

Movement Coordination in Human–Robot Teams: A Dynamical Systems Approach

Tariq Iqbal, *Student Member, IEEE*, Samantha Rack, *Student Member, IEEE*,
and Laurel D. Riek, *Senior Member, IEEE*

Abstract—In order to be effective teammates, robots need to be able to understand high-level human behavior to recognize, anticipate, and adapt to human motion. We have designed a new approach to enable robots to perceive human group motion in real time to anticipate future actions and synthesize their own motion accordingly. We explore this within the context of joint action, in which humans and robots move together synchronously. In this paper we present an anticipation method, which takes high-level group behavior into account. We validate the method within a human–robot interaction scenario, in which an autonomous mobile robot observes a team of human dancers and then successfully and contingently coordinates its movements to “join the dance.” We compared the results of our anticipation method to move the robot with another method that did not rely on high-level group behavior and found that our method performed better both in terms of more closely synchronizing the robot’s motion to the team and exhibiting more contingent and fluent motion. These findings suggest that the robot performs better when it has an understanding of high-level group behavior than when it does not. This study will help enable others in the robotics community to build more fluent and adaptable robots in the future.

Index Terms—Coordination, human–robot collaboration, human–robot joint action, movement analysis, synchronization.

I. INTRODUCTION

As technology advances, autonomous robots are becoming more involved in human society in a variety of roles. Robotic systems have long been involved in assembly lines automating and increasing efficiency of monotonous or dangerous factory procedures [1]. However, as robots leave controlled spaces and begin to work alongside people in teams, many things taken for granted in robotics concerning perception and action do not apply, as people act unpredictably, e.g., they “break the rules” when it comes to what a robot can expect *a priori*. In order for robots to effectively perform their tasks and integrate into human social environments (HSEs), they must be able to comprehend high-level social signals and respond appropriately [2].

While working alongside humans, a robot might encounter people performing various social actions, such as engaging in social activities or performing synchronous movements [3]. For example, Ros, Baroni, and Demiris [4] used a humanoid robot to play the role of a dance instructor with children, and Fasola and

Mataric [5] designed a socially assistive robot to engage older adults in physical exercise. Others have used robots to dance and play cooperatively with children in therapeutic settings [6], [7]. Koenemann, Burget, and Bennewitz demonstrated a system that enabled humanoid robots to imitate complex human whole-body motion [8].

However, sometimes it can be difficult for a robot to perceive and understand all of the different types of events involved during these activities to make effective decisions, due to sensor occlusion, unanticipated motion, narrow field of view, etc. On the other hand, if a robot is able to make better sense of its environment and understand high-level group dynamics, then it can make effective decisions about its actions. If the robot has this understanding of its environment, then its interactions within the team might reach to a higher level of coordination, resulting in a fluent meshing of actions [9]–[12].

Human activity recognition from body movement is an active area of research across many fields [13]–[17]. These activities involve a wide range of behaviors, from gross motor motion (e.g., walking, lifting) to manipulation (e.g., stacking objects). All of these experiments showed impressive results in recognizing activities, either performed by an individual or dyad.

However, the focus of most of these methods is to recognize the activity of a single human, rather than to understand the team’s dynamics and how it might affect behavior. This understanding is critical in human–robot interaction scenarios, as the “one human, one robot” paradigm is rarely seen in ecological settings [18], [19]. To make informed decisions, robots need to understand this context [10].

Many disciplines have investigated interaction dynamics within groups, which include sociology, psychology, biology, music, and dance [20]–[31]. For example, Nagy *et al.* [20], [21] investigated collective behavior on animals and developed automated methods for assessing social dominance and leadership in domestic pigeons. Their investigation explored the effect of social hierarchical structure on dominance and leadership. Their results indicated that dominance and leadership hierarchical structures were independent from each other.

Inspired from bird flocks and fish schools, Leonard *et al.* [22], [23] investigated how collective group motion emerges when basic animal flocking rules (i.e., cohesive and repulsive element) are applied to a group of human dancers. Using tracked trajectories of head positions of individual dancers, the authors developed a time-varying graph-based method to infer conditions under which certain dancers emerged as the leaders of the group.

Synchronous motion, or joint action, is a common type of high-level behavior encountered in human–human interaction.

Manuscript received August 01, 2015; revised January 31, 2016; accepted April 05, 2016. Date of publication June 20, 2016; date of current version August 18, 2016. This paper was recommended for publication by Editor A. Kheddar and Guest Editors D. Kulic and G. Venture upon evaluation of the reviewers’ comments.

The authors are with the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556 USA (e-mail: tiqbal@nd.edu; srack@nd.edu; lriek@nd.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TRO.2016.2570240

It is a naturally present social interaction that occurs when two or more participants coordinate their actions in space and time to make changes to the environment [32]. Understanding synchronous joint action (SJA) is important, as it helps to accurately understand the affective behavior of a team and also provides information regarding the group level cohesiveness [33], [34]. Thus, if a robot has the ability to understand the presence of SJA in a team, then it can use that information to inform its own actions to enable coordinated movement with the team. It also might learn advanced adaptive coordination techniques that human teams use, such as tempo adaptation or cross-training [1], [35].

Many approaches have been taken by researchers across different fields to measure the degree of synchronization in continuous time series data, including recurrence analysis [36], correlation [37], and phase difference approaches [33]. Other sets of methods work across categorical time series data, which may define discrete events [38]. However, these event-based methods only consider a single type of event while measuring synchronization. To address this gap, we created an event-based method that can successfully take multiple types of discrete task-level events into consideration while measuring the degree of synchronization of a system [39].

Recent work in robotics has focused on developing predictive methods for improving the fluency of a joint interaction between a robot and one or more humans. For example, Hawkins, Bansal, Vo, and Bobick [40] developed a method that determines an appropriate action for an assistive robot to take when providing parts during an assembly activity. They employ a probabilistic model that considers the presence of variability in the human's actions. Hoffman and Breazeal [41] proposed an adaptive action-selection mechanism for a robot, which could make anticipatory decisions based on confidence of their validity and their relative risks. Through a study the authors validated the model and presented an improvement in task efficiency when compared with a purely reactive model.

Additionally, Pérez-D'Arpino and Shah [42] proposed a data-driven approach to synthesize anticipatory knowledge of human motion, which they used to predict targets during reaching motions. Unhelkar, Pérez-D'Arpino, Stirling, and Shah [43] extended this concept for a human-robot conavigation task. This model used an "human turn signals" during walking as anticipatory indicators, in order to predict human motion trajectories. This knowledge was then used for motion planning in simulated dynamic environments.

While this study will improve the ability of robots to have fluent interactions within HSEs, most of these methods are best suited for dyadic interaction and dexterous manipulation contexts. In contrast, we seek to explore methods for robots that will work robustly in groups, as well as for tasks involving gross motion with mobile robots.

In our prior work, we explored the problem of automatically modeling and detecting SJA in human teams, using both fixed and mobile sensors. We introduced a new nonlinear dynamical method, which performed more accurately and robustly than existing methods [10], [39].

