Maximum A Posteriori estimation (MAP) maximizing the posterior probability $P(\mathbf{w}|\mathcal{D})$ over the parameter space, and compute the predictive distribution $P(y_*|\mathbf{x}_*; \mathbf{w})$, for a test point $(\mathbf{x}_*, y_*)$.

Gaussian Processes (GPs) offer a generalization to this probabilistic inference approach by leveraging an integration over the entire parameter space, which eliminates the need for specifying a single set of parameters for inference. Instead, GPs take every possible parameter combination into account and weighs them into the final model based on their posterior probability

$$P(y_*|\mathbf{x}_*, \mathcal{D}) = \int_{\mathbf{w}} P(y_*|\mathbf{x}_*; \mathbf{w}) P(\mathbf{w}|\mathcal{D}) d\mathbf{w} \tag{4}$$

Given that a Gaussian prior was chosen for $\mathbf{w}$, the resulting predictive distribution $P(y_*|\mathbf{x}_*, \mathcal{D})$ can be shown to be also a Gaussian distribution

$$P(\mathbf{y}_*|\mathbf{x}_*, \mathcal{D}) \sim \mathcal{N}\big(\mu, \mathbf{\Sigma}\big) \tag{5}$$

with mean μ and covariance $\mathbf{\Sigma}$.

This implies that if one can find a function to compute correlations between the training points and the test point(s) to effectively compute the covariance $\mathbf{\Sigma}$, the parametric regression problem can be generalized to a non-parametric one which considers infinitely many functions and define "a distribution over functions" [135].

While solving a GPR problem, the covariance matrix associated with the training and the test outputs (or "labels") ($\mathbf{y}$ and $\mathbf{y}_*$ respectively) is broken down in the following form

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}_* \end{bmatrix} \sim \begin{bmatrix} \mathbf{K} + \sigma_n^2 \mathbf{I} & \mathbf{K}_* \\ \mathbf{K}_*^T & \mathbf{K}_{**} \end{bmatrix} \tag{6}$$

where $\mathbf{K}$ is the covariance matrix for the training inputs, $\mathbf{X}$; $\mathbf{K}_*$ is the covariance matrix for the union of the training and the test inputs, $[\mathbf{X}, \mathbf{X}_*]$; $\mathbf{K}_{**}$ is the variance of the test inputs, $\mathbf{X}_*$. This matrix can then be used to compute the predictive distribution for any set of test points given the training points. Assuming that the data is subject to a Gaussian noise with variance σ_n^2, it can be shown that the final predictive distribution is then

$$\begin{aligned} P(\mathbf{y}_*|\mathbf{X}_*, \mathcal{D}) &\sim \mathcal{N}\big(\mathbf{K}_*^T(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}\mathbf{y}, \quad \mathbf{K}_{**} - \mathbf{K}_*^T(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}\mathbf{K}_*\big) \\ &\sim \mathcal{N}\big(\mu, \mathbf{\Sigma}\big) \end{aligned} \tag{7}$$

which not only gives a mean prediction (μ) for a given test point, but also provides a confidence region around that prediction ($\mathbf{\Sigma}$) that is extremely useful in applications where uncertainty has an impact on the outcome of the inference.

The individual elements in the matrices in Eq. 6 (and 7), which are effectively kernels, can be calculated using a kernel function such as a squared exponential kernel:

$$k(x, x') = \sigma^2 exp\bigg(-0.5\Big(\frac{x - x'}{l_s}\Big)^2\bigg) \tag{8}$$

where l_s is a parameter referred to as *lengthscale*, which defines how quickly the output can change with respect to the changes in the input data, and σ is a scaling factor.

Even though GPR is referred to as a "non-parametric" technique, there are three parameters required to define a GPR model. These are the noise variance σ_n^2, which has the dimensionality of the output space; and the length scale l_s and the scaling factor σ, which share the dimensionality of the input space for a 1D output space. Conveniently, similar to many other regression approaches, these parameters can be learned using some training data and a loss function based on (negative) marginal log likelihood.

Even in its raw form, a Gaussian process can be used to model trajectory dynamics and make confidence-bound predictions for previously unseen observations. Figure 10-Left demonstrates the output of a Gaussian process following a toy example of a point agent moving in a 1D world. This example takes the time axis (t) as the feature (input) space and the position (x) as the output space. So the modeling approach based on a Gaussian process $\mathcal{GP}$, makes the assumption that

$$f(t) \sim \mathcal{GP}(m(t), k(t, t')) \tag{9}$$

where f is the function of time defining the behavior of the point agent, m and k are the mean and covariance functions of the underlying Gaussian process, respectively. Using this process and the formulation summarized in Equations 6-8, one can plug in a set of (unseen) test inputs $\mathbf{t}_*$, which in this example is simply a fine quantization of the whole *t-axis* range, and finally compute the predictive distribution over this feature space

$$P(\mathbf{x}_*|\mathbf{t}_*, \mathcal{D}) \sim \mathcal{N}\big(\mu, \mathbf{\Sigma}\big) \tag{10}$$

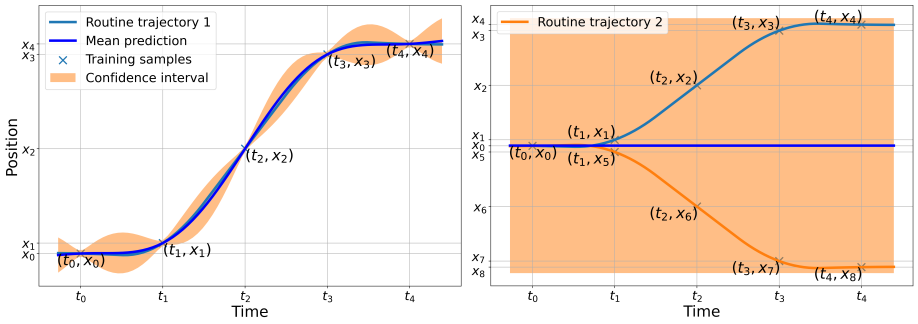

Figure 10: Gaussian process regression on two sets of data. **Left:** Regression is based on 5 observations from a single routine behavior trajectory. **Right:** Regression is based on two highly different routine trajectories. While the regression results are quite accurate in the case on the Left, the model cannot successfully fit in the case on the Right where the data has a high variance.

where $\mathcal{D} = \big\{(t_0, x_0), (t_1, x_1), (t_2, x_2), (t_3, x_3), (t_4, x_4)\big\}$ for the example in Figure 10-Left. The mean function (or the *prediction*) $\mu(t)$ and the confidence region which is computed using $\mathbf{\Sigma}$ is also demonstrated in this figure.

The primary issue with this raw form of Gaussian process regression arises when there are major variations in the agent's behavior (Figure 10-Right). The 1D point agent sometimes may choose to follow a trajectory (e.g., Figure 10-Right *Routine trajectory 1*) and a different one in other times (e.g., Figure 10-Right *Routine trajectory 2*). In this case, the predicted trajectory will not be successful in regressing onto neither of these ground truth trajectories.

An observation can be made here, with a reference to temporal *causality*, that the input **value** (t_2) is equally indicative of (x_2) and (x_6) and is not useful to make successful modeling/prediction. The **point** (t_2, x_2), however, is indicative of the future (x_3), and the point (t_2, x_6) is indicative of the future (x_7). Building on this observation, one can imagine to augment the 1D feature space into a 2D one and build the regression model on features, $\mathcal{F}$, from this augmented feature space similar to the one demonstrated in Figure 9. For the given system $f(t) = x_t$ and a fixed prediction horizon t_d, the new underlying system $f^\dagger(\cdot, \cdot)$ becomes:

$$\begin{aligned} f^\dagger(t, x_t) &= f(f^{-1}(x_t) + t_d) + \epsilon_n \\ &= x_{t+t_d} + \epsilon_n \\ \epsilon_n &\sim \mathcal{N}(0, \sigma_n^2) \end{aligned} \tag{11}$$

Given that t_d is the fixed time difference between the observations in the example

in Figure 10, the system $f^{\dagger}(\cdot,\cdot)$ will have the following data points:

$$\mathcal{F} = \begin{bmatrix} t_0 & x_0 \\ t_1 & x_1 \\ t_2 & x_2 \\ t_3 & x_3 \\ t_0 & x_0 \\ t_1 & x_5 \\ t_2 & x_6 \\ t_3 & x_7 \end{bmatrix}, \mathcal{X} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \end{bmatrix} \tag{12}$$

This augmented representation of the original dataset, $\mathcal{D}$, is not going to suffer from the issue demonstrated in Figure 10-Right at the cost of adding an extra dimension to the underlying Gaussian process. A nice property of Gaussian processes is that it is fairly straightforward to apply them to higher dimensional input-output spaces. An example visualization of data augmented this way is given in Figure 9-top-left. And Figure 11 demonstrates the outcome of GPR on the example dataset from Equation 12. The two confidence regions depicted in this image reside on the slices from the distribution surface created by the 3D data points. These slices are created by computing the intersection of the surfaces created by *Routine trajectory 1* and *Routine trajectory 2* with the distribution surface, and projecting this intersection on the $(Time, FuturePosition)$ plane. Notice that in Figure 11 at time t_0, the agent is at x_0 (not visible in the figure) and predicted to be on either x_1 or x_5 in the next time step. Depending on whether the agent moves to x_1 or x_5 at t_1, the model then can successfully predict that the agent is going to be either at x_2 or x_6 in the next time step, respectively.

Even though this example delivers the point of how the feature space can be altered to get the process to be able to handle major variations in the observations, in a real-world application, using a fixed horizon t_d might be problematic as the data may not have the exact granularity needed or it may prove to be restrictive as the required prediction horizon might be significantly different than t_d for each prediction task. A possibly finer approach to augmenting the state space to handle the variations in observed trajectories can leverage a continuous horizon *axis*, such that the original system $f(t)$ can be modified to

$$\begin{aligned} f^{\ddagger}(t, x, h) &= f(f^{-1}(x_{t+h})) + \epsilon_n \\ &= x_{t+h} + \epsilon_n \end{aligned} \tag{13}$$

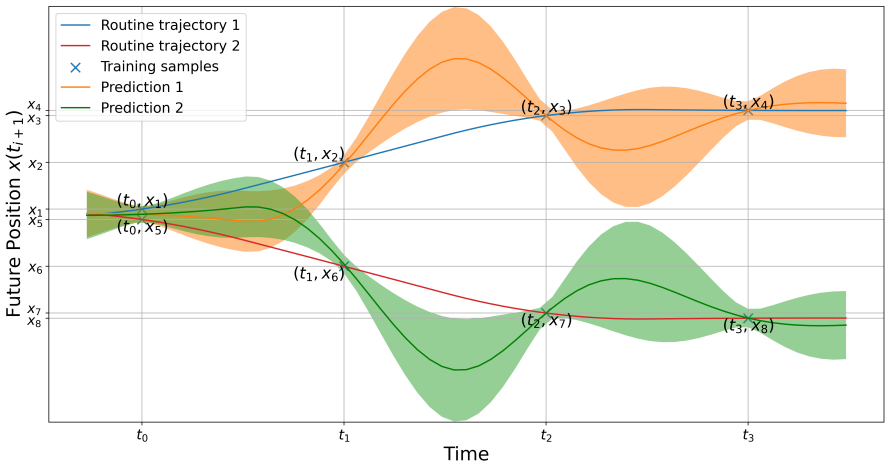

Figure 11: 2D projection of the Gaussian process regression after re-structuring the feature space. The model is now able to distinguish between the two dissimilar routines.

and the underlying feature and output matrices become

$$\mathcal{F} = \begin{bmatrix} t_0 & x_0 & 0 \\ t_0 & x_0 & h_0 \\ t_0 & x_0 & h_1 \\ \dots & \dots & \dots \\ t_1 & x_1 & 0 \\ t_1 & x_1 & h_0 \\ t_1 & x_1 & h_1 \\ \dots & \dots & \dots \\ . & . & . \end{bmatrix}, \mathcal{X} = \begin{bmatrix} f(t_0) \\ f(t_0 + h_0) \\ f(t_0 + h_1) \\ \dots \\ f(t_1) \\ f(t_1 + h_0) \\ f(t_1 + h_1) \\ \dots \\ . \end{bmatrix} \tag{14}$$

for any given horizon $h = [h_0, h_1, ...]$.

Assuming $\mathcal{F}$ and $\mathcal{X}$ are used as the input and output training data, respectively, for the example data in Equation 14, the GPR model's parameters can be learned to compute a distribution over a predicted trajectory, $\mathbf{x}_*$, using a desired horizon $\mathbf{h} = [h_0, h_1, h_2, ...]^T$ and current position x_c. Adjusting the notation of the Equation 7 for this example, this predictive distribution is then

$$P(\mathbf{x}_*|\mathbf{F}_*, (\mathcal{F}, \mathcal{X})) \sim \mathcal{GP}\Big(\mathbf{m}\big(t, x, h\big), \mathbf{k}\big((t, x, h), (t', x', h')\big)\Big) \tag{15}$$

where,

$$\mathbf{F}_* = \begin{bmatrix} ... & ... & ... \\ t_{c-2} & x_{c-2} & 0 \\ t_{c-1} & x_{c-1} & 0 \\ t_c & x_c & 0 \\ t_c & x_c & h_0 \\ t_c & x_c & h_1 \\ t_c & x_c & h_2 \\ ... & ... & ... \end{bmatrix} \tag{16}$$

and the horizon axis in this feature matrix can be collapsed onto the time axis to be paired with the resulting mean timeseries trajectory prediction over the horizon for each data point in the input data, $\mathbf{F}_*$,

$$\begin{aligned} FutureTrajectory = [&(t_c + 0, \mu_{x_*|t_c,x_c,0}), \\ &(t_c + h_0, \mu_{x_*|t_c,x_c,h_0}), \\ &(t_c + h_1, \mu_{x_*|t_c,x_c,h_1}), \\ &(t_c + h_2, \mu_{x_*|t_c,x_c,h_2}), \\ &(t_c + h_3, \mu_{x_*|t_c,x_c,h_3}), \\ &...] \end{aligned} \tag{17}$$

This resulting trajectory prediction will

1 Be based on the GP parameters obtained during training on previously observed trajectories

2 Retain and leverage the correlation information that already exists in the model's covariance matrix, $\mathbf{K}$

3 Be further informed by the history of the currently observed trajectory (top half of $\mathbf{F}_*$) since the data points in $\mathbf{F}_*$ will be used to compute $\mathbf{K}_*$, and this will help handle minor variations in the current trajectory compared to the previously observed ones

4 Be able to query the future trajectory based on every the data point available to the model in $\mathcal{F}$, $\mathcal{X}$, and $\mathbf{F}_*$

5 Result in a completely data-driven (long horizon) trajectory prediction even after observing and training on a **single** routine demonstration.

This reformed data space can be used in Gaussian process regression to predict long-horizon human trajectories without any theoretical limitations on the horizon length. The only practical limitation is the availability computational resources that are able to keep data matrices $\mathcal{F}$, $\mathcal{X}$, and $\mathbf{F}_*$ in the memory. This limitation can be managed by appropriately down-sampling routine trajectories while making sure the sample size is sufficient to represent the underlying data (e.g., sampling at Nyquist rate), or by using approximate Gaussian process methods [52]. When there are major variations in the observed routine

trajectories, if a current trajectory resembles any of the past observations during inference time, the model will be able to untangle it from other dissimilar behaviors without a need for a tool that is able to distinguish and cluster dissimilar trajectories and training several different models for each distinct behavior. Moreover, the minor variations that might be occur between otherwise identical routines can be modeled by the noise term, σ_n^2, of the GPR model, and the model then can generalize across these variations with the learned distribution.

