Physical Collaboration of Human-Human and Human-Robot Teams

Kyle B. Reed, *Member*, *IEEE*, and Michael A. Peshkin, *Member*, *IEEE*

Abstract—Human partners working on a target acquisition task perform faster than do individuals on the same task, even though the partners consider each other to be an impediment. We recorded the force profile of each partner during the task, revealing an emergent specialization of roles that could only have been negotiated through a haptic channel. With this understanding of human haptic communication, we attempted a "Haptic Turing Test," replicating human behaviors in a robot partner. Human participants consciously and incorrectly believed their partner was human. However, force profiles did not show specialization of roles in the human partner, nor enhanced dyadic performance, suggesting that haptic interaction holds a greater subconscious subtlety. We further report observations of a nonzero dyadic steady-state force perhaps analogous to cocontraction within the limb of an individual, where it contributes to limb stiffness and disturbance rejection. We present results on disturbance rejection in a dyad, showing lack of an effective dyadic strategy for brief events.

Index Terms—Human-human, human-machine, human-robot interaction, physical cooperation, collaboration, pHRI.

1 Introduction

Humans have long worked together in pairs or groups and, presumably, can communicate task-relevant information via physical interaction. Klingspor et al. [1] suggest that human-robot physical communication in a shared task should be designed to follow implicit human-human communication standards. However, little is known about what the implicit physical communication standards in humans might be. As human-machine interfaces become more common, systems that are more fluent in the language of physical communication with people should be more intuitive to use, and require less training or adaptation on the part of their users. Sebanz et al. [2], in a review of coordinated interaction, say that it may not be possible to fully understand how humans operate by solely studying people working in isolation. Understanding the haptic interactions between two people should elucidate how an individual works alone as well as lead to an increased understanding of how a human can intuitively and cooperatively work with a robot on physical tasks. This report presents three motion and disturbance rejection experiments investigating the interactions of two people physically cooperating. How do the resources of two people combine to complete a task?

2 Background

Much of human-human interaction is mediated by vision. When performing a task within sight of another person, for instance swinging a leg while seated, two people will tend to synchronize their actions [3]. While watching another person, mirror neurons in the brain develop a representation of how to perform that same action [4], [5]. Although human-human communication is a large research endeavor, most studies focus on visual and auditory channels of communication. Very few have studied communication through the physical exchange of force and motion.

Physical interaction between humans relies upon the cooperative force and past experience of participants, thus understanding the motion of individuals acting alone is a necessary, but not sufficient, aspect of understanding group actions. From birth, humans begin building up a base of experience for interacting with objects in the world through repeated interactions [6]. Internal models are constantly updated, so, for instance, humans are able to anticipate the weight of an object based upon its size [7] and can produce a precomputed force that will lift the object along a desired path [8]. The learned trajectories of human hands tend to be roughly straight [9] and smooth [10] with low jerk. Smooth motion, quick adaptation, and past knowledge allow humans to manipulate known and unknown objects proficiently.

Humans are also able to adapt to perturbations from the environment, for instance from an external force or another person. One method of resisting external perturbations is to cocontract opposing muscles on the same joint, thus increasing limb stiffness. Cocontraction is a common strategy when interacting with unstable force fields [11], [12].

Models have been developed that are able to predict motion and performance quite well, for instance animating human athletes [13], relating performance to the task [14], and recognizing human intent [15], [16], but there are many areas that are not yet well understood. For our area of interest, it is unclear how the forces and motions of two people combine. Do the models of performance and motion that well describe individual motions extend to dyads? Do dyads respond to perturbation forces similarly to the way an individual does or do dyads develop a new strategy?

Gentry [17] studied how two people work together while dancing. The two partners communicate through their hands while the rhythm of the music aids in synchronizing

- *K.B. Reed is with the Laboratory for Computational Sensing and Robotics, Johns Hopkins University, 3400 N. Charles Street, CSEB 113, Baltimore, MD 21218. E-mail: reedkb@alumni.northwestern.edu.*
- *M.A. Peshkin is with the Department of Mechanical Engineering, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208-3111. E-mail: peshkin@northwestern.edu.*

Manuscript received 3 Jan. 2008; revised 6 June 2008; accepted 21 Aug. 2008; published online 4 Sept. 2008.
Recommended for acceptance by L. Jones.
For information on obtaining reprints of this article, please send e-mail to: toh@computer.org, and reference IEEECS Log Number TH-2008-01-0001.
Digital Object Identifier no. 10.1109/ToH.2008.13.