In this paper we explore how a robot can use these models to synthesize SJA in order to coordinate its movements

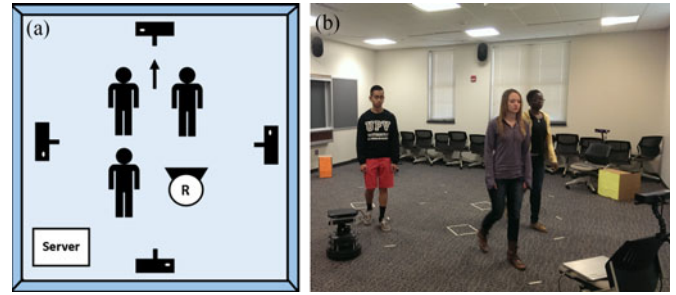


Fig. 1. (a) Data acquisition setup. (b) Three participants are dancing along with a Turtlebot robot.

with a human team. The main contribution of this study is the introduction of a new method for anticipation of robot motion that takes human group dynamics into account. We validated our method within a human-robot interaction scenario, where an autonomous mobile robot observes a team of human dancers and then successfully and contingently coordinates its movements to "join the dance." We compared the results of our anticipation method with another method that does not rely on high-level group behavior. Our method performed better both in terms of more closely synchronizing the robot's motion to the team and exhibiting more contingent and fluent motion.

The outline of this paper is as follows. Section II describes the experimental testbed for studying SJA and the system architecture. Then, Section III provides details of the two anticipation methods. Section IV describes the experimental procedure. Sections V and VI discuss how the data were preprocessed and the experimental results. Finally, Section VII discusses the implication of these findings for the robotics community.

II. SYSTEM ARCHITECTURE AND EXPERIMENTAL TESTBED

In order to explore how a robot can use human group dynamics to synthesize SJA with a mixed team, we needed an experimental testbed in which a robot could perform tasks synchronously with humans. We also required a group activity in which each member's actions would have an impact on others' actions, as well as have an impact on the dynamics of the group overall.

Therefore, we designed a movement task in which a team of humans and a robot could coordinate their motion in real time. Specifically, we explored SJA within the context of synchronous dance. In concert with an experienced dancer, we choreographed a routine to the song Smooth Criminal by Michael Jackson, which is in 4/4 time. The dance is iterative and performed cyclically in a counterclockwise manner [see Fig. 1(a)] There are four iterations in a dance session, corresponding to each of the cardinal directions (north, west, south, and east). Each iteration includes the dancers taking the following steps in order: move forward and backward twice, then clap and turn 90° (see Fig. 2) [3].

A. Data Acquisition Process

Fig. 1(a) shows the data acquisition setup. Three human participants and a Turtlebot v.2 robot were arranged in two rows. Four Microsoft Kinect v.2 sensors were positioned approximately three feet above the ground at each of the cardinal




Fig. 2. Five high-level events were detected during the dance using skeletal data from participants. One iteration of the dance consists of two passes of the start moving forward, stop moving forward, start moving backward, stop moving backward events, and a clap event in order.

directions. Each sensor was connected to a computer (client) to capture and process the depth, infrared, and skeletal data from the Kinect. All four clients and the server ran Windows 8 on an Intel Core i5 processor at 1.70 Hz with 12 GB of RAM.

As we are studying synchronous activity, it was critical that all clients and the robot maintained a consistent time reference. Thus, we created a server to manage communication and global time synchronization. Synchronization architecture details can be found in [44].

Each client performed real-time processing of the raw data in order to detect dance events (e.g., move forward, stop, etc.), which it sent to the server with a timestamp. When the server received data from the clients, it generated predictions for how the robot should move using one of two anticipation methods, which are described in Section III. The server was also responsible for determining the active client, which refers to which of the four sensors the dancers were facing during a given iteration.

In order to allow for offline analysis, the clients also recorded time-synchronized depth, infrared, audio, and skeletal data using an automated interface with Kinect Studio. The server and robot also kept detailed logs of all communication, odometry information, events received from the active client, and information about the dancers.

B. Client-Side Data Processing

We extracted five high-level events from the participants' movements during the dance: start moving forward, stop moving forward, start moving backward, stop moving backward, and clap. The start moving forward event is detected when a participant begins approaching the Kinect, and stop moving forward when they stop moving. Similarly, when a participant moves away from the sensor (backward), that is identified as a start moving backward event, and when they stop, stop moving backward is identified. We also detected participants' clap events, which occurred at the end of each iteration. See Fig. 2.

To detect these events from participants' body movements, clients used the skeletal positions provided by the Kinect. Clients calculated forward and backward motion onsets along the z -axis primarily using the spine base body joint position, as it is the most stable and reliable joint position when participants are in motion.

However, there were times when participants did not move their spine base, but did move their mid-spine, shoulders, or neck, to signal the onset of motion. Therefore, clients also used these positions, again along the z -axis, to detect four additional events: early start moving forward, early stop moving forward,

early start moving backward, and early stop moving backward. For these early events, clients calculated joint change positions by comparing the current and previous frame. If at least half of the joint positions changed, then it indicated the participant had started moving. To detect clap events, clients used the x and y coordinates from the 3-D skeletal position of the left and right hand and shoulder joints. Claps occurred when the ratio of the distance between the hands and the distance between the shoulder joints was less than a threshold (0.6) and when this ratio value reaches a local minima.

C. Robot Command Generation and Execution

After the server determines which movement the robot should make, which it does by using an anticipation methods described in Section III, it sends a movement command to the robot. These commands include: move forward, move backward, stop, and turn. The server translated the clap commands into rotation commands while sending it to the robot, since the robot cannot clap.

The robot, which ran the robot operating system (ROS) version Hydro on Ubuntu version 12.04, accepted commands from the server, parsed the commands, and used an ROS publisher to send movement commands to the controller. The robot is capable of forward and backward movement and can rotate on its vertical axis in either direction.

III. EVENT ANTICIPATION METHODS

For this work, we created two anticipation methods to move the robot. The first method, i.e., synchronization-index based anticipation (SIA), is inspired by our prior SJA detection work [39]. It calculates the synchronicity of the group in real time, determines who the most synchronous dancer is, and uses that information to move the robot. For the second method, event cluster-based anticipation (ECA), we created to establish a reasonable comparison anticipation method for SIA that does not rely on group dynamics. ECA is a straightforward method that involves averaging the times participants moved during a previous iteration of the dance. Fig. 3 gives a visual comparison of how the two methods work in practice and they are described textually below.