A final consideration before making this approach feasible for a practical application is the input/output dimensions. So far the examples given have followed a 1D point agent. However, this is not practical for a real-life trajectories, which are 2D, at best, assuming the goal agent is moving on a planar space. To handle this, a trivial extension to the original function, $f(t)$, can be applied to consider it as a mapping from time to 2D coordinates instead of 1D coordinates:

$$f(t) \rightarrow (x, y) \tag{18}$$

and following the same notion of leveraging the current position and a horizon axis as well as the time, we can construct the new input/output space to represent a mapping from time, current position and the horizon to 2D coordinates:

$$f^{\ddagger}(t, x, y, h) \rightarrow (x, y) \tag{19}$$

which presents a "multi-task" GPR since now the output is 2D.

There are methods available in the literature to handle the multi-task GPR problem [9]. After observing during the initial experimentation with synthesized data that the multi-task approach did not result in any significant performance improvements for the application intended in this thesis work, it was decided to use two GPR models for each output dimension of the regression problem for simplicity. The two predictive probability distributions associated with each GP for $x-$ and $y-axis$ outputs of the trajectory predictions are

$$\begin{aligned} P(\mathbf{x}_*|\mathbf{F}_*, (\mathcal{F}, \mathcal{X})) &\sim \mathcal{GP}_x\Big(\mathbf{m}\big(t, x, y, h\big), \mathbf{k}\big((t, x, y, h), (t', x', y', h')\big)\Big) \\ P(\mathbf{y}_*|\mathbf{F}_*, (\mathcal{F}, \mathcal{Y})) &\sim \mathcal{GP}_y\Big(\mathbf{m}\big(t, x, y, h\big), \mathbf{k}\big((t, x, y, h), (t', x', y', h')\big)\Big) \end{aligned} \tag{20}$$

which share the same input data tensors, $\mathcal{F}$, but separate GP parameter sets $(\sigma_n^2, \mathbf{l}_s, \sigma)_x$ and $(\sigma_n^2, \mathbf{l}_s, \sigma)_y$ learned by training on separate output tensors $\mathcal{X}$ and $\mathcal{Y}$, respectively. The resulting mean predictions μ_{x_*} and μ_{y_*} are then combined to create a final 2D trajectory prediction:

$$\begin{aligned} \mathit{Future\ 2D\ Trajectory} = [&(t_c + 0, \mu_{x_*|t_c,x_c,y_c,0}, \mu_{y_*|t_c,x_c,y_c,0}), \\ &(t_c + h_0, \mu_{x_*|t_c,x_c,y_c,h_0}, \mu_{y_*|t_c,x_c,y_c,h_0}), \\ &(t_c + h_1, \mu_{x_*|t_c,x_c,y_c,h_1}, \mu_{y_*|t_c,x_c,y_c,h_1}), \\ &(t_c + h_2, \mu_{x_*|t_c,x_c,y_c,h_2}, \mu_{y_*|t_c,x_c,y_c,h_2}), \\ &(t_c + h_3, \mu_{x_*|t_c,x_c,y_c,h_3}, \mu_{y_*|t_c,x_c,y_c,h_3}), \\ &...] \end{aligned} \tag{21}$$

5.2 Algorithms

In this section, the the algorithms used during the implementation of the GPR-based trajectory prediction described in the previous section are given in pseudo-code. Note that these algorithms are given assuming a 1D output space, instead of a 2D one, for notational simplicity.

To begin with, Algorithm 1 demonstrates the overview of the method used to learn the kernel parameters for the GPR. The input for this algorithm is a dataset of trajectories which are not "clipped" into "snippets" or segments of routines for the sake of completeness. During the experiments, the dataset was already collected in segments of routines rather than a full trajectory spanning multiple routines. The clipping subroutine within this context can be exemplified by assuming the dataset $\mathcal{D}$ is a week-long trajectory observation of an agent, and each routine segment R corresponds to individual days within this dataset.

Algorithm 1 Learning Gaussian Process Regression kernel parameters

Require:
Timeseries data $\mathcal{D}$
Maximum training loss threshold $loss_{max}$
Ensure:
Optimal lengthscales, l_s; scaling factors, σ; noise level σ_n
0: **procedure** LEARNKERNELPARAMETERS($\mathcal{D}, loss_{max}$)
1: $\{R_1, R_2, ...\} \leftarrow ClipRoutineSnippets(\mathcal{D})$
2: $\{S_1, S_2, ...\} \leftarrow DownSampleSnippets(\{R_1, R_2, ...\})$
3: $(\mathcal{F}, \mathcal{X}) \leftarrow FeaturizeSnippets(\{S_1, S_2, ...\}, \{R_1, R_2, ...\})$
4: **while** $loss > loss_{max}$ **do**
5: $loss = -MarginalLogLikelihood(\mathcal{GP}(\sigma, l_s, \sigma_n|\mathcal{F}), \mathcal{X})$
6: $(\sigma, l_s, \sigma_n) \leftarrow loss.backward(\sigma, l_s, \sigma_n)$
7: **return** (σ, l_s, σ_n)

Assuming that the segments $\{S_i\}$ are composed of data points sampled at a frequency much higher than necessary, then it is good practice to down-sample them to mitigate issues that may rise due to data tensors growing beyond practical computational limits. Although there can be much smarter ways to conduct this down-sampling operation, such as by first finding the "key points" in the trajectories that would result in the best representation of the trajectory with minimum amount of data points, a simpler approach using equidistant sampling in both time and position dimensions was preferred since maximum representational efficiency was not a primary point of concern in this work. This equidistant down-sampling method was assumed to provide the *sufficient* representation of the full trajectory, and is provided below in Algorithm 2 for the sake of completeness.

The "featurization" mentioned in Algorithm 1 Line 3 refers to the transformation of the original input-output space to create the spatiotemporal feature

Algorithm 2 Down-sampling subroutine used for the experiments

Require:
Raw routine snippet $R = \big\{(t_0, x_0), (t_1, x_1), ..., (t_n, x_n)\big\}$
Time quantization t_q
Position quantization x_q

Ensure:
Down-sampled snippet S

0: **procedure** DownSampleSnippet(R,t_q,x_q)
1: $S \leftarrow \{(t_0, x_0)\}$ ▷ Initialize snipped trajectory
2: $t_{last} = t_0$
3: $x_{last} = x_0$
4: **for** $i \in [1, ..., n]$ **do**
5: $t_d = t_i - t_{last}$ ▷ Time difference
6: $x_d = x_i - x_{last}$ ▷ Position difference
7: **if** $(t_d > h_{tq})$ OR $(x_d > h_{xq})$ **then**
8: $S \leftarrow S \cup \{(t_i, x_i)\}$
9: $t_{last} = t_i$
10: $x_{last} = x_i$
11: **return** S

space points, $\mathcal{F}$, and the output points, $\mathcal{X}$, as shown in Equation 14. The pseudo-code in Algorithm 3 demonstrates the computation of $\mathcal{F}$ and $\mathcal{X}$. In the simplest terms, this algorithm looks over a fixed horizon from each sample point in the down-sampled segments, and uses the same equidistant sampling approach used in the down-sampling subroutine to sample additional points from this horizon that will populate the elements in $\mathcal{F}$ where $h \neq 0$ (Equation 14).

Once the dataset is featurized and the GP parameters are learned using this augmented dataset, then the model is ready to make trajectory predictions based on new–unseen–trajectory observations following the steps given in Algorithm 4. This algorithm demonstrates how new observations can be added to the training feature dataset $\mathcal{F}$ repeatedly to inform the model's predictions with the most recent observations. This is done solely by recomputing the covariance matrices $\mathbf{K}$, $\mathbf{K}_*$, and $\mathbf{K}_{**}$ based on these observations, **without** re-training the model parameters. The input feature matrix for the prediction, $\mathbf{F}_*$, is simply created by appending the desired future time points [h_1,h_2,...] to the most recent observation (t_i, x_i), appending the resulting vector into the feature matrix as rows. The final feature matrix's first and second columns are then t_i and x_i for all rows, and its last column is a series of points representing the prediction horizon $[0, h_1, h_2, ..., h_{max}]^T$. The predictive distribution computed by the $\mathcal{GP}$ for this input feature matrix will then represent the distribution of the possible trajectories that the agent will take following the point (t_i, x_i) and will have the form $\big[(t_i + h_j, x_j)\big]$, where j indexes the rows of the feature matrix.

Since the GPR generates a predictive distribution over a set of trajectories, it can be used to measure *the quality of fit* of a recently observed trajectory

segment to this distribution based on the previous observations. If an observed segment lies outside the confidence bounds of the current distribution, then it means that this segment either represents a new major routine variation, or it is a previously observed routine with a minor variation but the model's noise parameter cannot handle this level of variation well-enough yet. This observation can be used to create triggers to tell the system that it likely needs more training based on the newly acquired data, and is a first glance at how the proposed method can be used for "lifelong learning" purposes after it's been deployed in an environment. Even though this feature of GPR is not made use of during the experiments within the scope of this thesis, a pseudo-code on how this can be done is given in Algorithm 5 for the sake of completeness.

Algorithm 3 Feature space transformation subroutine

Require:

Raw snippet $R = \{(t_{r0}, x_{r0}), (t_{r1}, x_{r1}), ..., (t_{rn}, x_{rn})\}$
Down-sampled snippet $S = \{(t_0, x_0), (t_1, x_1), ..., (t_k, x_k)\}$
Prediction horizon $h_{max} < t_{rn}$
Horizon time quantization h_{tq}
Horizon position quantization h_{xq}

Ensure:

A set of features F and labels X for training and inference

procedure FEATURIZESNIPPET($S, R, h_{max}, h_{tq}, h_{xq}$)
  $F \leftarrow \{\ \}$ ▷ Features from snippet S
  $X \leftarrow \{\ \}$ ▷ Labels from snippet S
  for $i \in [0, ..., k]$ **do**
    $F \leftarrow F \cup \{(t_i, x_i, 0)\}$
    $X \leftarrow X \cup \{x_i\}$
    $t_{last} = t_i$
    $x_{last} = x_i$
    $j = Index(Where(S[i] == R[\cdot]))$
    while $t_d < h_{max}$ **do** ▷ Iterate over the raw trajectory
      $j = j + 1$
      $t_d = t_{rj} - t_{last}$
      $x_d = x_{rj} - x_{last}$
      if $(t_d > h_{tq})\ \ OR\ \ (x_d > h_{xq})$ **then**
        $F \leftarrow F \cup \{(t_{rj}, x_{rj}, t_d)\}$
        $X \leftarrow X \cup \{x_{rj}\}$
        $t_{last} = t_{rj}$
        $x_{last} = x_{rj}$
  return (F, X)

Algorithm 4 Predicting trajectories given a GPR

Require:
Timeseries data stream $TS(t) \rightarrow x$
Desired prediction horizon h_{max}
GP parameters $(\sigma_n^2, \mathbf{l}_s, \sigma)$
Past featurized observations $\mathcal{F}$

procedure PREDICTTRAJECTORY($\mathcal{GP}, \mathcal{F}, h_{max}$)
  $\mathbf{F}_* \leftarrow \{\ \}$
  for each data point (t_i, x_i) in data stream **do**
    $\mathcal{F} \leftarrow \mathcal{F} \cup \{(t_i, x_i, 0)\}$ ▷ Add latest observation to data tensors
    $\mathbf{F}_* \leftarrow \{(t_i, x_i, 0), (t_i, x_i, h_1), (t_i, x_i, h_2), ..., (t_i, x_i, h_{max})\}$
    Recompute $\mathbf{K}, \mathbf{K}_*, \mathbf{K}_{**}$ using $\mathcal{F}$ and $\mathbf{F}_*$
    $P\Big(\mathbf{x}_*|\mathbf{F}_*, \big(\mathcal{F}, (\sigma_n^2, \mathbf{l}_s, \sigma)\big)\Big) \leftarrow \mathcal{GP}(\cdot)$
    $T_{pred} \leftarrow \big[(t_i + 0, \mu_{x_*|t_i,x_i,0}),\ \ldots\ , (t_i + h_{max}, \mu_{x_*|t_i,x_i,h_{max}})\big]$

Algorithm 5 Lifelong learning for the proposed GPR approach

Require:
Timeseries data stream $TS(t) \rightarrow x$
Prediction horizon h_{max}
GP parameters $(\sigma_n^2, \mathbf{l}_s, \sigma)$
Past featurized observations $\mathcal{F}$

while forever **do**
  $S \leftarrow \{\}$ ▷ Current trajectory segment
  while current segment not over **do**
    for each data point (t_i, x_i) in current segment data stream **do**
      $S \leftarrow S \cup (t_i, x_i)$
      $(P,\ T_{pred}) \leftarrow PredictTrajectory(\mathcal{GP}, \mathcal{F}, h_{max})$
      if S out of confidence bounds of P **then**
        re-train $\leftarrow$ true
  $\mathcal{D} \leftarrow \mathcal{D} \cup S$
  if re-train **then**
    do in parallel
      $LearnKernelParameters(\mathcal{D})$
    re-train $\leftarrow$ false
  continue

5.3 Proof of Concept Experiments using Synthesized Data

The GPR backbone of behavior modeling and prediction framework described throughout this section was implemented using GPyTorch open-source library [38]. First, an initial prototype was created to prove the functionality of the concept using synthesized data. This data composed of several 2D routine trajectories with major variations, which were generated using linear interpolation between a series of arbitrarily defined "keypoints," which represented the main

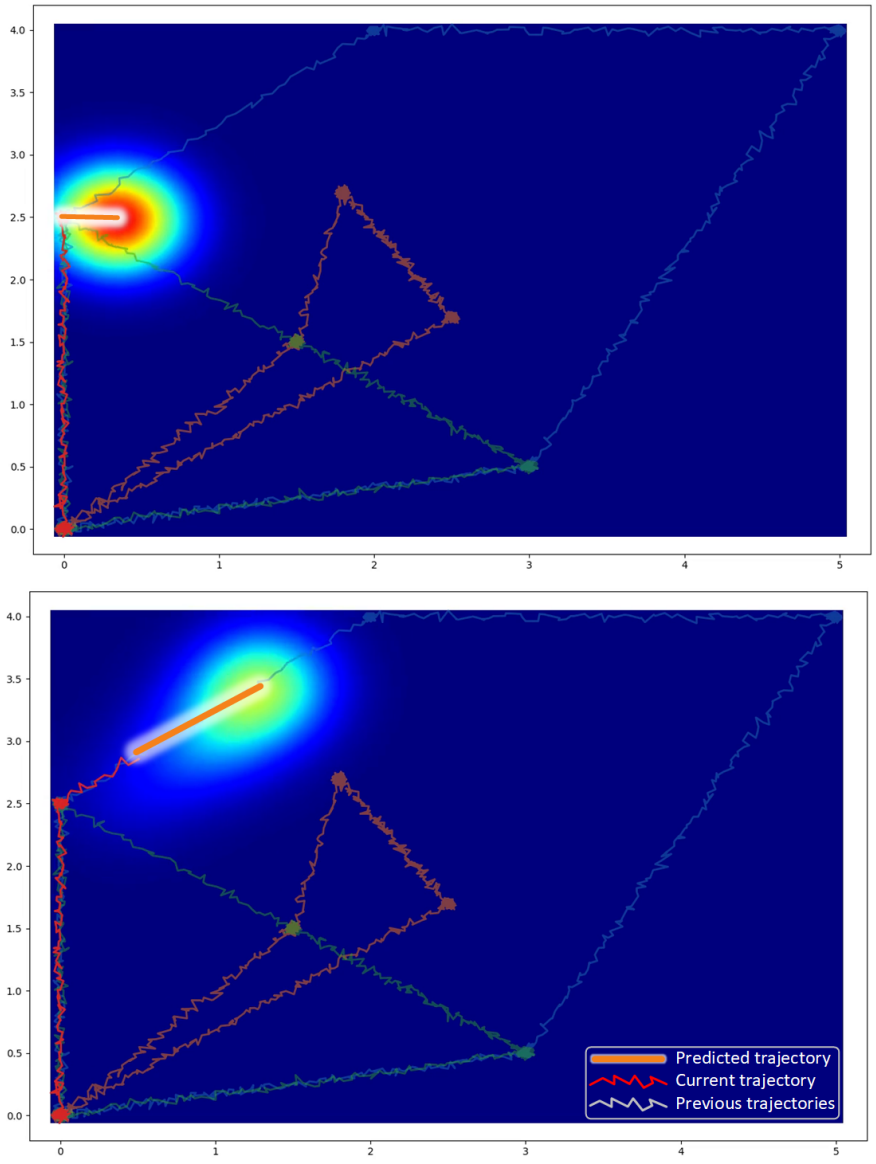

Figure 12: Demonstration of predicted trajectories and the probability heatmaps associated with the predicted position at the end of the current prediction horizon. The Gaussian process model parameters were learned and the covariance matrix $\mathbf{K}$ was computed based on previously observed trajectories. The "current trajectory" was not seen during training. **Top:** The agent is at an intersection and the model cannot decide which direction the agent will go and makes the average prediction of the two options. **Bottom:** Once the agent starts moving one way from the intersection, the model then immediately figures out which routine the current trajectory belongs to and corrects its prediction.