1939-1412/08/$25.00 © 2008 IEEE Published by the IEEE CS, RAS, & CES

their motions. Gentry found that couples could perform equally well while blindfolded, which indicates that a significant amount of communication is through their hands and not through visual channels. Wegner and Zeaman [18] found that groups could track a path better than individuals could, yet, after studying social facilitation, leadership, guidance, and skill learning/transfer, they could not reach a satisfying conclusion about why the performance improved. Neither of these studies examined the forces between the two participants, which we will show is a vital part of communication.

Airplane flight sticks that mechanically couple the pilot, copilot, and the plane's control surfaces provide haptic feedback to each pilot. In a survey of commercial pilots, Field and Harris [19] found that haptic interaction between the two pilots and also between the pilot and the plane was an important and useful channel of communication. The pilots noted the importance of using this haptic feedback to determine if the plane was working properly when flying on autopilot. Field and Harris also found that the desirable communication was lost when switching to Fly-By-Wire (FBW) designs. FBW systems eliminate the direct mechanical connection between the pilot and the plane's control surfaces and also between the two pilots. In a comparison between direct mechanical connections and nonhaptic FBW, Summers et al. [20] found a performance decrement when using nonhaptic FBW.

When two people cooperate on a physical task, there are multiple ways to combine the redundant controls that are available [21], [22]. In some FBW airplane designs, the command given to the airplane control surfaces is an electronic average of the position from the two pilots' flight controls [20]. Glynn and Henning [23] conducted an experiment where the motion of a mass was controlled by averaging two participants' force inputs. They found that haptic feedback decreased errors and improved performance when visual and haptic feedbacks were synchronous. Averaging the input command is a simple strategy but not necessarily the best combination since each individual's motion will be diluted. Imagine the effect if one pilot attempts to avoid an obstacle by turning to the left while the other to the right: the average effect is straight into the obstacle. A better solution to the redundant control problem likely consists of exploiting the redundant abilities of the dyad. Knoblich and Jordon [24] suggest that groups should be able to perform better than individuals since each person in a group can focus on a subset of the actions.

Shergill et al. [25] performed an experiment examining the forces applied between two people without motion. Each participant was told to push against a transducer at a force level matching that of the force just previously applied to them. Each participant was unaware of the instructions given to the other participant. Taking turns, each participant applied a force, which was then applied to the other participant. Shergill et al. found that in every case, the forces escalated from trial to trial. In a second set of experiments, Shergill et al. applied a brief constant force to the tip of a participant's finger and asked them to generate the same force with another finger on the other hand. The participants consistently generated too much force. Shergill et al. suggested that in both experiments, self-generated forces were perceived as weaker than externally generated forces, which can affect how one person perceives his/her partner and how two people cooperatively work on a task.

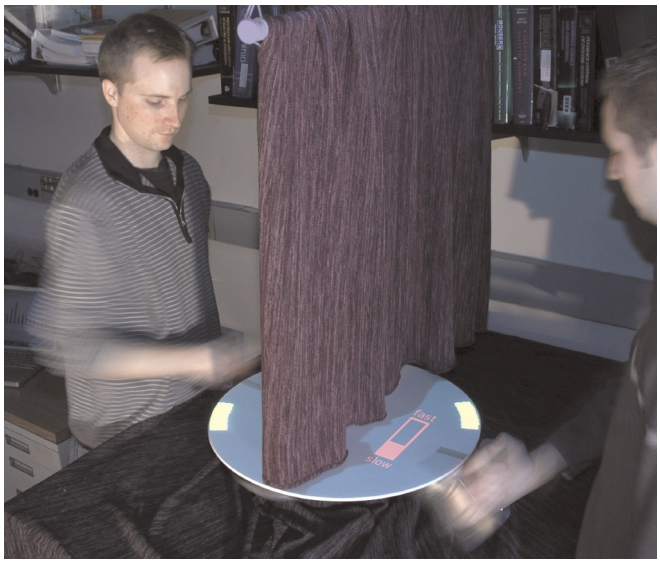

Fig. 1. **Experimental setup with a two-handled crank**. Two participants cooperate to complete a timed target acquisition task. They cannot see each other and are instructed not to speak, thus communication is limited to forces and motions transmitted through the rigid handles of the device. Picture digitally modified for clarity.