A. Synchronization Index-Based Anticipation

The SIA method takes a group's internal dynamics into account when generating robot movements. The main idea is that, for a given iteration, the participant who moves the most

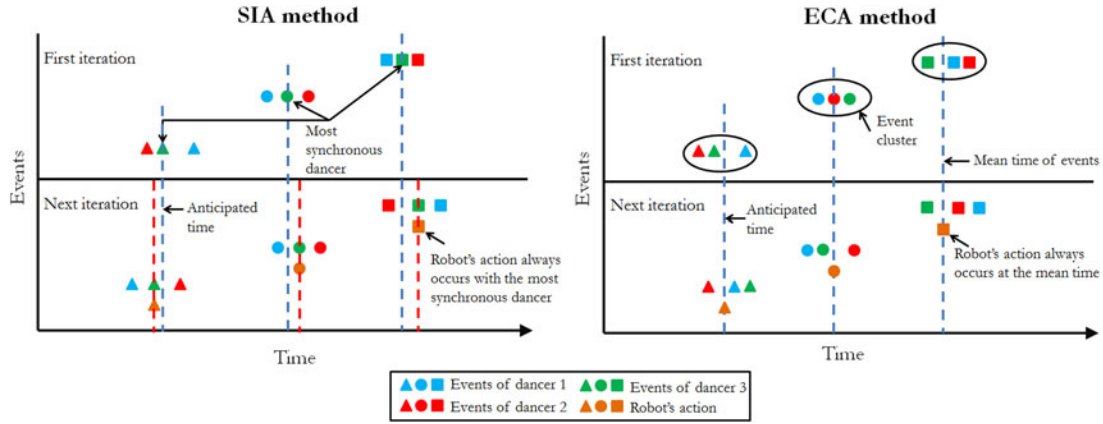


Fig. 3. Visualization of the two anticipation methods. (Left) SIA. (Right) ECA. The upper portion of the graph represents one iteration of the dance session and the lower portion represents the next iteration of the same session.

synchronously with the other dancers is a good model for the robot to follow in order to be well coordinated with the team. In addition, the method will adjust its identification of the most synchronous dancer after each iteration. Fig. 3(b) explains this method visually.

Thus, to generate future actions for the robot using this method, at the beginning of each iteration we measured the most synchronous person of the group using the nonlinear dynamical method we described in [39]. We will briefly describe the method in Sections III-A1 and III-A2 and then discuss in Section III-A3 how we used the method to assess the most synchronous dancer to inform how the robot should move.

1) *Measuring Synchronization of Events Across Two Time Series:* We can express the task-level events associated with each dancer as a time series. Suppose x_n and y_n are two time series, where $n = 1 \dots N$. Here, each time series has N samples. Suppose, m_x and m_y are the number of events occurring in time series x and y respectively, and E is the set of all events [39].

The events of both series are denoted by $e_x(i) \in E$ and $e_y(j) \in E$, where $i = 1 \dots m_x$, $j = 1 \dots m_y$. The event times on both time series are t_i^x and t_j^y ($i = 1 \dots m_x$, $j = 1 \dots m_y$), respectively [39].

In the case of synchronous events in both time series, the same event should appear roughly at the same time or within a time lag $\pm\tau$ [39].

Now suppose that $c^\tau(x|y)$ denotes the number of times a single type of event $e \in E$ appear in time series x shortly after they appear in time series y . Here

$$c^\tau(x|y) = \sum_i^{m_x} \sum_j^{m_y} J_{ij}^\tau \quad (1)$$

where

$$J_{ij}^\tau = \begin{cases} 1, & \text{if } 0 < t_i^x - t_j^y < \tau \\ \frac{1}{2}, & \text{if } t_i^x = t_j^y \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

Similarly, we can calculate $c^\tau(y|x)$ denoting the number of times a single type of event $e \in E$ appear in a time series y shortly after they appear in time series x .

Now, $Q_\tau(e)$ represents the synchronization of events in two time series, in which we are only considering a single type of event e in both time series. From $c^\tau(x|y)$ and $c^\tau(y|x)$, we can calculate $Q_\tau(e)$ as

$$Q_\tau(e) = \frac{c^\tau(x|y) + c^\tau(y|x)}{\sqrt{m_x m_y}}. \quad (3)$$

The value of $Q_\tau(e)$ should be in between 0 and 1 ($0 \leq Q_\tau(e) \leq 1$), as we normalize it by the number of events that appear in both time series. $Q_\tau(e) = 1$ shows that all the events of both time series are fully synchronized and appeared within a time lag $\pm\tau$ on both time series. On the other hand, $Q_\tau(e) = 0$ shows us that the events are asynchronous [39].

Now we extend the notion of synchronization of events in two time series for multiple types of events. Suppose we have n types of events $\{e_1, e_2, \dots, e_n\} \in E(n)$, where $E(n)$ is the set of all types of events. First, we calculate $Q_\tau(e_i)$ for each event type $e_i \in E(n)$. While calculating $Q_\tau(e_i)$, we will not consider any other type of event, except e_i [39].

Now let $m_x(e_i)$ be the number of events of type e_i occurring in time series x and $m_y(e_i)$ is the number of events of type e_i occurring in time series y . To measure synchronization of multiple types of events between two time series, we take the average of $Q_\tau(e_i)$ for each event type e_i , weighted by the number of events of that type. We will call this the synchronization index of that pair [39].

Therefore, the overall synchronization of events in time series x and y of that pair is

$$\forall e_i \in E(n) : Q_\tau^{xy} = \frac{\sum [Q_\tau(e_i) \times [m_x(e_i) + m_y(e_i)]]}{\sum [m_x(e_i) + m_y(e_i)]}. \quad (4)$$

If all events are synchronous in both time series, then the value of Q_τ^{xy} will be 1. If no synchronous are synchronous, the value of Q_τ^{xy} will be 0 [39].

2) *Measuring the Individual Synchronization Index:* We calculated the pairwise synchronization index for each pair.

Suppose we have H number of time series. The time series data are represented as s_1, s_2, \dots, s_H . First, we calculate the pairwise event synchronization index for each pair. So, we have the value of $Q_\tau^{s_1 s_2}, Q_\tau^{s_1 s_3}, \dots, Q_\tau^{s_{(H-1)} s_H}$ [39].

We modified our process slightly from the description in [39]. After calculating the pairwise synchronization index, we built a directed weighted graph from these indices, in which each time series is represented by a vertex. However, in [39], after calculating the pairwise synchronization index, an undirected weighted graph was built. In a fully connected situation, the directed and the undirected graph represents the same connectivity.

Therefore, if the time series are s_1, s_2, \dots, s_H , then there is a vertex in the graph that will correspond to a time series. We connect a pair of vertices with a weighted edge, based on their synchronization index value. In this case, there will be an incoming and an outgoing edge for each pair of vertices. We will refer to this graph as the group topology graph (GTG) [39].