activity regions. These hand-crafted keypoints were designed to provide a deeper understanding of how the proposed modeling approach would behave in various scenarios. The interpolation between these points represented the mobility behavior between activities. A Gaussian noise was added on these trajectories to replicate minor variation sources, such as sensor noise. Additionally, each routine trajectory was regenerated multiple times with *activity noise* to simulate further minor routine variations. This noise was based on hand-crafted rules which incorporated randomness, and would manipulate (1) the agent's velocity between two consecutive keypoints, (2) the time that the agent spends on a certain keypoint, and (3) the time at which the agent arrives at a keypoint. These trajectories were used to train the GPR model, while other similarly generated trajectories were used to create and evaluate trajectory predictions. Figure 12 demonstrates probability heatmaps associated with the predicted future position of an agent in a 2D space. I.e.,

$$p\Big((t_{now}, x_{now+t_d}, y_{now+t_d}) = (x, y)\big|\mathbf{F}_*, \mathcal{F}, \mathcal{GP}\Big) \tag{22}$$

along with predicted trajectories that would get the agent there

$$\mathbb{E}\big[P|\mathbf{F}_*, \mathcal{F}, \mathcal{GP}\big], \forall t_d \in [0, t_{horizon}] \tag{23}$$

The analysis and observations from these experiments show that the model prototype is able to model both the minor and major variations in the routine trajectories. Possibly the most critical key takeaway from these experiments is the model's ability to distinguish between two routines. Figure 12-top demonstrates the agent at an intersection between two possible trajectories observed in the training data. At this moment, the model is not able to figure out which way the agent is going to go next, as expected. However, once a signal is received from the agent to indicate either one of the possible routes it can take in the form of a slight movement toward one of the routes, then the model can successfully predict the rest of this trajectory as demonstrated in Figure 12-bottom.

5.4 Proof of Concept Experiments using Motion Capture Data

As the next step for the proof of concept studies, real-life human motion data was used. Among several possible options to acquire human position data–such as using the robot's sensors or using an RGB-D-based pose estimation system–a motion capture system (MoCap) was chosen. This decision was based on MoCap systems' being significantly more precise and robust against lighting variations and occlusions, as well as their ability to provide higher "frame rates." The data acquisition environment, demonstrated in Figure 13, was set up using two HTC VIVE Pro virtual reality base stations and an HTC VIVE Tracker 3.0 [26]. The tracker was attached to a belt worn by the human during the experiments, and the 6-DoF pose information of the tracker was streamed to a computer using SteamVR and OpenVR API. This stream is used for real-time tracking,

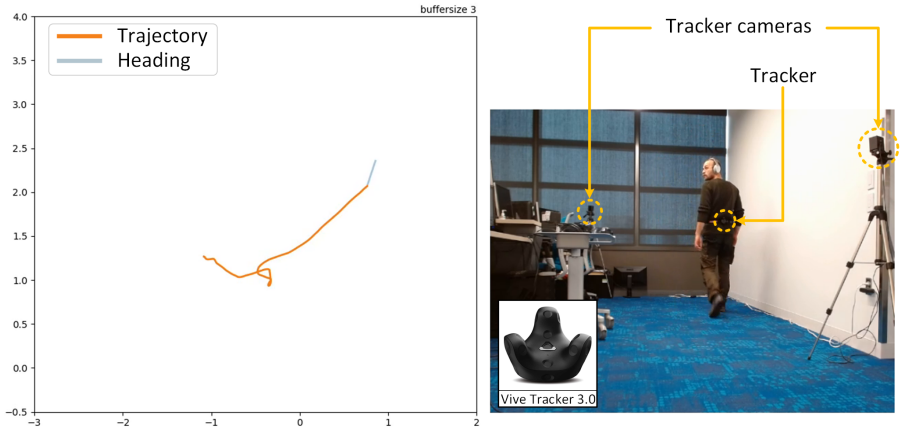

Figure 13: **Left:** The trajectory and the heading tracked by the motion capture system displayed in real-time. **Right:** Real-life in the meantime.

behavior modeling, and predictions as well as logging the data locally for later analyses.

Figure 14 shows the real-time predicted trajectories of a human wearing the MoCap belt and moving around the lab space on trajectories similar to previously observed ones[2]. In this example, the human followed three distinct patterns: (1) make a full counterclockwise rotation around the room (CCW), (2) start moving clockwise into an intersection, take a short pause, and finish the rotation counterclockwise (CW-Pause-CCW), and (3) start moving clockwise into the intersection, take a long pause, and finish the rotation clockwise (CW-LongPause-CW). In Figure 14(b), the model immediately predicts that this is a counterclockwise pattern once the human starts moving counterclockwise. In Figure 14(c), the human follows a clockwise path, makes it to the intersection point, and takes a pause. Once they start moving counterclockwise, the model can then predict that the rest of this pattern will be counterclockwise. In Figure 14(d), the human follows the same clockwise path into the intersection but this time takes a longer pause. The model correctly predicts based on the pause duration–even before the human starts moving–that this pattern will finish with a clockwise movement. Each trajectory in this example took between 28 and 42 seconds, and the prediction horizon was 30 seconds. Note that the proposed model was shown to be able to handle much longer predictions (>400 s) in later experiments (Figures 18).

[2]Code and video demonstrations are available at `github.com/terzioglan/GPTrajectoryPrediction`

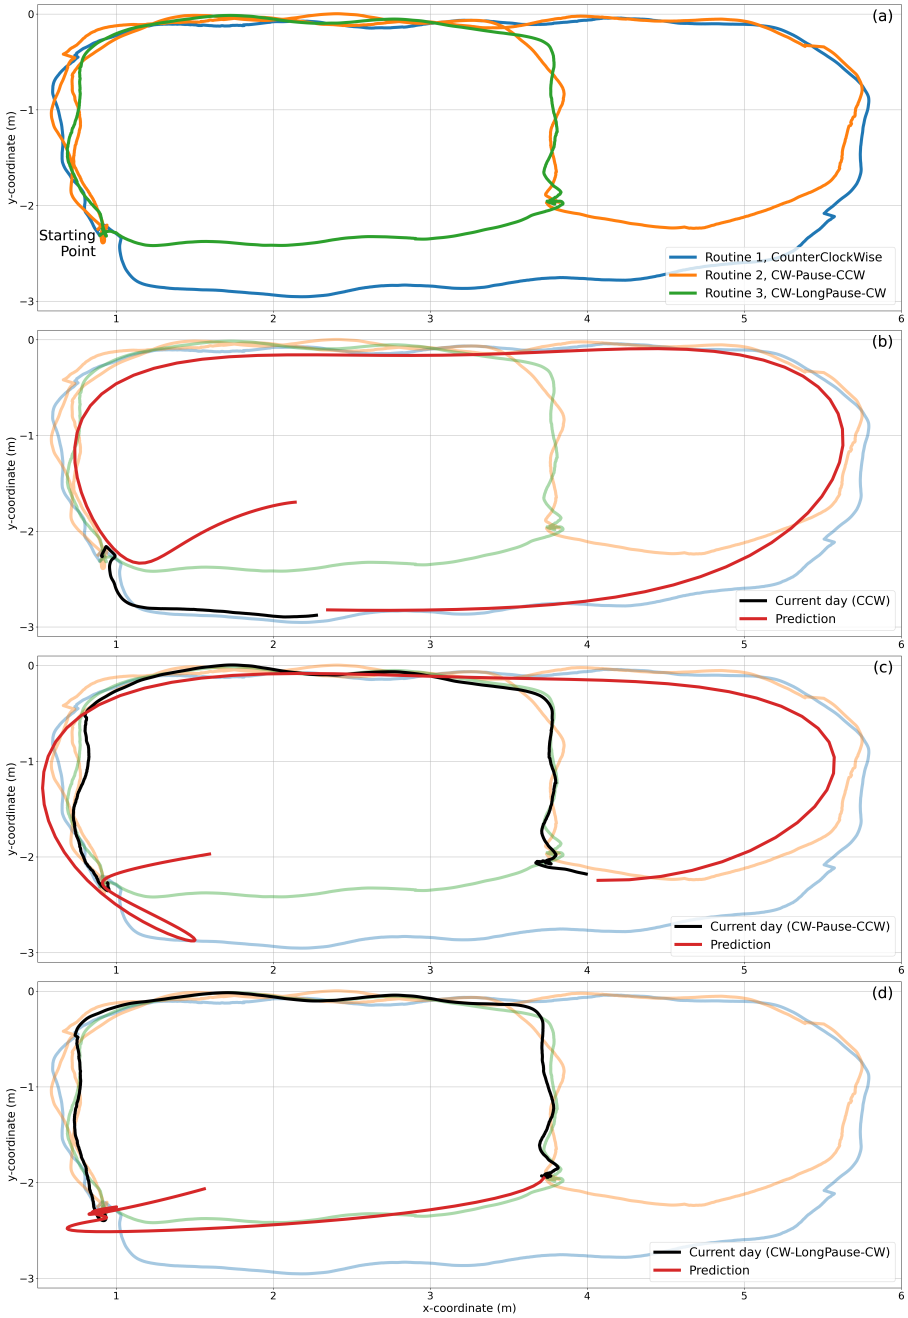

Figure 14: Demonstration of real-time trajectory prediction using the MoCap system. (a)Three trajectories following different path patterns are used to learn the Gaussian process parameters. (b-d)The proposed model was able to untangle these variations and make correct predictions for all three movement patterns in previously unseen trajectories.

The trajectory predictions on the MoCap data stream were performed at

100 Hz, which was the refresh rate of the tracker cameras (i.e., "Vive Base Stations"). The Gaussian process regression, covariance-matrix computation for each incoming data segment, and visualization of logged data and predictions were handled simultaneously with the prediction computation on separate processing threads running on an ordinary desktop computer equipped with an NVIDIA GeForce RTX4060Ti 8GB GPU. This same system was later used in the summative evaluation user studies (Section 6), in which the robot learned the participants' movement patterns within an artificial office space and moved proactively to engage in snack/drink recommendation interactions following the system structure presented earlier in Figure 8. Data collected during pilot studies with this system were also used to compute robot interaction poses that formed jointly-focused F-formations [71, 112, 72, 25, 73] with the human and the target item for delivering recommendations.

6 Summative Evaluation Study

To evaluate the effectiveness of the proposed trajectory prediction-based proactive robotic intervention framework, we designed a within-subjects user study and a simulated office environment (Figure 15). In this environment, the participants were prompted to imagine that this was their own office space, where their employer had deployed a robot to "accompany employees" during daily activities. Any further information regarding the robot's function was withheld from the participants until the study concluded. Each participant was asked to simulate 10 days of office activities within this space following a variety of schedules that they wrote on small cards before beginning to simulate their days. Through 5 of these 10 days, the robot's behavior was *Reactive*, which were the Control condition trials for this experiment; and in the other 5, the robot's behavior was *Proactive*, which served as the intervention condition trials (Section 6.5).

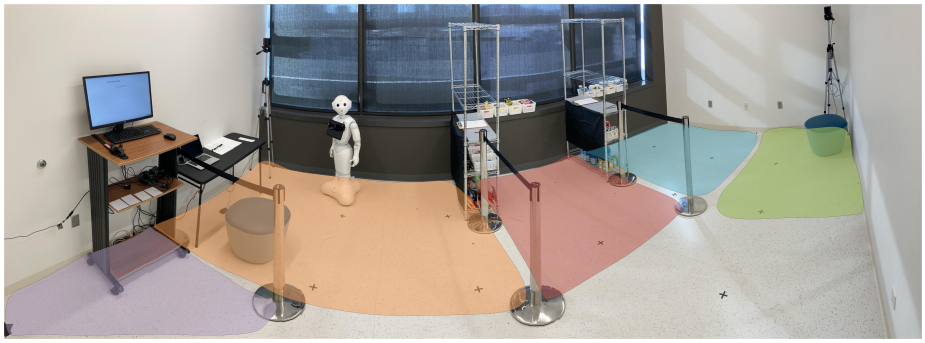

Figure 15: The simulated office environment used in the summative evaluation studies. The colored regions depict different "rooms" defined in this office.

6.1 Simulated Office Environment

The office space consisted of five regions, namely *Work Area*, *Rest Area*, *Snack Pantry*, *Drink Pantry*, and *Front Desk*. Participants were assigned a variety of tasks across each of these regions. Figure 16 shows the full layout of the office and its regions.

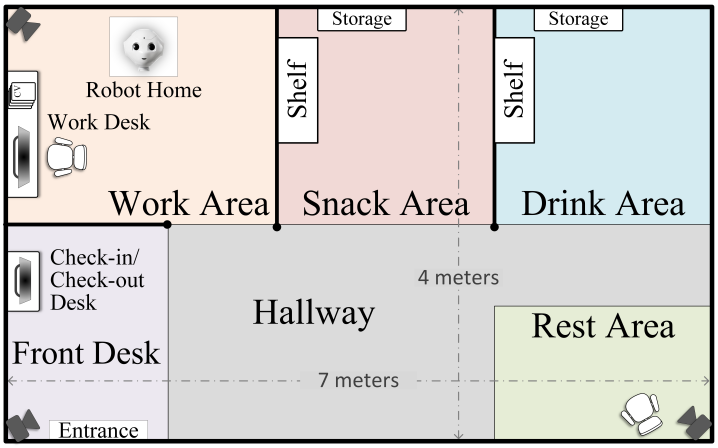

Figure 16: The layout of the simulated office environment.

The *Work Area* included of a desk, a chair, a laptop computer, and ten bundles of five curricula vitae (CVs). In this area, participants completed two work-related tasks: For *Work 1* task, participants selected a candidate from one of the CV bundles that matched a provided job description. As the second work task, *Work 2*, they wrote a short sentence on the computer explaining their reasoning for selecting the candidate. Each of these two tasks had to be completed once each day in the correct order, but they were not required to be completed consecutively.