The individual synchronization index ($I_\tau(s_i)$) depends on both the group composition and the size of the group. We assumed that, during this dance performance, each human participant may have some direct or indirect influences on the other human participants of the group [39]. $I_\tau(s_i)$ for a participant is measured as the average of the weight of the outgoing edges to the corresponding vertex in the topology graph. Therefore, the $I_\tau(s_i)$ of series s_i is

$$I_\tau(s_i) = \frac{\sum_{j=1, \dots, H, j \neq i} Q_\tau^{s_i s_j} \times f(s_i, s_j)}{\sum_{j=1, \dots, H, j \neq i} f(s_i, s_j)} \quad (5)$$

where

$$f(s_i, s_j) = \begin{cases} 1, & \text{iff edge}(s_i, s_j) \in \text{GTG} \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

3) Determining the Most Synchronous Dancer and Anticipating Their Next Movement: The person with the highest individual synchronization index during an iteration is considered to be the most synchronous person of the group. This is because a high individual synchronization index indicates close synchronization with the other group members. Thus, let this person be MSP.

Suppose, during itr_i , we determine $\text{MSP}(\text{itr}_i)$ as the most synchronous dancer of the group. Now, assuming that a similar timing pattern of events will occur during the next iteration ($\text{itr}_{(i+1)}$), if the robot follows the events of the $\text{MSP}(\text{itr}_i)$, then the group will become more synchronous.

We can describe this concept mathematically. To reach a synchronous state, all events must occur very closely in time, i.e., within a time lag $\pm\tau$. Thus, we want to minimize the difference between event timings for each pair of agents. Now, if Δt_{ij} represents the time difference of one event between agent i and j , then our goal is

$$\forall i, j \in H : \min \left(\sum \Delta t_{ij} \right). \quad (7)$$

Now, for our scenario, as shown in Fig. 4, suppose Dancer 2 was the most synchronous person during one iteration (itr_i) of the dance session, i.e., $\text{MSP}(\text{itr}_i)$ was Dancer 2. Now, during $\text{itr}_{(i+1)}$, a similar timing pattern holds and the timing of one

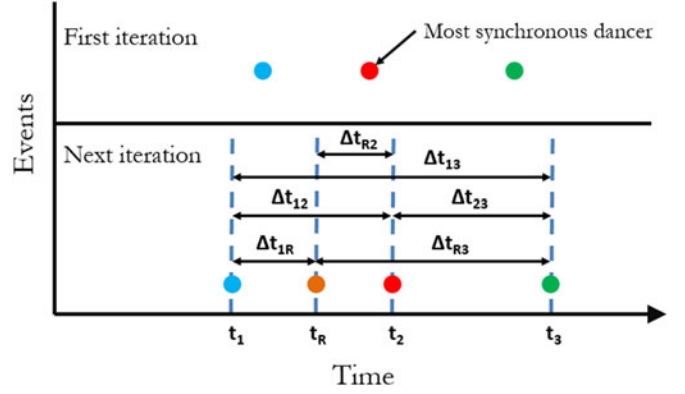


Fig. 4. Example timings from a single type of event during two consecutive iterations.

particular event of the three dancers and the robot are t_1, t_2, t_3 , and t_R , respectively. To reach a synchronous state, the following is required:

$$\min \left(\sum \Delta t_{12} + \Delta t_{23} + \Delta t_{1R} + \Delta t_{R3} + \Delta t_{13} + \Delta t_{R2} \right). \quad (8)$$

As Dancer 2 is the MSP, from Fig. 4 one can see $\Delta t_{12} + \Delta t_{23} = \Delta t_{13}$, and $\Delta t_{1R} + \Delta t_{R3} = \Delta t_{13}$. Thus, 8 becomes

$$\min \left(\sum \Delta t_{13} + \Delta t_{13} + \Delta t_{13} + \Delta t_{R2} \right). \quad (9)$$

As only the term Δt_{R2} depends on the robot's movement in (9), by minimizing Δt_{R2} we can minimize the equation. Thus, if the robot and the Dancer 2 (in this case, the MSP) perform the same event at the same time, then Δt_{R2} will become 0, which will minimize (9). This implies that if the robot can perform the events close to the timings of the most synchronous person, then the whole group will reach a more synchronous state.

Thus, for a given iteration, $\text{itr} \in \forall$ iterations, the server will determine $\text{MSP}(\text{itr}_i)$. Then, during the next iteration, $\text{itr}_{(i+1)}$, the server will track all movements of $\text{MSP}(\text{itr}_i)$. The server then processes the information by utilizing the early detected events (early start moving forward, early stop moving forward, early start moving backward, and early stop moving backward) following the method described in the next paragraph.

As we know the timing of the events during the previous iteration of the dance $\text{itr}_{(i-1)}$, our anticipation method assumed that the similar event will happen more or less at the same time during this iteration itr_i . Therefore, when it was close to the timing of events of $\text{MSP}(\text{itr}_{(i-1)})$ during itr_i , and the server received early detected events associated with $\text{MSP}(\text{itr}_{(i-1)})$, then the server anticipated those events as the indicator of the start of a movement. The server then sent appropriate commands to the robot to perform that movement.

For example, suppose Dancer 2 was the most synchronous person during iteration 1, i.e., $\text{MSP}(\text{itr}_1)$ was Dancer 2. Dancer 2 performed a start moving forward event three seconds from the start of itr_1 . Therefore, during itr_2 , it was assumed that the start moving forward would happen three seconds from the iteration's

start. Thus, if the server received a sufficient number of early start moving forward events around t_3 , then it notified the robot command generator to generate commands to execute forward movement. This process was similar for all other regular events, excluding the clap event.

B. Event Cluster-Based Anticipation Method

We created the ECA method to establish a reasonable comparison anticipation method for SIA that does not rely on group dynamics. ECA is theoretically simple but powerful in nature. As the dance is rhythmic and iterative in nature, the movement events for one iteration are similar to events that happened in the previous iteration. Thus, we averaged the events timing during one iteration to predict the timing of those same events for the next iteration. Fig. 3(a) explains this method visually.

First, for one iteration, we presented all the events associated with the dancers by a time series. Thus, this time series represented all the events of that iteration. Then, we clustered all the similar types of events together, those that happened within a time threshold, ϵ . For example, for a single event e , we calculated the timing of the event performed by three human participants, i.e., $t(\text{dancer}_1(\text{itr}_i), e)$, $t(\text{dancer}_2(\text{itr}_i), e)$, and $t(\text{dancer}_3(\text{itr}_i), e)$. Here, t represents the timing of an event and itr_i represents the iteration i .

After that, for each cluster of similar events, we calculated the average time of all the events and used that time as the timing of the event for the next iteration. These events and the times were the predicted events and timing for the next iteration of the dance. Thus, $t(\text{robot}(\text{itr}_{(i+1)}), e) = (t(\text{dancer}_1(\text{itr}_i), e) + t(\text{dancer}_2(\text{itr}_i), e) + t(\text{dancer}_3(\text{itr}_i), e))/3$. After the prediction of all the events for next iteration, the method sends a command to the robot command generator module to generate an appropriate movement for the robot.

IV. EXPERIMENTS

A. Pilot Studies

Before performing the main experiment to compare the performance between the two anticipation methods, we performed a series of pilot studies to test the setup of the system and set various parameters for the two anticipation methods. We conducted two sets of pilot studies, with a total of seven participants (three women and four men). Participants were opportunistically recruited and were compensated with a \$5 gift card for participating [3].

During the first set of pilots, a sole participant danced with the robot. Here, we sought to measure two things: how fast the robot received action messages and how accurately the robot performed with the human participant.

During the second set of pilot studies, a group of three participants danced with the robot. Here, we sought to establish appropriate parameters for the anticipation methods. To acquire these measurements, we recorded events generated from server logs as well as from odometry data from the robot. We compared the two, noting differences in velocity, distance, and event timings [3].

Results from the pilot study showed that the robot received messages from the server in a timely manner. We also analyzed the movement patterns of the robot when it coordinated its movements with the humans' and found it to be well coordinated. Based on these data, we felt confident that the robot was moving synchronously with participants and continued with the main experiment.

B. Main Experiment

We recruited a total of nine groups (27 participants in total, three persons per group) for our main experiment. Fourteen participants were women and 13 were men. Their average age was 22.93 years (s.d. = 3.98 years) and the majority were undergraduate and graduate students. Only three participants had prior dancing experience and 24 did not. Participants were recruited via mailing list and campus advertisement. Upon scheduling a timeslot, participants were randomly assigned to join a group with two others. Each participant was compensated with a \$8 gift card for their time.

After giving informed consent, participants viewed an instructional video of the choreographed dance and the experimenters explained the different movements. The participants then had time to practice the dance movements as a group as many times as they wanted. During this practice session, the robot did not dance with them.

Following the practice session, the group participated in three dance sessions. During the first session, only humans participated in the dance. During the last two sessions, the robot joined the group. In Sessions 2 and 3, the robot moved using either ECA then SIA, or SIA then ECA. (The order was counter-balanced to avoid bias.) Participants were blind as to which method was in use.

During the last two sessions, the four clients recorded depth, infrared, and skeletal data of the participants, and the server logged all event and timing data. A single camera mounted on a tripod recorded standard video of the experiment for manual analysis purposes only.

Following the experiment, participants completed a short questionnaire asking them to rate which of the two dance sessions they felt was more synchronous, a measure we have used in prior work [39]. Participants also reported which session they felt they were more synchronous with the rest of the group.

V. ROBOT DATA PREPROCESSING

The server provided the human movement data logs and the clients raw data during the experiment, as detailed in Sections II and IV-B. However, to conduct a complete comparison between the two anticipation methods, it is also necessary to determine how and when the robot actually moved during the two experimental sessions. To do this, we used the timestamped odometric data from the robot (x and y pose and angular- z orientation), as well as the server-side robot command logs.

We calculated the same events for the robot as for the humans (forward, backward, stop, and turn). Based on the changes in two consecutive x or y pose values and the robot's heading, we calculated whether the robot was moving forward or backward.

For example, when the robot faced the first Kinect sensor and moved forward, then the changes in two consecutive pose values would be positive and, if moving backward, negative. We detected turn events using changes greater than 0.4 in the z heading value of the Turtlebot's angular twist message. (Note that turn events are considered equivalent to the humans' clap events in our analysis.)

Stop events were determined when a difference less than 0.002 was detected between two consecutive poses. These stop events were classified as forward or backward depending on the heading of the robot.

After detecting all events for the robot, we manually checked the data files for any errors. During this process, we determined a 7% missing event rate. These missing events were evenly distributed across both of the anticipation methods. We manually checked the recorded video and odometric logs from the robot and determined that the robot actually moved correctly during the majority of those instances, so manually inserted the missing events into the files. There were a few instances (about 3.7% overall) when the robot did not perform the activity that it was instructed to perform, which was mostly due to network latency. We discarded those data from the analysis.

VI. DATA ANALYSIS AND RESULTS

To compare the performance and accuracy of the two anticipation methods, we first measured how synchronously the entire group, including the robot, coordinated their movements during both sessions. We then measured how appropriately timed the robot's motion was with its human counterparts.

A. Measuring Synchronization of the Group

Using the method described in [39] and discussed in Section III-A, we measured the degree of synchronization of the group for each iteration of the dance. First, we created an individual time series for each of the dancers and the robot. Events in the time series were start moving forward, stop moving forward, start moving backward, stop moving backward, and clap). Then we calculated the pairwise synchronization index for each pair using the method described in Section III-A2.

From the pairwise synchronization index we built a GTG and calculated the individual synchronization index for each human dancer, as described in Section III-A. As the humans physically stood very close in proximity, we assumed that each of the group members was influenced by all other members of the group across the course of an entire dance session. (In every iteration, participants rotated their position; therefore, a person in the front at itr_i will end up in the back by $\text{itr}_{(i+2)}$.) Thus, in the analysis, every human was connected in the graph with all other members of the group, including the robot.

When calculating the robot's individual synchronization index, we employed slightly different analyses between ECA and SIA. For ECA, because the robot's motion was based on the average of all dancers' motions in the previous iteration, when building the GTG all edges from the robot connected to all other human group members. However, for SIA, at any given itr_i the robot was only ever following $\text{MSP}(\text{itr}_{(i-1)})$ in real time. Thus,

during itr_i the robot was only influenced by that person, not by the other group members. Thus, it is logical to take only the pairwise synchronization index between the robot and that person into account while calculating the individual synchronization index of the robot and building the GTG for that iteration. Therefore, we only considered an outgoing edge from the robot to $\text{MSP}(\text{itr}_{(i-1)})$ in the GTG.

After measuring the individual synchronization index, we calculated the group synchronization index (GSI) for each group using the method described in [39]. Here we describe the method very briefly.

While calculating the GSI, both the individual synchronization index as well as the members' connectivity to the group was taken into consideration. For a given vertex in the GTG, the ratio of the number of outgoing edges connecting to it and the number of maximum possible edges in a very synchronized condition for that vertex is called the connectivity value (CV). Thus, we can define CV of series s_i as

$$\text{CV}(s_i) = \frac{\sum_{j=1, \dots, H, j \neq i} f(s_i, s_j)}{H-1}. \quad (10)$$

The CV represents how well an individual is synchronized with the rest of the group. First, we calculate each individual's synchronization index multiplied by their CV. Then, the overall GSI is computed by taking the average of this product [39]. Therefore, the overall GSI, G_τ , is computed by

$$G_\tau = \frac{\sum_{i=1, \dots, H} I_\tau(s_i) \times \text{CV}(s_i)}{H}. \quad (11)$$

While calculating the GSI we used $\tau = 0.25$ s. This value means we considered two events to be synchronous when the same types of events in two time series occurred within 0.25 s of one another.

Table I presents the GSIs for each group (three humans and one robot), across both anticipation methods (ECA and SIA) and across the four iterations per session). The table also presents the average GSI for each group in the rightmost column. Boldface is used to indicate which of the two methods yielded a higher GSI, and this is indicated for both the per-iteration GSI and the average GSI per group.