The *Rest Area* was decorated with a chair. In this area, the participants' task was to *Rest* by taking a seat, starting a one-minute timer on their phone as a reminder for the end of their "break," and spending this time however they would like to.

The *Snack Pantry* contained a shelf with various healthy and unhealthy snacks (e.g., dried berries, candy bars), empty baskets, and a log book. In this region, as the *Snack* task, the participants were asked to "take a look at the snacks and imagine" what they would get if they were currently in the office looking for a snack, grab the snack and place it in an empty basket, and log the snack they just got on the logbook.

The *Drink Pantry* had the identical structure to the *Snack Pantry*, except that it contained healthy and unhealthy drink items (e.g., sparkling mineral water, soda). The *Drink* task was also identical to the *Snack* task, except that participants chose a drink instead.

The *Front Desk* featured a computer on a standing desk, and this region was where each day began and ended. At the end of each day, the participants reported their experience with the robot for the corresponding day through a survey (Section 6.6) displayed on the computer screen as their *Check-out* task. Once the *Check-out* task was completed, the participants then clicked through the survey to be prompted to start their next day and picked any one of the schedule cards they created earlier to follow during their upcoming day. This procedure of starting a new simulated day was coded as the *Check-in* task.

While creating their schedules, participants were instructed that each day should begin with a *Check-in* and end with a *Check-out*. They were also instructed that each day must contain single instances of *Work 1* and *Work 2*, and they could schedule the remaining three tasks–*Rest*, *Snack*, and *Drink*–as many times as they pleased. Initially, they were asked to imagine what a typical work day would look like for them to fill out their first schedule. Afterwards, they were asked if they could imagine any variations of this routine based on their experience, and asked to fill as many additional schedule cards as they wished.

6.2 Data Acquisition and Processing

During the experiments, participants and the robot wore belts that carried a motion tracking marker (HTC VIVE Tracker (3.0)). The system acquired 6D pose information from both of these trackers at a 100 Hz sampling rate during the session, and logged them locally along with a timestamp for each acquired sample. The 3D pose information (x-y coordinates and orientation) for the robot's tracker was used to drive the robot around the office space between different regions of interest autonomously. The 2D position information (x-y coordinates) from the participant's tracker and the associated timestamps were used to train the GPR model, as well as make real-time trajectory predictions at 10 Hz during the sessions.

Following the scheme demonstrated in Figure 8, each observed participant trajectory was snipped into "daily" trajectory segments at the transition from a *Check-out* task to a *Check-in* task at the end of each day. Figure 17 shows an example of 10 days of trajectory data collected in the simulated office environment, along with the defined activity regions within the laboratory space.

6.3 Task Assignment and Activity Recognition

For the user study sessions, the robot was assigned with the task of delivering a healthy option recommendation during *Snack* and *Drink* activities (Figure 8-"Task Assignment"). The instantaneous nature of such health-related decision making activities makes them highly relevant to the goals of an opportunistic health intervention system: To intervene at the time and place a health-related decision is being made. Such interventions–referred to as "stimulus control"–are shown to be the one of the most effective ways to facilitate health behavior change, [99]. To simulate a hypothetical activity recognition tool to suit the

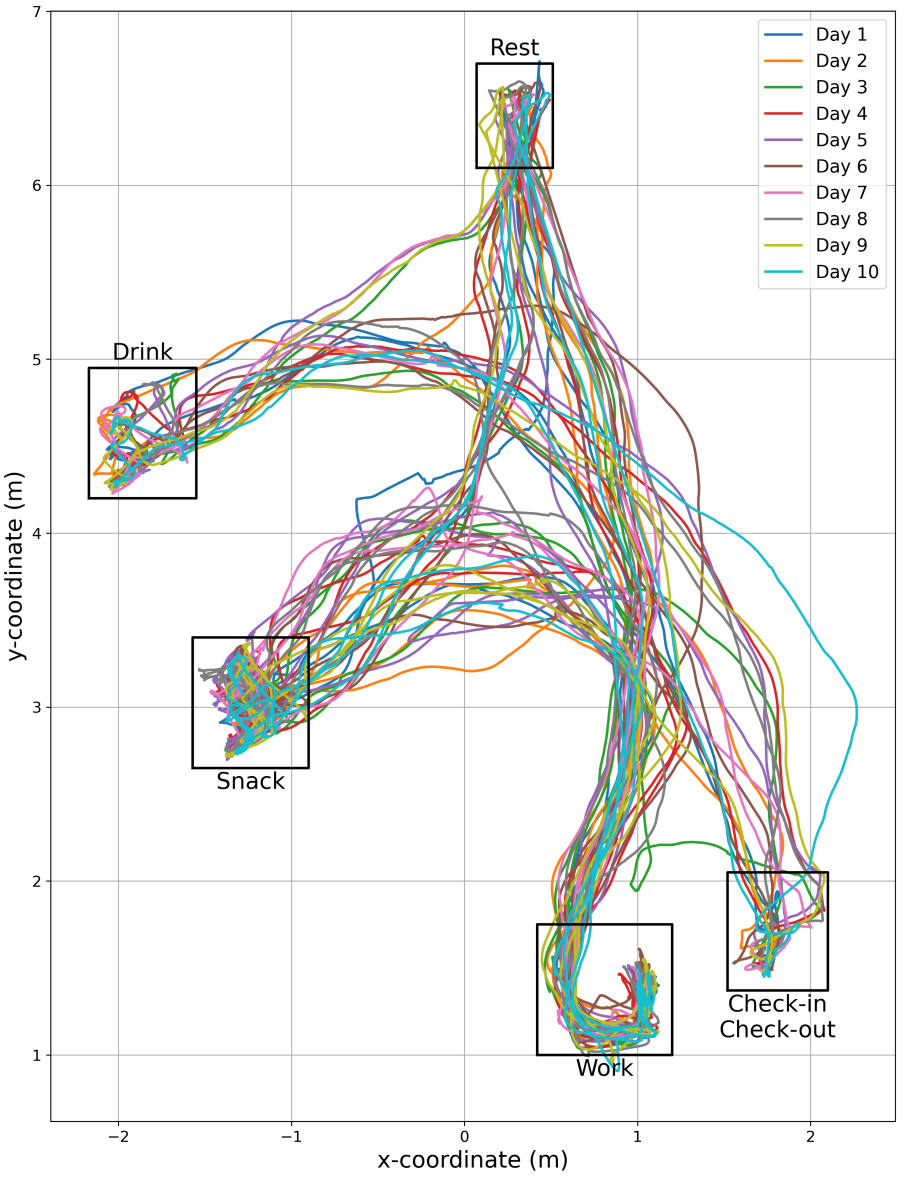

Figure 17: An example 10-day trajectory data in the simulated office study. These trajectories span approximately 58 minutes of total duration. The boxes represent the activity regions defined in the office. Note that this is only a 2D projection of the otherwise 3D observation space, whose third axis is the time.

purpose and the scope of our studies, and follow the scheme demonstrated in Figure 8, we leveraged time and coordinate-based bounding boxes around the activity regions. These spatiotemporal bounding boxes were used to label the intersecting sections of the current or predicted time series position trajectories with the associated activities. If an observed or predicted trajectory intersected with one or more of these bounding boxes for at least 3 seconds, than those activities were referred to as being "recognized" in the corresponding trajectory.

6.4 Model Training and Trajectory Prediction

For each participant, a new GPR model was trained from scratch during their experiment session. Depending on the current size of the data tensors–which varied from day to day as well as from user to user–the training cycles took around 5 to 20 seconds and were run on an NVIDIA GeForce RTX 4060 Ti 8 GB GPU.

Although re-training the model after each simulated day would have yielded better representations of the routine trajectories, we refrained from doing so in order to save crucial seconds spent during training–possibly minutes in total–and maximize the utility from the duration of the study sessions. For this purpose, the GPR model was only re-trained if the participant exhibited a previously unseen daily routine. These routines were detected by a helper algorithm that checked the bounding boxes for all 7 activities throughout the current day, and computed the current routine as a sequence of activities. If the sequence of activities during any given day had not been previously demonstrated by the participant, than that day's trajectory was tagged as a training trajectory. At the end of such a day, the system featurized this trajectory following Algorithm 3, added these new features and labels to the model tensors, and re-trained the model.

The trajectory featurization process inherited a downsampling step in order to keep the model tensor size from growing beyond computational limits (Algorithm 2). This downsampling was performed using equidistant regimes in all feature axes, t, x, y, h. Every consecutive point with 50 centimeters or 15 seconds separation in between–whichever came first–were selected as key feature point, (t_i, x_i, y_i), from the original trajectory. Then, following the trajectory from each of these key feature points over a 2-minute horizon and sampling more points to populate the data points across the horizon axis using the same equidistant regime, the 4D feature matrix and the 2D label matrix was generated as demonstrated in Equations 14,19, and 20. Although the proposed trajectory prediction method does not impose a strict upper limit on the prediction horizon (as long as the number of sampled data points remained tractable), a 2-minute prediction horizon was found to be sufficiently long during pilot experiments for our study's purposes.

In order to make trajectory predictions, new feature matrices for the current day's trajectory were generated and added to the model's tensors as exemplified in Equation 16 and Algorithm 4 periodically. To control the tensor size growth, only the last 20 seconds of trajectory history was retained in the model tensors.

The model tensor was updated at 1 Hz and the trajectory predictions were generated at approximately 10 Hz. Intersections of the predicted trajectories with the spatiotemporal bounding boxes of the activities were then computed, and the "persistent" intersections were output as predicted activities by the system. This output consisted of the activity's name, when it will take place relative to the current time, and the duration it will last for. The "persistence" in this context was defined as the same activity being recognized within the same time window several times consecutively ($N = 30$ used for the experiments). Example 400-second-long trajectory predictions and corresponding recognized activity time windows are demonstrated Figure 18.

6.5 Robot Behavior

During the experiments, we used the mobile humanoid robot platform Pepper (Figure 1; [107]). The two levels of the within-subjects robot behavior manipulation were *Reactive* and *Proactive*. The *Reactive* robot acted only after detecting that the participant was currently engaged in one of the target activities–either *Snack* or *Drink*–for at least the past 3 seconds. The *Proactive* robot acted only upon predicting, at least 40 seconds in advance, that the participant will be engaged in a target activity for at least 3 seconds.

In either condition, once the robot was triggered to act and attempt to deliver a *Snack* or *Drink* intervention, it autonomously navigated to the associated activity region. If the robot was able to join with the participant at the target task location before the participant picked up a snack or drink item from the shelves, it was then triggered to deliver the intervention. At this time, the system selected a random healthy option from the inventory, and the robot uttered a pre-scripted propositional statement recommending the participant to consider getting the proposed item, accompanied by a supporting health benefit argument (e.g., "*What about a banana? Bananas provide potassium to help your kidneys and heart function properly.*"). If the robot failed to meet with the participant at the target task location before the participant picked up an item, then the system sent it either back to the home position or to the next predicted target task location without uttering any statements. During idle periods in both conditions, the robot remained stationary in its home position next to the work desk (Figure 16) while running idle animations, and it did not interact with the participant.

During 5 of the 10 simulated days in the office, the robot's behavior was set to be *Reactive*; and during the other 5, it was set to be *Proactive*. The first day for each participant was considered an acclimation day where they got used to the tasks and following the schedule cards. This day's trajectory data was not used for training the GPR trajectory prediction model, and the robot was set to *Reactive* condition. The second day's data was used as the first training data for the *Proactive* robot, and the robot's behavior was again set to *Reactive* since the system was not able to make trajectory predictions yet. The system was able to make trajectory predictions for the remaining 8 days. These days were randomly assigned to either *Proactive* or *Reactive* conditions, making sure

that the total day count for each condition summed up to 5.

During the experiments, the experimenter observed the session from a different room. Before starting a navigation task, the system required confirmation from the experimenter as an added safety measure. Once a navigation task to a target activity region was completed, the system asked the experimenter for confirmation of whether to deliver the intervention utterance or not. The experimenter either confirmed the request if the participant was in the activity region and had not picked an item yet, and rejected it otherwise. During the *Proactive* days, the experimenter postponed responding to an utterance request for as long as the current target task was still being predicted to happen by the system or until the participant showed up at the intervention region. If the task was no longer being predicted and the participant still did not show up, then the experimenter rejected the utterance request. If the participant showed up in the task region and joined with the robot while the activity was still in the prediction scope, then the experimenter confirmed the utterance request, and the intervention was delivered. This intervention delivery trigger was originally designed to be handled autonomously by using face detection capabilities of the Pepper robot. However, during the pilot studies, it was decided that this capability was not reliable enough to guarantee high repeatability across intervention attempts, and the aforementioned manual triggering approach was adopted to prevent confounding effects.

6.6 Measures

The primary outcome for the user studies was the Intervention Success, which was measured at two levels: True or False for each intervention attempt. During an intervention attempt, if the robot was able to meet the participant to deliver the recommendation utterance before the participant already picked up a snack or drink item, then this intervention was considered successful. Additionally, at the end of each day during the *Check-out* task, the participants responded to three Likert scale survey questions to rate their perception of the robot in the past day. These questions were single-item scales selected from the three subscales of the Robot Social Attributes Scale [21], namely Warmth, Competence, and Discomfort. The selected scales were Social, Competent, and Awkward. We preferred to use only a subset of the full scale in order to mitigate the survey fatigue since each participant completed this survey ten times during a single session. Moreover, we scored the snack and drink items with -1 for each unhealthy item and +1 for each healthy item, and computed snack, drink, and total basket scores for each day. We also counted the total number of healthy and unhealthy items in the baskets for each day. During the exit interviews, the participants were asked for their preference between the *Reactive* and the *Predictive* robot. The interviews were transcribed and coded for analysis.

6.7 Offline Validation

After the experiments were completed, the logged trajectory data from the sessions were used for prediction success evaluations. For these evaluations, the system's task assignment was to predict every occurrence of all 7 activities in each day. Among these 7, *Work 1*, *Work 2*, *Check-in*, and *Check-out* occurred at most once each day, whereas *Drink*, *Snack*, and *Rest* could occur multiple times. Spatiotemporal bounding boxes for each of these activities were made available to the system as "recognition" parameters.

Using the routine detection helper algorithm (Section 6.3), ground truth activity sequences for each day were computed. The days where each novel routine sequence occurred for the first time for each participant were selected as training data segments for the corresponding participant. If a detected routine occurred only once throughout the 10 days for a participant, then that data segment was discarded and was neither used for training nor testing. The remaining data segments were used as testing dataset for the corresponding participant.

A GPR model was trained from scratch for each participant using the trajectories in the training set. Then, for each day segment in the testing set, the model made trajectory predictions over a 2-minute horizon every 10 milliseconds throughout the day. Each trajectory prediction was checked for intersections with the spatiotemporal bounding boxes of every activity, and the recognized activities were inserted to a data structure containing the sequence of predicted activities sorted by their predicted onset time within the day. The identical consecutive predicted activities in this data structure were later aggregated into a single prediction if their predicted time windows overlapped with one another. If there was no overlap in the predicted time windows, then these predicted activities were treated as different instances of the same activity being predicted, and no aggregation was performed.

The resulting sequence of activities at the end of a day was coded into "sentences" composed of single-letter "words". For example, if the predicted activity sequence in a day was "*Check-in*→ *Drink*→ *Work 1*→ *Snack*→ *Work 2*→ *Check-out*", this was coded into "C D W S W C". This resulting predicted sequence sentence was then compared to the ground truth sequence for the corresponding day using the ROUGE-L metric [82]. Additionally, for each predicted sequence, we generated 1000 random sequences of equal length and scored them using the same metric. The averages of these random scores were compared to the predicted sequence scores using paired samples t-test.