For 22 out of 36 total individual dance iterations, the SIA method yielded a higher GSI than the ECA method. And, in seven out of nine trials, the SIA method yielded a higher GSI than the ECA method.¹

Using a discrete analogue scale, we asked participants to rate on a scale from 1 to 5 how synchronous they thought the robot was with the other humans during the sessions. Based on their responses, we measured the more synchronous session of that trial, for which two out of three dancers agreed on their rating. For 7 of 9 trials, this collective rating matched with the more synchronous session of the trials determined by our method (see the last two columns in Table I).

¹Note that, due to a small sample size ($n = 36$), it would be dubious to run statistical means comparisons and one should not accept a p -value with certainty [45]. Instead, we agree with Gelman [45] that reliable patterns can be found by averaging, as reported here.

TABLE I
GSIS FOR ALL GROUPS

Group No.	Iteration No.	GSI		Mean GSI	
		ECA	SIA	ECA	SIA
1	1	0.39	0.40	0.29	0.45
	2	0.30	0.26		
	3	0.15	0.63		
	4	0.33	0.52		
2	1	0.28	0.66	0.46	0.56
	2	0.39	0.37		
	3	0.63	0.51		
	4	0.54	0.71		
3	1	0.37	0.40	0.45	0.52
	2	0.48	0.52		
	3	0.59	0.62		
	4	0.37	0.55		
4	1	0.30	0.20	0.33	0.29
	2	0.37	0.38		
	3	0.37	0.30		
	4	0.28	0.28		
5	1	0.31	0.24	0.38	0.43
	2	0.39	0.44		
	3	0.52	0.59		
	4	0.30	0.46		
6	1	0.50	0.43	0.41	0.41
	2	0.41	0.45		
	3	0.40	0.32		
	4	0.33	0.46		
7	1	0.41	0.25	0.42	0.47
	2	0.42	0.29		
	3	0.35	0.66		
	4	0.52	0.68		
8	1	0.78	0.32	0.50	0.52
	2	0.40	0.48		
	3	0.41	0.72		
	4	0.41	0.56		
9	1	0.56	0.35	0.45	0.44
	2	0.35	0.34		
	3	0.42	0.46		
	4	0.46	0.59		

Each group includes three people and one robot.

B. Measuring Robot Timing Appropriateness

For both anticipation methods, we aimed to have the robot perform its actions (events) as closely as possible in time to its human counterparts. Thus, we measured how close the robot's actual movement was to what the humans were doing at that time.

Thus, as a measure of timing appropriateness (TA) of the robot, we calculated the absolute time difference between the time when the robot performed an event and the ideal timing of that event. As a measure of the ideal timing of an event, we took the average timing of an event performed by the humans. This measure is similar to the absolute offset measure used in [46]; however, the TA measure used here is within the context of a group.

First, we represented all events associated with the humans during an iteration by a time series. Then, we clustered all the similar types of events together with those that were performed by the dancers within a time threshold, ϵ . For example, for




Fig. 5. TA calculation for the robot's movement.




Fig. 6. TA measure for the robot with 95% confidence interval for both methods, SIA and ECA.

a single event e , we calculated the timing of the event performed by three human participants within ϵ , i.e., $t(\text{dancer}_1, e)$, $t(\text{dancer}_2, e)$, and $t(\text{dancer}_3, e)$. We also calculated the timing of that event performed by the robot, $t(\text{robot}, e)$. Then, to calculate the ideal timing for the robot, we take the average of these times of this event performed by the humans. Thus, $t(\text{robot}_{\text{ideal}}, e) = (t(\text{dancer}_1, e) + t(\text{dancer}_2, e) + t(\text{dancer}_3, e))/3$. Then we calculated the TA of that event performed by the robot as, $\text{TA}(e) = |t(\text{robot}, e) - t(\text{robot}_{\text{ideal}}, e)|$. Fig. 5 presents an example calculation of TA for event e .

After calculating TA for each event during all the trials, we created two histograms, one for each anticipation method. We used a bin size of 0.1 s, starting at 0 s and going to 2.5 s. Then we calculated the frequency of the events for which the TA falls within that time span.

In Figs. 7(a) and (b), we present histograms representing the TA measure and the cumulative percentage of event frequencies, for the ECA and SIA methods, respectively. Fig. 7(a) (ECA) shows that the robot was able to perform 81.88% of its events within 1.2 s and 90% of its events within 1.6 s of the appropriate timing. Fig. 7(b) (SIA) shows that the robot performed 81.65% of the events within 0.8 s and 90.82% of the events within 1.2 s of the appropriate timing.

Fig. 7(c) presents the cumulative percentage of events for both methods together. One can find that the robot performed the events more appropriately during the SIA method than compared with the ECA method.




Fig. 7. Event frequency distribution and the cumulative percentage distribution of the TA measure for the two anticipation methods: SIA (left) and ECA (right). The rightmost graph shows the distribution of the TA measure for both methods.

For the SIA method, the mean for the TA measure was 0.54 s (s.d. = 0.59 s) (see Fig. 6). For the ECA method, the mean TA measure was 0.70 s (s.d. = 0.50 s) (see Fig. 6). While these data did not have a normal distribution, as is visible from the graph and a normality test, they did have a sufficient number of means to compare statistically. We conducted a Wilcoxon Signed Rank Test and found that the TA values for the ECA method were significantly larger than for the SIA method, $z = -4.399$, $p < 0.05$, $r = -0.18$. This means that when using the SIA method the robot moved more appropriately in time than when using the ECA method.

VII. DISCUSSION AND FUTURE WORK

The results suggest that the human–robot team was more synchronous using SIA than using the ECA method. Moreover, when SIA was used, the robot was able to perform its actions significantly closer to the appropriate timing of the event. This supports the idea that SIA is well suited to provide movement coordination information to a robot during an SJA scenario.

Additionally, these results might support the robustness of the SIA method over the ECA method, as the SIA method is more dynamic and adaptable to change within the group. In our study, the SIA method chose the most synchronous dancer in the group and used that to inform the robot’s actions in real time. However, relying on a method like ECA would mean that if a dancer was moving asynchronously within the group, the robot’s motion could be adversely affected (as it is following everyone). SIA is robust to handle this phenomenon, as a person who performed asynchronous movements within the group is unlikely to ever be chosen as the most synchronous person.

This study shows that taking team dynamics into account can be useful for robots when conducting coordinated activities with teammates. This study can lead others in the robotics community toward further investigating the role of a group on behavior, rather than just focusing on individuals. This has implications not only for human–robot interaction, but also for multirobot systems research. We are currently exploring the effect of different anticipation methods in multihuman multirobot scenarios [47], [48].

One limitation of this study is how event detection is calculated. In the current setup, a predefined set of human activities were detected by the system to understand the group dynamics. Building on this foundation, our future work will include in-

corporating human gross motion directly to the synchronization measurement step, instead of using pre-labeled events. Moreover, we are also planning to incorporate a decision module for robots, which will use the perceived knowledge to select the best decision from a set of options, based on the context [19], [49].