The proposed GPR model's trajectory predictions are expected to become more refined as the day progresses–especially for the cases such as the one exemplified in Figure 18-b,c. In order to evaluate the effects of the current trajectory history on the activity sequence prediction success, several aggregate sequence predictions were collected beginning at different time points throughout each test day. These time points, $T0, T1, T2, ..., TN$, were defined as the timings where the current trajectory transitioned into a new activity region at each new test day. To clarify, the $T0$ prediction sequence is collected starting right at the beginning of a day when the current position is within the spatial boundaries of

the first activity region. Once the current position moved into the second activity region, a new sorted activity sequence data structure, $T1$, was created and populated with predictions through the rest of the day alongside $T0$ predictions. In these settings, for the example sequence given above where the ground truth for $T0$ is "C D W S W C," the ground truth sequences for $T1$ and $T2$ are "D W S W C" and "W S W C," respectively. All of these sub-sequence predictions were evaluated and compared to random sequences of equal length separately.

6.8 Study Procedure

We recruited participants from a social media platform internal to our institution as well as through word of mouth within the surrounding laboratories. The participants were required to be at least 18 years old and proficient in spoken and written English. The study was approved by our institution's IRB, and the participants were compensated for their time. Experiment sessions took approximately 90 minutes and were video- and audio-recorded.

Upon arriving at our laboratory, the participants were briefly introduced to the study and consented. Upon consent, the participants completed a sociodemographic survey and a robot attitude survey [21]. The participants were then led to the simulated office room and were given a detailed briefing about the experimental tasks. During the briefing, the robot was turned on, running idle animations in its home position, and it did not interact with the participant. Once the briefing was complete, the participants filled out as many routine cards as they wished. After the participant finished filling out their routine cards, the experimenter prompted them to start with their first day and left the room. Once the participants were done with their 10 days of activity simulation, they left the room to meet the experimenter and then completed a post-experiment survey. Finally, the participants were interviewed about their experience and debriefed.

7 Results

We recruited 15 participants for this study (8 male, 7 female) between ages 20-37 ($Median = 34, Mean = 30.7, SD = 5.1$). 7 participants had no prior experience with any kind of robot. 3 participants have previously interacted with a social robot and 4 participants previously worked with an industrial robot manipulator. The survey data from one session was excluded from the analysis because the study log containing the mapping between the day number and trial condition was lost during file export due to a runtime error. Trajectory data acquisition from all 15 sessions was completed successfully.

In total, these sessions yielded 150 simulated days and 1191 activities within the simulated office. Among these 1191 activities, 410 were *Snack* and *Drink* tasks, which were the intervention opportunities for the robotic system. The *Reactive* robot had 211 total opportunities, and the *Proactive* robot had 199

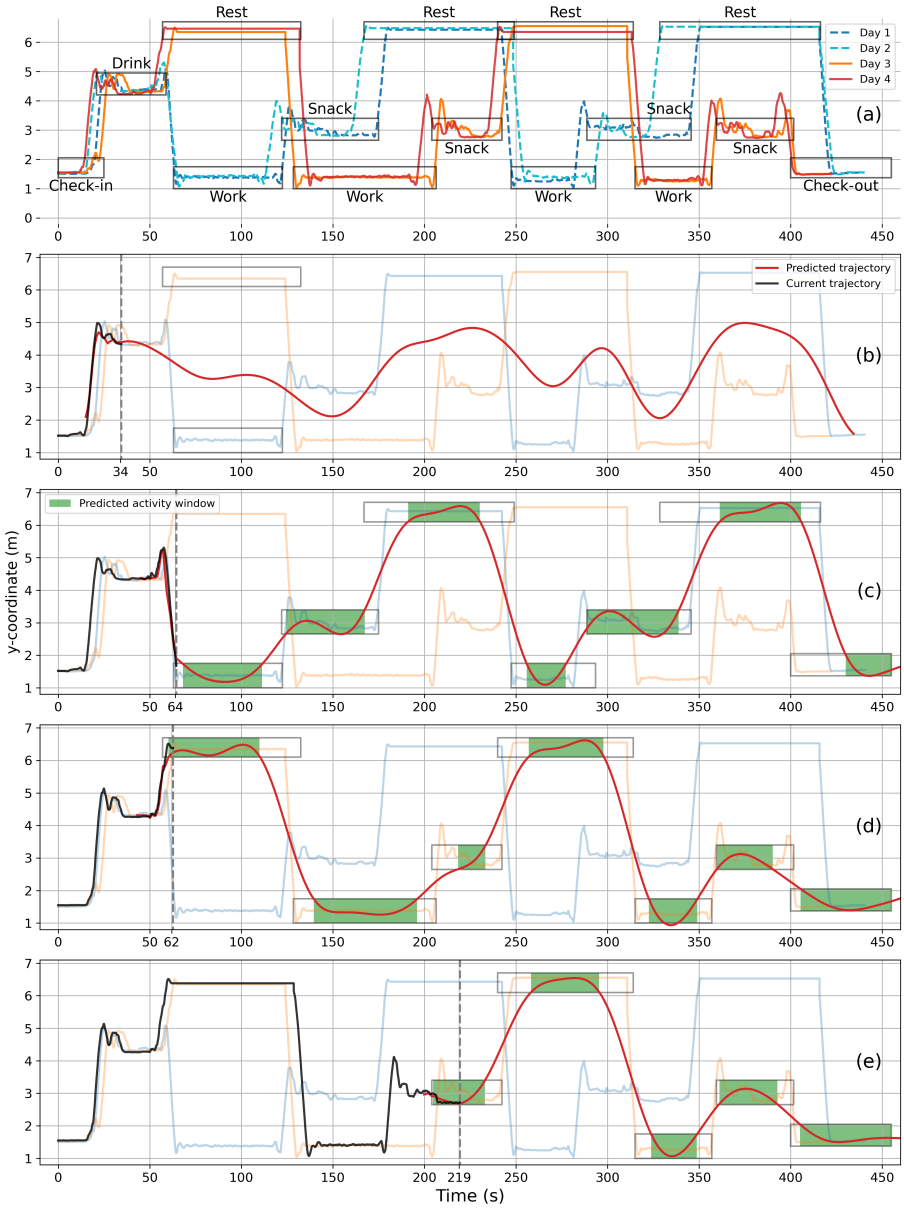

Figure 18: (a)Demonstrations of observed minor and major routine variations in the collected data and spatiotemporal activity windows. (b-e)400-second-long predicted trajectories during previously unseen days. (c-e)Predicted activity windows. Note that these plots demonstrate only a single 2D snapshot from the underlying 6D inference problem for observational simplicity. In (b), the model is yet uncertain whether the user will advance towards *Rest* or *Work* activities and follow the associated routine. Once the trajectory takes a turn indicating either routine, then the remainder of the trajectory as well as the routine sequence can be successfully predicted (c, d).

total opportunities during the sessions. 75 of the 150 simulated days was with the *Reactive* robot and the other 75 were with the *Proactive* robot.

For the offline validations, 150 days of trajectories were analyzed in individual clusters per participant, and the activity sequences were extracted. The activity sequences in 15 days were observed only once during the corresponding participant's session, and were discarded from offline validation dataset. The remaining data segments yielded 34 days of training data in total. Each participant's GPR model was trained using 2 to 3 days of data from their own sessions. The total number of test segments for all participants summed up to 101 days.

In Figure 19, a breakdown of the number of routines *defined* by each participant and the number of routines *executed* by each participant are given. Although all participants defined 2 to 3 routines, during the sessions, 2 to 5 different routines were executed by the participants. This was primarily due to participants getting confused while following the routine cards and either failing to remember which was the last activity they did, confusing the order of *Drink* and *Snack* they had on their cards, or confusing the locations of these tasks within the lab space. Each participant completed 6 to 10 activities per day, Figure 20, and each day took about 7 minutes on average ($M = 392s, SD = 103s, min = 190s, max = 668s$), Figure 21.

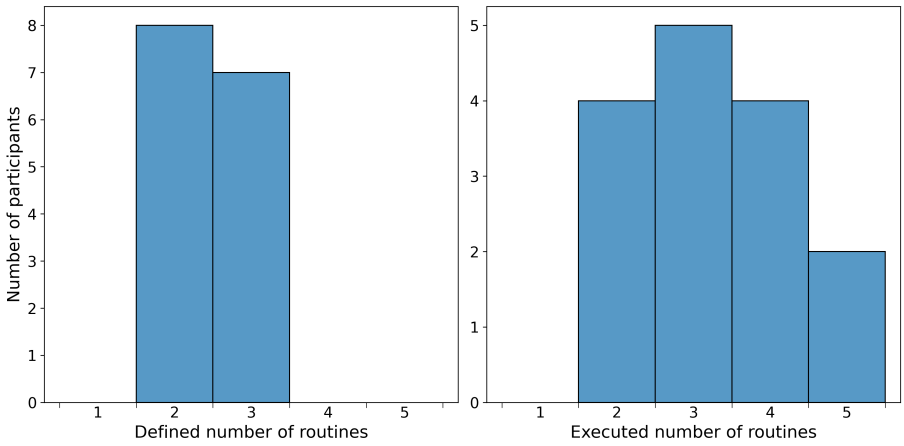

Figure 19: Participants defined 2 to 3 routines and executed 2 to 5 different routines during the sessions. Even though they followed the cards with their day's routines written on them, the participants sometimes forgot which step they were at or confused the order of *Snack* and *Drink* tasks.

7.1 Primary Outcome

The primary outcome from this study was the number of successful interventions based on the robot behavior condition. Each day, both *Reactive* and *Proactive*

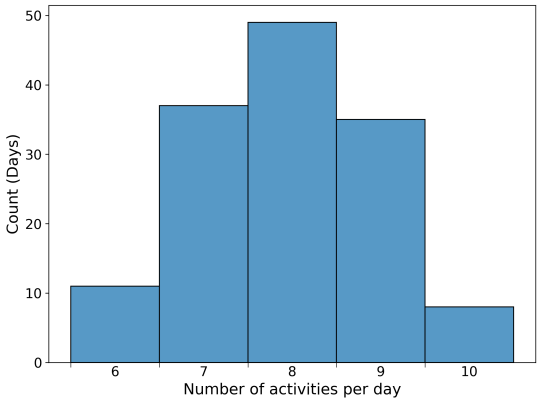

Figure 20: Participants engaged in 6 to 10 activities per day.

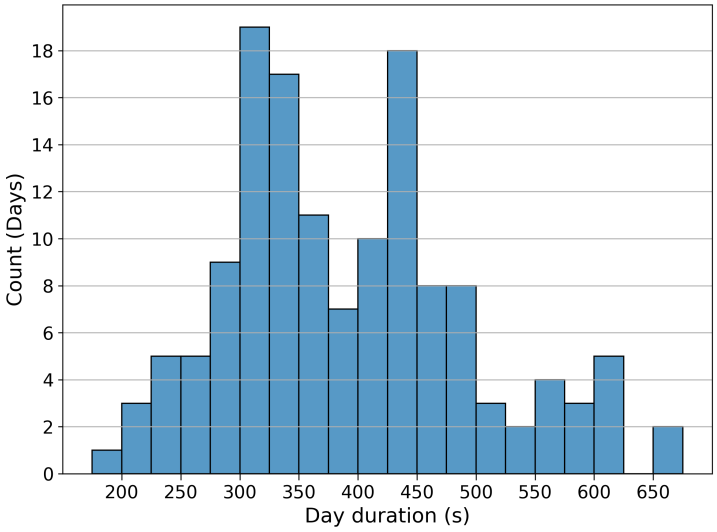

Figure 21: A simulated day took 392 seconds in average. The shortest day took 190 seconds and the longest day took 668 seconds. The durations of the days varied from day-to-day, routine-to-routine, as well as participant-to-participant.

robots attempted to address one or more intervention opportunities. Figure 22 shows a breakdown of these attempts by the robot condition.

The *Reactive* robot had a total of 211 intervention opportunities. It successfully delivered an intervention in 8 of these opportunities. In 141 of them, the *Reactive* robot failed to make it to the intervention location in time. The remaining 62 opportunities were "skipped" by this robot. These were the opportunities when the reactive robot was busy navigating to another point of

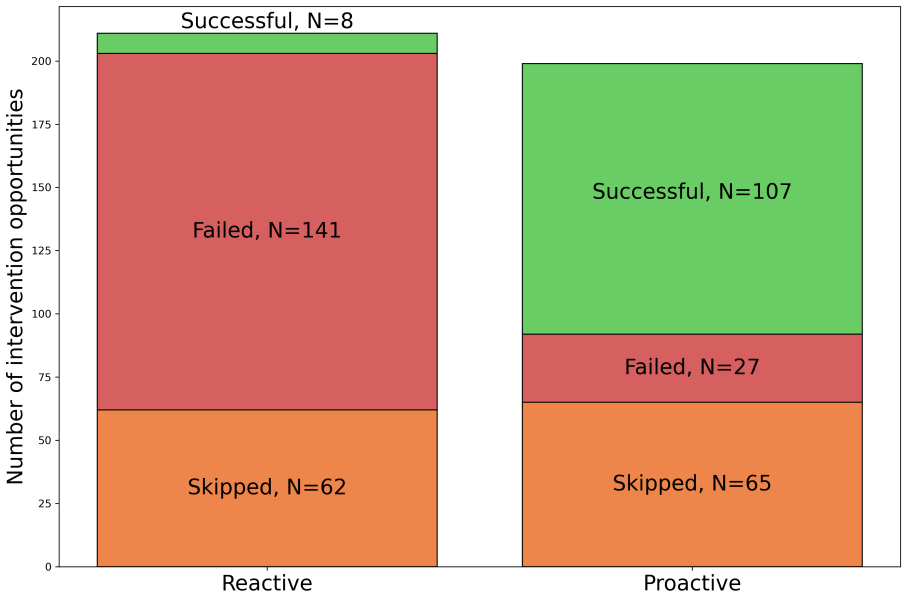

Figure 22: The breakdown of intervention opportunities per robot condition. The *Reactive* robot failed on the majority of its intervention attempts, and the *Proactive* robot was successful in the majority its attempts.

interest in the office space, which was either the *Snack* or *Drink* zone, or its home position to where both robots returned if there currently were no open intervention opportunities.

The *Proactive* robot had a total of 199 intervention opportunities. It successfully delivered a total of 107 *Snack* and *Drink* interventions. 27 of the 199 opportunities were failed to be addressed by the *Proactive robot*. These failures predominantly occurred during the days where the participant was demonstrating a new routine for the first time. During 2 trials, the robot's collision avoidance algorithm blocked the robot erroneously, which resulted in the *Proactive* robot to arrive at the zone late, and the intervention attempt to fail. 65 trials were skipped by the *Proactive* robot due to multiple reasons: (1) The robot was in transit to another point of interest in the office when the intervention opportunity was predicted and it was too late to act by the time the current navigation was completed, (2) The system failed to predict the opportunity, (3) The opportunity was predicted to happen in less than 40 seconds which causes the system to not act on the opportunity (Section 6.5).

The number of successful interventions per day ($Median = 1, Mean = 0.82 (SD = 0.93), min = 0, max = 3$) condition was analyzed using Mann-Whitney U test based on robot behavior condition. The analysis showed that the *Proactive* robot delivered significantly more successful interventions com-

pared to the *Reactive* robot ($U = 292.5, n_1 = n_2 = 70, p < .0001$, two-tailed). Overall, the *Proactive* robot was able to deliver a successful intervention in 79.85% of the attempted opportunities, and the *Reactive* robot's success rate was 5.37% for its attempts.

7.2 Secondary Outcomes

7.2.1 Perceived Social Attributes

The first group of secondary outcomes included single-item questions from the sub-scales of the Robotic Social Attributes Scale [21]. These items were *Social*, *Capable*, and *Awkward*. In order to investigate possible novelty effects, these scales were analyzed in conjunction with the day number. Additionally, to investigate possible carryover effects, these scales were also analyzed based on whether the response was collected before or after the first successful robot intervention. The responses to these scales were first transformed using Aligned Rank Transform (ART) [140] to make the data compatible with the intended analysis method. Then, a factorial analysis of variance (ANOVA) was conducted on these two sets of transformed data.

These analyses have shown that the *Proactive* robot was perceived as significantly more capable ($F(1, 139) = 28.159, p < .0001$), significantly more social ($F(1, 139) = 22.996, p < .0001$), and marginally less awkward ($F(1, 139) = 3.393, p = .068$) than the *Reactive* robot. The robot was perceived as significantly more capable as the days went by ($F(1, 139) = 2.194, p = 0.027$). There were no significant effects of the day number on the *Social* ($F(1, 139) = 1.677, p = .102$) nor on the *Awkward* ($F(1, 139) = 0.874, p = .550$) scores. The analysis did not show any significant interactions between the robot condition and the day number. The raw–non-transformed–measurements from the social attributes scale by robot condition and day number are demonstrated in the regression plots in Figure 23.

The analyses on the second set of transformed data have shown that the robot was perceived as significantly more capable ($F(1, 139) = 7.168, p = .008$) and significantly more social ($F(1, 139) = 5.587, p = .020$) after the first successful intervention. The timing of the measurement did not have any significant effect on the perceived awkwardness of the robot ($F(1, 139) = 1.604, p = .207$). No significant interactions were found between the robot condition and the measurement timing on the social attribute measures.

7.2.2 Snack and Drink Quality

The robot condition's effect as well as the timing of the measurement–whether it was before or after the first intervention–on meal item quality were analyzed using Mann-Whitney U test for scale variables; and chi-squared test for categorical variables. The analyses have shown that participants included significantly fewer unhealthy items in their daily baskets under the *Proactive* robot condition ($U = 1922.5, n_1 = n_2 = 70, p = .016$, two-tailed). There were no significant

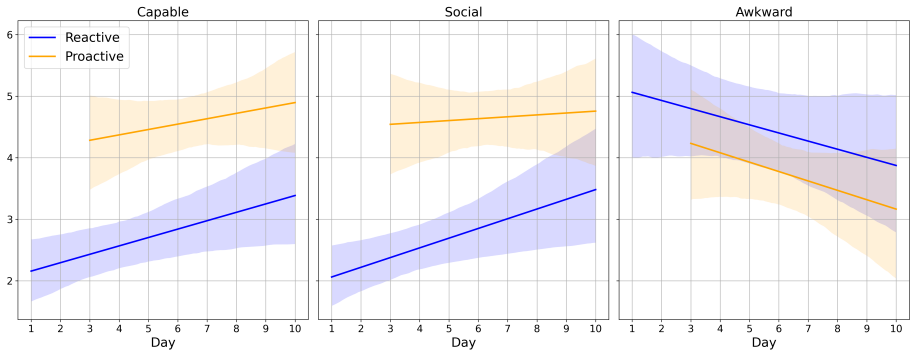

Figure 23: Social attributes scores by day and the robot condition.

effects of the robot condition on number of healthy items ($U = 2235.0, n_1 = n_2 = 70, p = .353$, two-tailed), daily basket score ($U = 2116.5, n_1 = n_2 = 70, p = .157$, two-tailed), snack ($U = 2149.5, n_1 = n_2 = 70, p = .197$, two-tailed), nor drink ($U = 2323.0, n_1 = n_2 = 70, p = .583$, two-tailed) scores. The chi-squared tests revealed a marginal correlation between the robot condition and the basket containing at least one healthy food item, in favor of the *Proactive* robot ($\chi^2(1) = 3.739, p = .053$). No such correlation was found between the robot condition and the basket containing at least one healthy drink item ($\chi^2(1) = 1.346, p = .246$).

The meal quality metrics were primarily affected by whether the first successful intervention was delivered or not. Participants included significantly more healthy items ($U = 1419.0, n_1 = 35, n_2 = 105, p = .037$, two-tailed), and significantly fewer unhealthy items ($U = 1296.5, n_1 = 35, n_2 = 105, p = .004$, two-tailed) in their baskets after the first successful robot intervention. Similarly, the basket scores ($U = 1355.0, n_1 = 35, n_2 = 105, p = .018$, two-tailed), snack scores ($U = 1422.0, n_1 = 35, n_2 = 105, p = .039$, two-tailed), and drink scores ($U = 1500.5, n_1 = 35, n_2 = 105, p = .092$, two-tailed) were higher after the first successful intervention. The participants were more likely to include at least one healthy drink item ($\chi^2(1) = 3.191, p = .074$) and a snack item ($\chi^2(1) = 7.179, p = .007$) in their baskets in the days following the first successful intervention.

7.2.3 Compliance versus Snack and Drink Quality

In order to confirm whether complying with the robot's recommendations actually improved the snack and drink quality scores, an additional analysis was performed. The "compliance" was coded as True in the trials where the participant included the recommended item in their corresponding baskets, and False otherwise. The analyses have shown that complying with a recommendation significantly increased the number of healthy items ($U = 102.5, n_1 = 14, n_2 = 58, p <$

.0001, two-tailed) and significantly decreased the number of unhealthy items ($U = 242.0, n_1 = 14, n_2 = 58, p = .008$, two-tailed) in the baskets. Additionally, overall basket scores ($U = 121.0, n_1 = 14, n_2 = 58, p < .0001$, two-tailed) as well as the snack scores ($U = 129.0, n_1 = 14, n_2 = 58, p < .0001$, two-tailed) were significantly improved by complying with the robot's recommendations. No such improvement was found for drink scores ($U = 314.0, n_1 = 14, n_2 = 58, p = .172$, two-tailed). The participants were significantly more likely to include at least one healthy snack item ($\chi^2(1) = 23.533, p < .0001$) and marginally more likely to include at least one healthy drink item ($\chi^2(1) = 2.937, p = .087$) if they complied with the robot's recommendation.

7.2.4 Exit Interviews

During the exit interviews, the participants were asked for their preference for either the *Reactive* or the *Proactive* robot. 14 out of 14 participants stated that they would prefer to have the *Proactive* robot in their office space instead of the *Reactive* robot. Participants explained this preference as being due to the *Reactive* robot being "too late" and the participants having "already" decided what to pick before the robot arrived, or the participants having to "wait" for the *Reactive* robot to arrive at the intervention location. Two participants stated a nuanced preference over the different robot behaviors depending on whether the robot was a "colleague", in which case they had a stronger preference for the *Proactive* robot; or a "pet", in which case they had a preference for the *Reactive* robot because the robot, behaving this way and following them around, was perceived to be "cute".

7.3 Offline Validation Results

The trajectory data collected from all 15 participants was used in the offline validation studies. These trajectories contained 1191 activities in total, all of which were assigned as a target task to the system and were attempted to be predicted by the system. The GPR model was trained from scratch for each participant using the first occurrence of each unique routine as training data, and the rest of the days as testing data for the associated participant. The same training and prediction parameters used during the sessions (Section 6.4) were used during offline validations. The only difference was that the data tensors were updated with each trajectory sample at around 100 Hz instead of 1 Hz, since computation time wasn't critical during offline predictions. With every new sample, a new trajectory prediction was made, yielding approximately 100 predictions per second instead of ∼10 predictions per second–which was the maximum rate that the system could keep up with during the experiment sessions. This yielded approximately 5.5 million trajectory predictions in total for all 15 participants.

Figure 24 shows the Rouge-L scores for each day in the test dataset overlaid with a box plot to outline the distribution. A score of 1.0 indicates the ground truth sequence was predicted perfectly. The mean scores of the predicted ac-

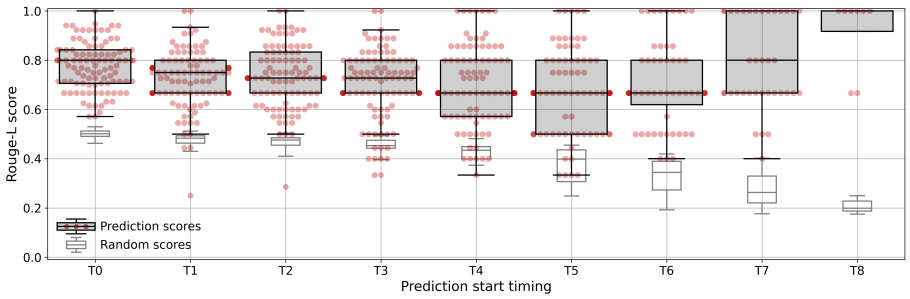

Figure 24: Rouge-L scores for the predicted activity sequences during offline validation studies.

tivity sequences through all 9 timings (T0,T1,...) were 0.75 and the standard deviation was 0.15. The sequence prediction score of the model was significantly higher than random for all timings (paired samples t-test $p < .00001$, two-tailed). The mean random scores decreased consistently through the timings. This is because the length of the remaining ground truth sequence decreased by one at each consecutive timing, while the number of possible activities to randomly populate this sequence remained constant.

Mean prediction score remained relatively stable compared to random throughout the timings. The spread of the prediction scores increased with the timing. The increase in the spread can be attributed to two observations: If the participant followed their routine with relatively small time variations compared to a previously observed identical routine that was used for training, then the model would refine and improve its predictions as it received more context about the current routine as the day progressed. This resulted in more high-score predictions in later timings. However, if this was not the case and a repetition of a previously observed training routine had significant differences, then the model would get more and more lost as the day progressed.

8 Discussion

The presented approach was shown to be able to successfully make 2-minute-long trajectory predictions, which were then used to drive the robot around a simulated office through simulated days which lasted several minutes. However, it should be straightforward to extend these predictions over longer horizons in longer routines that spanned hours–maybe days. As long as the scale of relative variations between the output and the input signals remained comparable, the exact model presented in this work can be used to make predictions in domains with completely different time-scales. Moreover, the 2-minute horizon used during the studies can be increased to be equal to the total duration of the routines. The only reason that the prediction horizon was kept relatively short

during the sessions was to keep the data tensors as small as possible, and not risk running into memory issues during the studies.

An interesting finding from the summative evaluation study was the lack of an impact of the robot condition on several meal quality metrics. One possible explanation for this might be the simplicity of +1/-1 scoring scheme (Section 6.6), which was used to evaluate the meal quality. Combined with the relatively small number of items picked from the shelves during each *Snack/Drink* activity, this scoring scheme might have failed to achieve the necessary scoring granularity at the trial level. Another possible explanation is that the carryover effects were indeed strongly influencing the participants' decisions following the first successful intervention by the robot. Regarding this, two participants explicitly stated during the exit interviews that they felt the robot was asking them to pick healthy items even when it was "staring" at them from across the room without saying anything–which clearly did not happen during a successful intervention trial. The perception of such mutual gaze due to the positioning of the robot in its home position and there being no obstacles in the office to break this gaze connection with the robot might have further strengthened the carryover effects.

9 Limitations

The primary limitation of the trajectory prediction-based intervention system presented in this thesis is its scalability. This system was implemented using an exact GPR model, which is known to have strict limitations on the amount of data that it can handle. If such a system were intended to be used in real-life longitudinal settings, then this model would likely not be suitable to scale with the possible varieties in the users' routines through the weeks or months the system was deployed for. The second limitation is with regard to the trajectory prediction generated by the system and the lack of physical constraints that are guiding these predictions. In the current state, the predicted trajectories might go through physically impossible terrain (e.g., walls)causing infeasible outcomes and possibly causing the loss of otherwise important context about the nature of the behaviors. Third, even though the presented system was evaluated in an environment with variations in the observed behaviors through the usage of routine cards, the true stability or variability of real-life routines were probably not reflected during the summative evaluation studies, and it still remains as future work to verify whether real-life routines have similar characteristics to the routines evaluated in this thesis.

The primary limitation regarding the experimental design of the summative evaluation studies was due to convenience sampling. Even though this participant pool provided a sufficiently large dataset of activity sequences for the offline validations, it was not large enough to reflect the behavioral characteristics and the preferences of the general public. Moreover, the short duration of the one-hour study inside a room with the robot was likely confounded with novelty effects, and it is unclear if the advantage of the proactivity over reactive robot

behavior would persist after the novelty effects have worn off in longitudinal settings.

10 Conclusion

The work described in this manuscript aimed to push towards robot proactivity with a use case of in-situ opportunistic health interventions using a mobile humanoid robot to deliver recommendations during a health-related decision-making task. In an attempt to generalize the desiderata for such systems beyond this initial case study, the overarching key aspects of such systems were identified as spatiotemporal causality, adaptability, proactivity, and generalizability. Through the proposed approach, it has been shown that temporal **causality** can be retained by the proposed feature space reconstruction method which makes it possible to infer possible future positions based on the current and the past observations (Sect. 5.1). By leveraging assumptions regarding the underlying dynamics of the spatiotemporal system, which routine behavior is modeled as through Gaussian processes, and making use of such model's predictive power, the proposed approach was useful with very small amounts of data–as small as observations only from a single routine trajectory–making it highly **adaptable** to different users and environments (Sect. 5.3). The ability to model and predict routine behavioral trajectories combined with the proposed system architecture as demonstrated in Figure 8 made it possible to render the robotic platform to predict and act upon future activities of study participants and achieved **proactivity** (Sect. 4.4). The capability of the proposed method and structure to **generalize** to other task domains that incorporated some other form of time series data was not discussed, and remains a part of future work.

The findings from the summative evaluation study (Section 6) have shown that the proposed proactive system was significantly more successful in delivering successful interventions than a reactive system, which did not have predictive capabilities. The proactive robot was found to be more social and capable, and perceived as less awkward by the participants, and participants showed a strong preference for the proactive robot. Even though the proactive behavior had significant effects only on a small subset of the meal quality metrics, the importance of delivering a successful intervention revealed itself as carryover effects (Section 7.2.2): The meal quality measurements were significantly higher after delivering one successful intervention with the robot. This signified the importance of being able to deliver successful interventions as such interactions led the participants to become more aware and adjust their preferences accordingly with this awareness.

11 Future Work

In the technical direction of the underlying approach, the future work should investigate how such a framework can be deployed longitudinally in life-long

learning settings, and be able to scale with the amount of data collected from daily observations. If such a system were to be deployed for a very long time period, using approximate GPR methods to handle the scalability of the representation learning system would likely not be sufficient. Additionally, the model would likely require additional algorithms to sample observations from new behaviors and "forget" previously learned but no longer relevant contextual connections between observations and behaviors.

A direction left for future work is demonstrating the generalizability of the proposed system in other domains that involved continuous-time series data. For example, the system can be further studied to model and predict biological signals, or be used to model human-object interactions and object manipulations.

A common property of any distribution is that it can be used to draw samples from. So, ideally, the learned representations in this work can be used to synthesize context-rich data based on real observations by sampling trajectories from the learned distribution. Such capability was out of the scope of this work. Nonetheless, it might be an interesting future direction for more relevant studies and applications.

Lastly, even though the connections between health behavior and spatiotemporal cues are shown by prior work in different contexts, it is necessary to conduct further studies specifically in human-robot health interactions to confirm whether these connections are still reliable to create proactive health interventions using robots in real-life settings. Moreover, longitudinal studies are necessary to confirm if such proactive intervention behaviors can retain their relevance to the users and effectiveness over long periods of time.