Another limitation of the current method is how it uses team metrics and task-related information. For example, the method does not yet incorporate dancer expertise, nor does it factor in the tempo or dynamics of the music. In the future, we plan to incorporate an understanding of these factors. For example, in a team of novice dancers, a robot could perhaps keep a team on tempo.

In the future, we also seek to explore the use of robot-centric vision and local coordination methods to calculate synchrony. This will enable robots to operate in more dynamic settings and lessen the need for external sensors. However, incorporating local sensor data will be more challenging as it might be more noisy due to occlusion and local movements. However, we will build on our prior multimodal fusion and others’ robot-centric perception work to overcome this challenge [19], [50].

We also will explore incorporating other synchronization methods humans employ, such as adapting to continuous tempo changes, within the SIA algorithm. Models like the adaptation and anticipation model have been proposed in the literature to computationally model this behavior in humans by combining adaptation and anticipation during an activity [35], [51]. It may be beneficial for a robot to have this ability both in human–robot and multirobot teams. This integration might make the SIA algorithm more robust in anticipating and synthesizing future activities more accurately.

We also hope to extend our methods to work beyond SJA activities, such as timed but varied collaborative tasks within industrial settings. A human–robot team working in an industrial setting has specific sequences of activities to perform overtime, some of which might be independent and might not have to happen synchronously. However, the events do have to happen contingently; therefore, some of our anticipatory methods may be applicable.

Movement coordination is an important, emerging research area in robotics, neuroscience, biology, and many other fields [13]–[17], [52]. Our work helps to enable robots to have a better understanding of how to coordinate with the environment. This can be useful both for solving problems in robotics and, perhaps, also in fields beyond.