References

[1] Henny Admoni and Brian Scassellati. Social eye gaze in human-robot interaction: a review. *Journal of Human-Robot Interaction*, 6(1):25–63, 2017.

[2] Chris Baker, Rebecca Saxe, and Joshua Tenenbaum. Bayesian models of human action understanding. In *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2005.

[3] Nikola Banovic, Tofi Buzali, Fanny Chevalier, Jennifer Mankoff, and Anind K. Dey. Modeling and Understanding Human Routine Behavior. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 248–260, New York, NY, USA, May 2016. Association for Computing Machinery.

[4] Joseph E Bauer. *Health Behavior and Health Education: Theory, Research, and Practice (3rd edn.)*. Medknow, 2004.

[5] Geetanjali Bhola and Dinesh Kumar Vishwakarma. A review of vision-based indoor har: state-of-the-art, challenges, and future prospects. *Multimedia Tools and Applications*, 83(1):1965–2005, 2024.

[6] Timothy W Bickmore. *Relational agents: Effecting change through human-computer relationships*. PhD thesis, Massachusetts Institute of Technology, 2003.

[7] Andreea Bobu, Dexter R. R. Scobee, Jaime F. Fisac, S. Shankar Sastry, and Anca D. Dragan. LESS is More: Rethinking Probabilistic Models of Human Behavior. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '20, pages 429–437, New York, NY, USA, March 2020. Association for Computing Machinery.

[8] Indu P Bodala, Nikhil Churamani, and Hatice Gunes. Creating a robot coach for mindfulness and wellbeing: A longitudinal study. *arXiv preprint arXiv:2006.05289*, 2020.

[9] Edwin V Bonilla, Kian Chai, and Christopher Williams. Multi-task gaussian process prediction. *Advances in neural information processing systems*, 20, 2007.

[10] Oliver Brdiczka, Norman Makoto Su, and James Bo Begole. Temporal task footprinting: identifying routine tasks by their temporal patterns. In *Proceedings of the 15th international conference on Intelligent user interfaces*, IUI '10, pages 281–284, New York, NY, USA, February 2010. Association for Computing Machinery.

[11] Cynthia Breazeal. *Designing sociable robots*. MIT press, 2004.

[12] Cynthia Breazeal, Michael Steven Siegel, et al. *Persuasive robotics: how robots change our minds*. PhD thesis, Massachusetts Institute of Technology, 2009.

[13] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

[14] Rebecca Brown. Habitual health-related behaviour and responsibility. *Responsibility and Healthcare*, page 210, 2024.

[15] Dietmar Bruckner and Rosemarie Velik. Behavior Learning in Dwelling Environments With Hidden Markov Models. *IEEE Transactions on Industrial Electronics*, 57(11):3653–3660, November 2010. Conference Name: IEEE Transactions on Industrial Electronics.

[16] Tamara Bucher, Clare Collins, Megan E Rollo, Tracy A McCaffrey, Nienke De Vlieger, Daphne Van der Bend, Helen Truby, and Federico JA Perez-Cueto. Nudging consumers towards healthier choices: a systematic review of positional influences on food choice. *British Journal of Nutrition*, 115(12):2252–2263, 2016.

[17] Tamara Bucher, Klazine Van der Horst, and Michael Siegrist. The fake food buffet–a new method in nutrition behaviour research. *British Journal of Nutrition*, 107(10):1553–1560, 2012.

[18] Tamara Bucher, Klazine Van der Horst, and Michael Siegrist. Fruit for dessert. how people compose healthier meals. *Appetite*, 60:74–80, 2013.

[19] John T Cacioppo and Richard E Petty. The elaboration likelihood model of persuasion. *ACR North American Advances*, 1984.

[20] Sonja Caraian and Nathan Kirchner. Influence of robot-issued joint attention cues on gaze and preference. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 95–96, Tokyo, Japan, March 2013. IEEE.

[21] Colleen M Carpinella, Alisa B Wyman, Michael A Perez, and Steven J Stroessner. The robotic social attributes scale (rosas) development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on human-robot interaction*, pages 254–262, 2017.

[22] Kaixuan Chen, Dalin Zhang, Lina Yao, Bin Guo, Zhiwen Yu, and Yunhao Liu. Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. *ACM Computing Surveys (CSUR)*, 54(4):1–40, 2021.

[23] Vijay Chidambaram, Yueh-Hsuan Chiang, and Bilge Mutlu. Designing persuasive robots: how robots might persuade people using vocal and

nonverbal cues. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, HRI '12, pages 293–300, New York, NY, USA, March 2012. Association for Computing Machinery.

[24] Robert B Cialdini and Noah J Goldstein. Social influence: Compliance and conformity. *Annu. Rev. Psychol.*, 55:591–621, 2004.

[25] T Matthew Ciolek and Adam Kendon. Environment and the spatial arrangement of conversational encounters. *Sociological Inquiry*, 50, 1980.

[26] HTC Corp. Vive pro eye. www.vive.com/sea/product/vive-pro-eye/overview/, Last accessed April 2024.

[27] Angelo Costa, Stella Heras, Javier Palanca, Jaume Jordán, Paulo Novais, and Vicente Julián. Argumentation Schemes for Events Suggestion in an e-Health Platform. In Peter W. de Vries, Harri Oinas-Kukkonen, Liseth Siemons, Nienke Beerlage-de Jong, and Lisette van Gemert-Pijnen, editors, *Persuasive Technology: Development and Implementation of Personalized Technologies to Change Attitudes and Behaviors*, Lecture Notes in Computer Science, pages 17–30, Cham, 2017. Springer International Publishing.

[28] Nils Dahlbäck, Arne Jönsson, and Lars Ahrenberg. Wizard of oz studies - wizard. *Knowledge-based systems*, 6(4):258–266, 1993.

[29] Qi Deng and Dirk Söffker. Modeling and Prediction of Human Behaviors based on Driving Data using Multi-Layer HMMs. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 2014–2019, October 2019.

[30] Daniel C Dennett. *The intentional stance*. MIT press, 1989.

[31] Abubaker Elbayoudi, Ahmad Lotfi, Caroline Langensiepen, and Kofi Appiah. Modelling and simulation of activities of daily living representing an older adult's behaviour. In *Proceedings of the 8th ACM International Conference on PErvasive Technologies Related to Assistive Environments*, PETRA '15, pages 1–8, New York, NY, USA, July 2015. Association for Computing Machinery.

[32] David Ellis, Eric Sommerlade, and Ian Reid. Modelling pedestrian trajectory patterns with Gaussian processes. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 1229–1234, September 2009.

[33] Ghadeer Eresha, Markus Häring, Birgit Endrass, Elisabeth André, and Mohammad Obaid. Investigating the influence of culture on proxemic behaviors for humanoid robots. In *2013 IEEE RO-MAN*, pages 430–435, August 2013.

[34] Juan Fasola and Maja J Matarić. A socially assistive robot exercise coach for the elderly. *Journal of Human-Robot Interaction*, 2(2):3–32, 2013.

[35] David Feil-Seifer and Maja J Matarić. Socially assistive robotics. *IEEE Robotics & Automation Magazine*, 18(1):24–31, 2011.

[36] BJ Fogg. Persuasive computers: perspectives and research directions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '98, pages 225–232, USA, January 1998. ACM Press/Addison-Wesley Publishing Co.

[37] Benjamin Gardner, Phillippa Lally, and Jane Wardle. Making health habitual: the psychology of 'habit-formation'and general practice. *The British Journal of General Practice*, 62(605):664, 2012.

[38] Jacob R Gardner, Geoff Pleiss, David Bindel, Kilian Q Weinberger, and Andrew Gordon Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems*, 2018.

[39] Aimi S. Ghazali, Jaap Ham, Emilia I. Barakova, and Panos Markopoulos. Effects of Robot Facial Characteristics and Gender in Persuasive Human-Robot Interaction. *Frontiers in Robotics and AI*, 5:73, June 2018.

[40] Aimi S Ghazali, Jaap Ham, Emilia I Barakova, and Panos Markopoulos. Poker Face Influence: Persuasive Robot with Minimal Social Cues Triggers Less Psychological Reactance. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 940–946, Nanjing, August 2018. IEEE.

[41] Aimi Shazwani Ghazali, Jaap Ham, Emilia Barakova, and Panos Markopoulos. Assessing the effect of persuasive robots interactive social cues on users' psychological reactance, liking, trusting beliefs and compliance. *Advanced Robotics*, 33(7-8):325–337, April 2019.

[42] Marta C Gonzalez, Cesar A Hidalgo, and Albert-Laszlo Barabasi. Understanding individual human mobility patterns. *nature*, 453(7196):779–782, 2008.

[43] Michael A Goodrich and Dan R Olsen. Seven principles of efficient human robot interaction. In *SMC'03 Conference Proceedings. 2003 IEEE International Conference on Systems, Man and Cybernetics. Conference Theme-System Security and Assurance (Cat. No. 03CH37483)*, volume 4, pages 3942–3948. IEEE, 2003.

[44] Horst-Michael Gross, Steffen Mueller, Christof Schroeter, Michael Volkhardt, Andrea Scheidig, Klaus Debes, Katja Richter, and Nicola Doering. Robot companion for domestic health assistance: Implementation, test and case study under everyday conditions in private apartments. In

2015 IEEE/RSJ international conference on intelligent robots and systems (IROS), pages 5992–5999. IEEE, 2015.

[45] Edmund T Hall and Edward T Hall. *The hidden dimension*, volume 609. Anchor, 1966.

[46] J. Ham and C. J. H. Midden. A Persuasive Robot to Stimulate Energy Conservation: The Influence of Positive and Negative Social Feedback and Task Similarity on Energy-Consumption Behavior. *International Journal of Social Robotics*, 6(2):163–171, April 2014.

[47] Jaap Ham, Raymond H. Cuijpers, and John-John Cabibihan. Combining Robotic Persuasive Strategies: The Persuasive Power of a Storytelling Robot that Uses Gazing and Gestures. *International Journal of Social Robotics*, 7(4):479–487, August 2015.

[48] Mojgan Hashemian, Marta Couto, Samuel Mascarenhas, Ana Paiva, Pedro A. Santos, and Rui Prada. Persuasive Social Robot Using Reward Power over Repeated Instances of Persuasion. In Raian Ali, Birgit Lugrin, and Fred Charles, editors, *Persuasive Technology*, Lecture Notes in Computer Science, pages 63–70, Cham, 2021. Springer International Publishing.

[49] Mojgan Hashemian, Ana Paiva, Samuel Mascarenhas, Pedro A Santos, and Rui Prada. Social power in human-robot interaction: Towards more persuasive robots. In *AAMAS*, 2019.

[50] Fritz Heider. *The psychology of interpersonal relations*. Psychology Press, 2013.

[51] Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282–4286, May 1995. Publisher: American Physical Society.

[52] James Hensman, Alexander Matthews, and Zoubin Ghahramani. Scalable variational gaussian process classification. In *Artificial Intelligence and Statistics*, pages 351–360. PMLR, 2015.

[53] Sarita Herse, Jonathan Vitale, Daniel Ebrahimian, Meg Tonkin, Suman Ojha, Sidra Sidra, Benjamin Johnston, Sophie Phillips, Siva Leela Krishna Chand Gudi, Jesse Clark, William Judge, and Mary-Anne Williams. Bon Appetit! Robot Persuasion for Food Recommendation. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 125–126, Chicago IL USA, March 2018. ACM.

[54] Tahera Hossain, Wanggang Shen, Anindya Das Antar, Snehal Prabhudesai, Sozo Inoue, Xun Huan, and Nikola Banovic. A Bayesian Approach for Quantifying Data Scarcity when Modeling Human Behavior via Inverse Reinforcement Learning. *ACM Transactions on Computer-Human Interaction*, July 2022. Just Accepted.

[55] Matthew Huggins, Sharifa Alghowinem, Sooyeon Jeong, Pedro Colon-Hernandez, Cynthia Breazeal, and Hae Won Park. Practical guidelines for intent recognition: Bert with minimal training data evaluated in real-world hri application. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, pages 341–350, 2021.

[56] Tâm Huynh, Mario Fritz, and Bernt Schiele. Discovery of activity patterns using topic models. In *Proceedings of the 10th international conference on Ubiquitous computing*, UbiComp '08, pages 10–19, New York, NY, USA, September 2008. Association for Computing Machinery.

[57] Ifrah Idrees, Siddharth Singh, Kerui Xu, and Dylan F. Glas. A Framework for Realistic Simulation of Daily Human Activity. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 30–37, August 2023. ISSN: 1944-9437.

[58] Amazon.com Inc. Amazon alexa. apps.apple.com/us/app/amazon-alexa/id944011620, Last accessed April 2024.

[59] Amazon.com Inc. Meet astro, a home robot unlike any other. www.aboutamazon.com/news/devices/meet-astro-a-home-robot-unlike-any-other, Last accessed April 2024.

[60] Apple Inc. Use siri on all your apple devices. support.apple.com/en-us/105020, Last accessed April 2024.

[61] Universal Robots Inc. Ur5e lightweight, versatile cobot. www.universal-robots.com/products/ur5-robot/, Last accessed April 2024.

[62] iRobot Corp. irobot corporation: Our history. web.archive.org/web/20120103091646/http://www.irobot.com/sp.cfm?pageid=203, Last accessed April 2024.

[63] Roohollah Jahanmahin, Sara Masoud, Jeremy Rickli, and Ana Djuric. Human-robot interactions in manufacturing: A survey of human behavior modeling. *Robotics and Computer-Integrated Manufacturing*, 78:102404, 2022.

[64] Siddarth Jain and Brenna Argall. Probabilistic human intent recognition for shared autonomy in assistive robotics. *ACM Transactions on Human-Robot Interaction (THRI)*, 9(1):1–23, 2019.

[65] Sooyeon Jeong, Laura Aymerich-Franch, Kika Arias, Sharifa Alghowinem, Agata Lapedriza, Rosalind Picard, Hae Won Park, and Cynthia Breazeal. Deploying a robotic positive psychology coach to improve college students' psychological well-being. *User Modeling and User-Adapted Interaction*, 33(2):571–615, 2023.

[66] Charmi Jobanputra, Jatna Bavishi, and Nishant Doshi. Human activity recognition: A survey. *Procedia Computer Science*, 155:698–703, 2019.

[67] Wendy Ju and David Sirkin. Animate Objects: How Physical Motion Encourages Public Interaction. In *Persuasive Technology*, volume 6137, pages 40–51. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. Series Title: Lecture Notes in Computer Science.

[68] Rudolph E Kalman and Richard S Bucy. New results in linear filtering and prediction theory. 1961.

[69] Takayuki Kanda, Dylan F. Glas, Masahiro Shiomi, Hiroshi Ishiguro, and Norihiro Hagita. Who will be the customer? a social robot that anticipates people's behavior from their trajectories. In *Proceedings of the 10th international conference on Ubiquitous computing*, UbiComp '08, pages 380–389, New York, NY, USA, September 2008. Association for Computing Machinery.

[70] Harold Harding Kelley and John Thibaut. *The social psychology of groups*. Transaction Publishers, 1959.

[71] Adam Kendon. Studies in the behavior of social interaction. *(No Title)*, 1977.

[72] Adam Kendon. Goffman's approach to face-to-face interaction. *Erving Goffman: Exploring the interaction order*, 1988.

[73] Adam Kendon. Spacing and orientation in co-present interaction. *Development of Multimodal Interfaces: Active Listening and Synchrony: Second COST 2102 International Training School, Dublin, Ireland, March 23-27, 2009, Revised Selected Papers*, pages 1–15, 2010.