REFERENCES

- [1] R. Wilcox, S. Nikolaidis, and J. A. Shah, "Optimization of temporal dynamics for adaptive human-robot interaction in assembly manufacturing," presented at the *Robot., Sci. Syst. Conf.*, Sydney, Australia, 2012.
- [2] L. D. Riek, "The social co-robotics problem space: Six key challenges," in *Proc. RSS, Robot. Challenges Visions Workshop*, 2013.
- [3] S. Rack, T. Iqbal, and L. Riek, "Enabling synchronous joint action in human-robot teams," in *Proc. 10th Annu. ACM/IEEE Int. Conf. Human-Robot Interaction Extended Abstr.*, 2015, pp. 153–154.
- [4] R. Ros, I. Baroni, and Y. Demiris, "Adaptive human-robot interaction in sensorimotor task instruction: From human to robot dance tutors," *Robot. Auton. Syst.*, vol. 62, pp. 707–720, 2014.
- [5] J. Fasola and M. Mataric, "A socially assistive robot exercise coach for the elderly," *J. Hum.-Robot Interact.*, vol. 2, pp. 3–32, 2013.
- [6] M. P. Michalowski, S. Sabanovic, and H. Kozima, "A dancing robot for rhythmic social interaction," in *Proc. ACM/IEEE Int. Conf. Human-Robot Interaction*, 2007, pp. 89–96.
- [7] H. W. Park and A. M. Howard, "Understanding a child's play for robot interaction by sequencing play primitives using hidden Markov models," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2010, pp. 170–177.
- [8] J. Koenemann, F. Burget, and M. Bennewitz, "Real-time imitation of human whole-body motions by humanoid," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2014, pp. 2806–2812.
- [9] L. D. Riek, T. C. Rabinowitch, P. Bremner, A. G. Pipe, M. Fraser, and P. Robinson, "Cooperative gestures: Effective signaling for humanoid robots," in *Proc. 5th ACM/IEEE Int. Conf. Human-Robot Interaction*, 2010, pp. 61–68.
- [10] T. Iqbal, M. J. Gonzales, and L. D. Riek, "Joint action perception to enable fluent human-robot teamwork," in *Proc. 24th IEEE Int. Symp. Robot Human Interactive Commun.*, 2015, pp. 400–406.
- [11] G. Hoffman and C. Breazeal, "Cost-based anticipatory action selection for human-robot fluency," *IEEE Trans. Robot.*, vol. 23, no. 5, pp. 952–961, Oct. 2007.
- [12] M. Cakmak, S. S. Srinivasa, M. K. Lee, S. Kiesler, and J. Forlizzi, "Using spatial and temporal contrast for fluent robot-human hand-overs," in *Proc. 6th Int. Conf. Human-Robot Interaction*, 2011, pp. 489–496.
- [13] J. Bandouch, O. C. Jenkins, and M. Beetz, "A self-training approach for visual tracking and recognition of complex human activity patterns," *Int. J. Comput. Vis.*, vol. 99, pp. 166–189, 2012.
- [14] H. Zhang and L. E. Parker, "4-dimensional local spatio-temporal features for human activity recognition," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2011, pp. 2044–2049.
- [15] H. Koppula and A. Saxena, "Learning spatio-temporal structure from RGB-D videos for human activity detection and anticipation," in *Proc. 30th Int. Conf. Machine Learning*, 2013, pp. 792–800.
- [16] F. Ofli, R. Chaudhry, G. Kurillo, R. Vidal, and R. Bajcsy, "Berkeley MHAD: A comprehensive multimodal human action database," in *Proc. IEEE Work. Appl. Comput. Vision*, 2013, pp. 53–60.
- [17] A. A. Samadani, A. Ghodsi, and D. Kulić, "Discriminative functional analysis of human movements," *Pattern Recog. Lett.*, vol. 34, pp. 1829–1839, 2013.
- [18] W. Burgard *et al.*, "Experiences with an interactive museum tour-guide robot," *Artif. Intell.*, vol. 114, pp. 3–55, 1999.
- [19] A. Nigam and L. Riek, "Social context perception for mobile robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2015, pp. 3621–3627.
- [20] M. Nagy, Z. Ákos, D. Biro, and T. Vicsek, "Hierarchical group dynamics in pigeon flocks," *Nature*, vol. 464, pp. 890–893, 2010.
- [21] M. Nagy, G. Vásárhelyi, B. Pettit, I. Roberts-Mariani, T. Vicsek, and D. Biro, "Context-dependent hierarchies in pigeons," *Proc. Nat. Acad. Sci.*, vol. 110, pp. 13049–13054, 2013.
- [22] N. E. Leonard *et al.*, "In the dance studio: Analysis of human flocking," in *Proc. Am. Control Conf.*, 2012, pp. 4333–4338.
- [23] N. E. Leonard *et al.*, "In the dance studio: An art and engineering exploration of human flocking," in *Controls and Art*. Switzerland: Springer, 2014.
- [24] M. Clayton, R. Sager, and U. Will, "In time with the music: The concept of entrainment and its significance for ethnomusicology," in *Proc. Eur. Meetings Ethnomusicol.*, 2005, pp. 1–82.
- [25] R. C. Schmidt and M. J. Richardson, "Dynamics of interpersonal coordination," in *Coordination: Neural, Behavioral and Social Dynamics*. Berlin, Germany: Springer, 2008.
- [26] E. Pereda, R. Q. Quiroga, and J. Bhattacharya, "Nonlinear multivariate analysis of neurophysiological signals," *Prog. Neurobiol.*, vol. 77, pp. 1–37, 2005.
- [27] S. Khalfa *et al.*, "Role of tempo entrainment in psychophysiological differentiation of happy and sad music?" *J. Psychophysiol.*, vol. 68, pp. 17–26, 2008.
- [28] T. Kreuz *et al.*, "Measuring synchronization in coupled model systems: A comparison of different approaches," *Phys. D, Nonlinear Phenomena*, vol. 225, pp. 29–42, 2007.
- [29] T. Himberg and M. R. Thompson, "Learning and synchronising dance movements in South African songs: Cross-cultural motion-capture study," *Dance Res.*, vol. 29, pp. 305–328, 2011.
- [30] D. Lakens, "Movement synchrony and perceived entitativity," *J. Exp. Soc. Psychol.*, vol. 46, pp. 701–708, 2010.
- [31] H. Hennig, "Synchronization in human musical rhythms and mutually interacting complex systems," *Proc. Nat. Acad. Sci.*, vol. 111, pp. 12974–12979, 2014.
- [32] N. Sebanz, H. Bekkering, and G. Knoblich, "Joint action: Bodies and minds moving together," *Trends Cogn. Sci.*, vol. 10, pp. 70–76, 2006.
- [33] M. J. Richardson, R. L. Garcia, T. D. Frank, M. Gergor, and K. L. Marsh, "Measuring group synchrony: A cluster-phase method for analyzing multivariate movement time-series," *Frontiers Physiol.*, vol. 19, p. 405, 2012.
- [34] G. Varni, G. Volpe, and A. Camurri, "A system for real-time multimodal analysis of nonverbal affective social interaction in user-centric media," *IEEE Trans. Multimedia*, vol. 12, no. 6, pp. 576–590, Oct. 2010.
- [35] M. M. van der Steen, N. Jacoby, M. T. Fairhurst, and P. E. Keller, "Sensorimotor synchronization with tempo-changing auditory sequences: Modeling temporal adaptation and anticipation," *Brain Res.*, vol. 1626, pp. 66–87, 2015.
- [36] N. Marwan, M. Carmenromano, M. Thiel, and J. Kurths, "Recurrence plots for the analysis of complex systems," *Phys. Rep.*, vol. 438, pp. 237–329, 2007.
- [37] I. Konvalinka *et al.*, "Follow you, follow me: Continuous mutual prediction and adaptation in joint tapping," *J. Exp. Psychol.*, vol. 63, pp. 2220–2230, 2010.
- [38] R. Q. Quiroga, T. Kreuz, and P. Grassberger, "Event synchronization: A simple and fast method to measure synchronicity and time delay patterns," *Phys. Rev. E*, vol. 66, p. 041904, 2002.
- [39] T. Iqbal and L. D. Riek, "A method for automatic detection of psychomotor entrainment," *IEEE Trans. Affective Comput.*, vol. 7, no. 1, pp. 3–16, Jan.–Mar. 2015.
- [40] K. P. Hawkins, S. Bansal, N. N. Vo, and A. F. Bobick, "Anticipating human actions for collaboration in the presence of task and sensor uncertainty," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2014, pp. 2215–2222.
- [41] G. Hoffman and C. Breazeal, "Cost-based anticipatory action selection for human-robot fluency," *IEEE Trans. Robot.*, vol. 23, no. 5, pp. 952–961, Oct. 2007.
- [42] C. Pérez-DArpino and J. Shah, "Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2015, pp. 6175–6182.
- [43] V. V. Unhelkar, C. Pérez-DArpino, L. Stirling, and J. Shah, "Human-robot co-navigation using anticipatory indicators of human walking motion," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2015, pp. 6183–6190.
- [44] T. Iqbal, M. J. Gonzales, and L. D. Riek, "A model for time-synchronized sensing and motion to support human-robot fluency," in *Proc. ACM/IEEE HRI, Workshop Timing HRI*, 2014.
- [45] A. Gelman, "Commentary: P values and statistical practice," *Epidemiology*, vol. 24, pp. 69–72, 2013.
- [46] G. Hoffman and G. Weinberg, "Synchronization in human-robot musicianship," in *19th Int. Symp. Robot Human Interactive Commun.*, 2010, pp. 718–724.
- [47] T. Iqbal and L. D. Riek, "Detecting and synthesizing synchronous joint action in human-robot teams," in *Proc. ACM Int. Conf. Multimodal Interaction*, 2015, pp. 581–585.
- [48] T. Iqbal and L. D. Riek, "Coordination dynamics in multi-human multi-robot teams," 2016, under review.
- [49] M. O'Connor and L. Riek, "Detecting social context: A method for social event classification using naturalistic multimodal data," in *Proc. IEEE Int. Conf. Workshops Autom. Face Gesture Recognit.*, 2015, pp. 1–7.
- [50] M. S. Ryoo and L. Matthies, "First-person activity recognition: What are they doing to me?" in *IEEE Conf. Comput. Vision Pattern Recog.*, 2013, pp. 2730–2737.
- [51] M. C. M. van der Steen and P. E. Keller, "The ADaptation and Anticipation Model (ADAM) of sensorimotor synchronization," *Front. Hum. Neurosci.*, vol. 7, p. 253, 2013.
- [52] A. Gelblum, I. Pinkoviezky, E. Fonio, A. Ghosh, N. Gov, and O. Feinerman, "Ant groups optimally amplify the effect of transiently informed individuals," *Nature Commun.*, vol. 6, p. 7729, 2015.



Tariq Iqbal (S'14) received the B.S. degree in computer science and engineering from Bangladesh University of Engineering and Technology, Dhaka, Bangladesh, and the M.S. degree in computer science from University of Texas at El Paso, El Paso, TX, USA. He is currently working toward the Ph.D. degree with the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN, USA.

His research interests include robotics, computer vision, machine learning, and social signal processing.



Samantha Rack (S'14) is currently working toward the B.S. degree in computer science with University of Notre Dame, Notre Dame, IN, USA.

She will begin working in industry following graduation. Her research interests include human-computer interaction and social signal processing.



Laurel D. Riek (M'00–SM'14) received the B.S. degree in logic and computation from Carnegie Mellon University, Pittsburgh, PA, USA, and the Ph.D. degree in computer science from University of Cambridge, Cambridge, U.K.

She is the Clare Boothe Luce Assistant Professor of computer science and engineering with University of Notre Dame, Notre Dame, IN, USA, and directs of the Robotics, Health, and Communication Laboratory. Her research interests include robotics, social signal processing, human behavior understanding, and healthcare informatics. She is the author of more than 75 technical articles in these areas.

Dr. Riek serves on the Editorial Board of IEEE TRANSACTIONS ON HUMAN MACHINE SYSTEMS, the Steering Committee of the ACM/IEEE Conference on Human-Robot Interaction, and numerous conference program committees.