[74] Yunkyung Kim and Bilge Mutlu. How social distance shapes human–robot interaction. *International Journal of Human-Computer Studies*, 72(12):783–795, December 2014.

[75] Michel C. A. Klein, Nataliya Mogles, Jan Treur, and Arlette van Wissen. A Computational Model of Habit Learning to Enable Ambient Support for Lifestyle Change. In Kishan G. Mehrotra, Chilukuri K. Mohan, Jae C. Oh, Pramod K. Varshney, and Moonis Ali, editors, *Modern Approaches in Applied Intelligence*, Lecture Notes in Computer Science, pages 130–142, Berlin, Heidelberg, 2011. Springer.

[76] K. L. Koay, K. Dautenhahn, S. N. Woods, and M. L. Walters. Empirical results from using a comfort level device in human-robot interaction studies. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-robot Interaction*, HRI '06, pages 194–201, New York, NY, USA, March 2006. Association for Computing Machinery.

[77] Susan M. Krebs-Smith, TusaRebecca E. Pannucci, Amy F. Subar, Sharon I. Kirkpatrick, Jennifer L. Lerman, Janet A. Tooze, Magdalena M. Wilson, and Jill Reedy. Update of the healthy eating index: Hei-2015.

Journal of the Academy of Nutrition and Dietetics, 118(9):1591–1602, 2018.

[78] Christos Kyrlitsias and Despina Michael-Grigoriou. Social interaction with agents and avatars in immersive virtual environments: A survey. *Frontiers in Virtual Reality*, 2:786665, 2022.

[79] Jennifer S Labrecque, Kristen M Lee, and Wendy Wood. Measuring context–response associations that drive habits. *Journal of the Experimental Analysis of Behavior*, 121(1):62–73, 2024.

[80] JC Lafferty and AW Pond. The desert survival problem. plymouth, mi: Human synergistics, 1974.

[81] Phillippa Lally, Cornelia HM Van Jaarsveld, Henry WW Potts, and Jane Wardle. How are habits formed: Modelling habit formation in the real world. *European journal of social psychology*, 40(6):998–1009, 2010.

[82] Chin-Yew Lin and Franz Josef Och. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the 42nd annual meeting of the association for computational linguistics (ACL-04)*, pages 605–612, 2004.

[83] Google LLC. A more helpful google assistant for your every day. blog.google/products/assistant/ces-2020-google-assistant, Last accessed April 2024.

[84] Abbas Saliimi Lokman and Mohamed Ariff Ameedeen. Modern chatbot systems: A technical review. In *Proceedings of the Future Technologies Conference (FTC) 2018: Volume 2*, pages 1012–1023. Springer, 2019.

[85] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2005.

[86] Magnus S. Magnusson. Discovering hidden time patterns in behavior: T-patterns and their detection. *Behavior Research Methods, Instruments, & Computers*, 32(1):93–110, March 2000.

[87] Ross Mead and Maja J Matarić. Proxemics and performance: Subjective human evaluations of autonomous sociable robot distance and social signal understanding. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5984–5991, September 2015.

[88] Ross Mead and Maja J. Matarić. Perceptual Models of Human-Robot Proxemics. In M. Ani Hsieh, Oussama Khatib, and Vijay Kumar, editors, *Experimental Robotics: The 14th International Symposium on Experimental Robotics*, Springer Tracts in Advanced Robotics, pages 261–276. Springer International Publishing, Cham, 2016.

[89] Albert Mehrabian et al. *Silent messages*, volume 8. Wadsworth Belmont, CA, 1971.

[90] Jonathan Mumm and Bilge Mutlu. Human-robot proxemics: Physical and psychological distancing in human-robot interaction. In *Proceedings of the 6th International Conference on Human-robot Interaction*, HRI '11, pages 331–338, New York, NY, USA, March 2011. Association for Computing Machinery.

[91] Tatwadarshi P Nagarhalli, Vinod Vaze, and NK Rana. A review of current trends in the development of chatbot systems. In *2020 6th International conference on advanced computing and communication systems (ICACCS)*, pages 706–710. IEEE, 2020.

[92] Stefanos Nikolaidis, David Hsu, and Siddhartha Srinivasa. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, 36(5-7):618–634, 2017.

[93] Mohammad Obaid, Eduardo B. Sandoval, Jakub Złotowski, Elena Moltchanova, Christina A. Basedow, and Christoph Bartneck. Stop! That is close enough. How body postures influence human-robot proximity. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 354–361, August 2016.

[94] Daniel J O'Keefe. The elaboration likelihood model. *The SAGE handbook of persuasion: Developments in theory and practice*, pages 137–149, 2013.

[95] Fotios Papadopoulos, Dennis Küster, Lee J. Corrigan, Arvid Kappas, and Ginevra Castellano. Do relative positions and proxemics affect the engagement in a Human-Robot collaborative scenario? *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, 17:321–347, 2016.

[96] Maithili Patel and Sonia Chernova. Proactive robot assistance via spatio-temporal object modeling. *arXiv preprint arXiv:2211.15501*, 2022.

[97] Alex Pentland and Andrew Liu. Modeling and Prediction of Human Behavior. *Neural Computation*, 11(1):229–242, January 1999.

[98] Emma Pierson, Tim Althoff, and Jure Leskovec. Modeling Individual Cyclic Variation in Human Behavior. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, pages 107–116, Republic and Canton of Geneva, CHE, April 2018. International World Wide Web Conferences Steering Committee.

[99] J Prochaska and W Velicer. The transtheoretical model of health behavior change. *American Journal of Health Promotion*, 12:38–48, 1997.

[100] James O Prochaska and Wayne F Velicer. The transtheoretical model of health behavior change. *American journal of health promotion*, 12(1):38–48, 1997.

[101] Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. VirtualHome: Simulating Household Activities via Programs. pages 8494–8502, 2018.

[102] Deepak Ramachandran and Eyal Amir. Bayesian inverse reinforcement learning. In *Proceedings of the 20th international joint conference on Artifical intelligence*, IJCAI'07, pages 2586–2591, San Francisco, CA, USA, January 2007. Morgan Kaufmann Publishers Inc.

[103] Kathleen Kelley Reardon. *Persuasion in practice*. Sage, 1991.

[104] Siddharth Reddy, Anca D. Dragan, and Sergey Levine. Where Do You Think You're Going?: Inferring Beliefs about Dynamics from Behavior, January 2019. arXiv:1805.08010 [cs, stat].

[105] Virginia P Richmond, Joan S Gorham, and James C McCroskey. The relationship between selected immediacy behaviors and cognitive learning. *Annals of the International Communication Association*, 10(1):574–590, 1987.

[106] Jaime A. Rincon, Angelo Costa, Paulo Novais, Vicente Julian, and Carlos Carrascosa. A new emotional robot assistant that facilitates human interaction and persuasion. *Knowledge and Information Systems*, 60(1):363–383, July 2019.

[107] SoftBank Robotics. Pepper the humanoid and programmable robot, October 2022.

[108] David L Ronis, J Frank Yates, and John P Kirscht. Attitudes, decisions, and habits as determinants of repeated behavior. In *Attitude structure and function*, pages 213–239. Psychology Press, 2014.

[109] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. Human motion trajectory prediction: A survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020.

[110] Shane Saunderson and Goldie Nejat. How robots influence humans: A survey of nonverbal communication in social human–robot interaction. *International Journal of Social Robotics*, 11:575–608, 2019.

[111] Shane Saunderson and Goldie Nejat. How Robots Influence Humans: A Survey of Nonverbal Communication in Social Human–Robot Interaction. *International Journal of Social Robotics*, 11(4):575–608, August 2019.

[112] Francesco Setti, Chris Russell, Chiara Bassetti, and Marco Cristani. F-formation detection: Individuating free-standing conversational groups in images. *PloS one*, 10(5):e0123783, 2015.

[113] Paschal Sheeran, Henk Aarts, Ruud Custers, Amanda Rivis, Thomas L Webb, and Richard Cooke. The goal-dependent automaticity of drinking habits. *British Journal of Social Psychology*, 44(1):47–63, 2005.

[114] Michael Steven Siegel. *Persuasive Robotics How Robots Change Our Minds*. PhD thesis, Massachusetts Institute of Technology, 2008.

[115] Mikey Siegel, Cynthia Breazeal, and Michael I. Norton. Persuasive Robotics: The influence of robot gender on human behavior. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2563–2568, St. Louis, MO, USA, October 2009. IEEE.

[116] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. Limits of predictability in human mobility. *Science*, 327(5968):1018–1021, 2010.

[117] James B Stiff and Paul A Mongeau. *Persuasive communication*. Guilford Publications, 2016.

[118] Feng-Tso Sun, Yi-Ting Yeh, Heng-Tze Cheng, Cynthia Kuo, and Martin Griss. Nonparametric discovery of human routines from sensor data. In *2014 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 11–19, March 2014.

[119] Luise Süssenbach, Nina Riether, Sebastian Schneider, Ingmar Berger, Franz Kummert, Ingo Lütkebohle, and Karola Pitsch. A robot as fitness companion: towards an interactive action-based motivation model. In *The 23rd IEEE international symposium on robot and human interactive communication*, pages 286–293. IEEE, 2014.

[120] Dag Sverre Syrdal, Kerstin Dautenhahn, Kheng Lee Koay, Michael L. Walters, and Wan Ching Ho. Sharing Spaces, Sharing Lives – The Impact of Robot Mobility on User Perception of a Home Companion Robot. In *Social Robotics*, volume 8239, pages 321–330. Springer International Publishing, Cham, 2013.

[121] Dag Sverre Syrdal, Kerstin Dautenhahn, Michael L Walters, and Kheng Lee Koay. Sharing spaces with robots in a home scenario-anthropomorphic attributions and their effect on proxemic expectations and evaluations in a live hri trial. In *AAAI fall symposium: AI in Eldercare: new solutions to old problems*, volume 2008, pages 116–123, 2008.

[122] Yunus Terzioğlu, Bilge Mutlu, and Erol Şahin. Designing social cues for collaborative robots: the role of gaze and breathing in human-robot collaboration. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*, pages 343–357, 2020.

[123] Yunus Terzioğlu, Keith Rebello, and Timothy Bickmore. Influencing health-related decision making and therapeutic alliance with robot mobility and deixis. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 627–633. IEEE, 2023.

[124] Yunus Terzioğlu, Keith Rebello, and Timothy Bickmore. Influencing health related decision making using robot counseling, mobility, and gesture. *Journal in review*, 2024.

[125] Enchanted Tools. Meet the mirokai-welcome to a post-robot era. https://enchanted.tools/, Last accessed April 2024.

[126] C-Y Ung, M Menozzi, C Hartmann, and M Siegrist. Innovations in consumer research: The virtual food buffet. *Food Quality and Preference*, 63:12–17, 2018.

[127] Agricultural Research Service U.S. Department of Agriculture. Usda food and nutrient database for dietary studies 2017-2018, October 2020.

[128] Beltsville Human Nutrition Research Center U.S. Department of Agriculture. Food patterns equivalents database per 100 grams of fndds 2017-2018 foods, October 2021.

[129] Elisabeth T Van Dijk, Elena Torta, and Raymond H Cuijpers. Effects of eye contact and iconic gestures on message retention in human-robot interaction. *International Journal of Social Robotics*, 5(4):491–501, 2013.

[130] Paul AM Van Lange and Daniel Balliet. *Interdependence theory.* American Psychological Association, 2015.

[131] M. L. Walters, D. S. Syrdal, K. L. Koay, K. Dautenhahn, and R. te Boekhorst. Human approach distances to a mechanical-looking robot with different robot voice styles. In *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, pages 707–712, August 2008.

[132] M.L. Walters, K. Dautenhahn, R. te Boekhorst, Kheng Lee Koay, C. Kaouri, S. Woods, C. Nehaniv, D. Lee, and I. Werry. The influence of subjects' personality traits on personal spatial zones in a human-robot interaction experiment. In *ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005.*, pages 347–352, Nashville, TN, USA, 2005. IEEE.

[133] Xuerui Wang, Andrew McCallum, and Xing Wei. Topical N-Grams: Phrase and Topic Discovery, with an Application to Information Retrieval. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pages 697–702, October 2007. ISSN: 2374-8486.

[134] Jürgen Wiest, Matthias Höffken, Ulrich Kreßel, and Klaus Dietmayer. Probabilistic trajectory prediction with Gaussian mixture models. In *2012 IEEE Intelligent Vehicles Symposium*, pages 141–146, June 2012. ISSN: 1931-0587.

[135] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

[136] Robert C Wilson and Anne GE Collins. Ten simple rules for the computational modeling of behavioral data. *Elife*, 8:e49547, 2019.

[137] Katie Winkle, Séverin Lemaignan, Praminda Caleb-Solly, Paul Bremner, Ailie J Turton, and Ute Leonards. In-situ learning from a domain expert for real world socially assistive robot deployment. In *Robotics: Science and Systems*, 2020.

[138] Katie Winkle, Severin Lemaignan, Praminda Caleb-Solly, Ute Leonards, Ailie Turton, and Paul Bremner. Effective Persuasion Strategies for Socially Assistive Robots. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 277–285, Daegu, Korea (South), March 2019. IEEE.

[139] Katie Winkle, Séverin Lemaignan, Praminda Caleb-Solly, Ute Leonards, Ailie Turton, and Paul Bremner. Effective persuasion strategies for socially assistive robots. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 277–285. IEEE, 2019.

[140] Jacob O Wobbrock, Leah Findlater, Darren Gergle, and James J Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 143–146, 2011.

[141] Xuhai Xu, Prerna Chikersal, Afsaneh Doryab, Daniella K. Villalba, Janine M. Dutcher, Michael J. Tumminia, Tim Althoff, Sheldon Cohen, Kasey G. Creswell, J. David Creswell, Jennifer Mankoff, and Anind K. Dey. Leveraging Routine Behavior and Contextually-Filtered Features for Depression Detection among College Students. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(3):116:1–116:33, September 2019.

[142] Xuhai Xu, Xin Liu, Han Zhang, Weichen Wang, Subigya Nepal, Yasaman Sefidgar, Woosuk Seo, Kevin S. Kuehn, Jeremy F. Huckins, Margaret E. Morris, Paula S. Nurius, Eve A. Riskin, Shwetak Patel, Tim Althoff, Andrew Campbell, Anind K. Dey, and Jennifer Mankoff. GLOBEM: Cross-Dataset Generalization of Longitudinal Human Behavior Modeling. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(4):190:1–190:34, January 2023.

[143] Anna K. Yanchenko, Di Daniel Deng, Jinglan Li, Andrew J. Cron, and Mike West. Hierarchical Dynamic Modeling for Individualized Bayesian Forecasting, January 2021. arXiv:2101.03408 [stat].

[144] Chao Zhang, Joaquin Vanschoren, Arlette van Wissen, Daniël Lakens, Boris de Ruyter, and Wijnand A. IJsselsteijn. Theory-based habit modeling for enhancing behavior prediction in behavior change support systems. *User Modeling and User-Adapted Interaction*, 32(3):389–415, July 2022.

[145] Brian D. Ziebart, J. Andrew Bagnell, and Anind K. Dey. Modeling Interaction via the Principle of Maximum Causal Entropy. June 2010. Publisher: Carnegie Mellon University.

[146] Brian D Ziebart, Andrew Maas, J Andrew Bagnell, and Anind K Dey. Human Behavior Modeling with Maximum Entropy Inverse Optimal Control. 2009.

[147] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.

[148] Philip G Zimbardo and Michael R Leippe. *The psychology of attitude change and social influence.* Mcgraw-Hill Book Company, 1991.