Telemanipulators allow a user to interact with a person or object over large distances. Interacting via teleoperation can lack fidelity if the device does not reproduce the other person's forces perfectly or introduces latencies [26], [27]. Much of the teleoperation research focuses on issues of how to recreate forces and less on the actual interaction between the two remote agents, thus we will not go into detail here.

The experiments presented in this report focus on the direct interaction of a person physically cooperating with another person and/or robot. The first experiment discusses the way in which the participants adopted unexpected cooperative strategies when working together. The second and third experiments attempted to replicate these cooperative strategies and the roles that commonly develop. The second experiment involved replacing one person with a robot programmed to perform one of the two specialized roles in order to test our understanding of the human behavior. The third experiment introduced perturbations during the task to examine an antagonistic force exerted by each member of a dyad.

3 Experiment One: Human-Human

In our investigation of dyadic physical communication, we devised the simplest task we could conceive: mutually acquiring a one DOF visual target. Fig. 1 illustrates the experimental setup. The two handles are connected via a rigid crank that is free to turn about its center, thus assuring a high fidelity mechanical haptic channel between the participants. A direct drive motor is attached under the table and is unknown to the participants. The motor was not used in the experiments we first describe. The angular

position of the crank is displayed to each participant as a bold black mark on the top disk, which moves with the crank. A projector above, aimed at the disk, displays a motivating performance measure, separate instructions to each participant, and a target in the same relative position for each participant. A curtain hangs between the participants to prevent visual communication and the participants are asked not to speak to each other.

3.1 Participants

Thirty students (10 male; two left-handed), age 18-24, from Northwestern University's Psychology participant pool participated with informed consent. One dyad was male-male, eight were female-male, and six were female-female.

3.2 Experimental Procedure

Two randomly selected participants stood on opposite sides of the crank. Participants were instructed to move the handle into the target as quickly as possible and to hold it there until a new target appeared (at a randomly selected time between 700 and 1,700 ms). The target changed color when the handle was inside the target. Each target subtended 6 degrees of the 50.5 cm diameter disk (2.6 cm at the perimeter of the disk) with a distance between consecutive targets of 70 $\pm$ 10 degrees (30.9 $\pm$ 4.4 cm). Five-sixths of the trials required a reversal of handle rotation from the previous trial; in one-sixth of the trials ("catch trials"), handle motion continued in the same direction. We discarded catch trials and the trials immediately following a catch trial. We included the variation of the target position and direction as well as the delay before a new target appeared to prevent participants from adapting to a predictable pattern. The unpredictable pattern minimized learning effects, thus allowing us to examine differences in their interaction.

The projector also displayed task information to each participant. Each participant could see a performance measure when they were performing the task. The performance measure encouraged participants to actively engage the task by showing a filled region proportional to their performance time, updated after each trial. The participants were told that this would display their performance time. The measure was not visible to participants when their partner performed alone. Additionally, to reduce the effects of experimenter bias, the projector displayed messages informing participants which block of trials they were to participate on.

An experimental run started with one of the individuals, or the dyad together, performing a block of 120 trials (target acquisitions). Half of the participant pairs completed one block of trials individually first, and then one block as a dyad (A, B, AB). The other half performed as the dyad first (AB, A, B). The sequence was performed twice (e.g., A, B, AB, A, B, AB). Presentation order made no significant difference. Each block of 120 trials included 100 opposite-direction trials and 20 same-direction catch trials. The experimental apparatus was identical when the participants were working as individuals and dyads, except that the small rotational inertia of the crank $(0.113\,\mathrm{kg}\cdot\mathrm{m}^2)$ was physically doubled in the dyad condition. Doubling the inertia normalized the results across individuals and dyads. If each individual applied the same forces alone as in a dyad, the performance would be identical in both cases. The entire experiment took less than 30 minutes, for a total of 720 nondiscarded trials (480 for each participant). Individual force and common motion were recorded at 1 kHz.

3.3 Performance

We reported results on the speed of task execution in Psychological Science [28] and will summarize only briefly here. We found that dyads completed the task faster than either one of the members could complete it when working individually. On average, dyads were 54.5 ms faster (8.5 percent improvement) than their constituent members working alone $(t(15) = 5.95, p < 0.01)$. The average completion time for dyads was 641.2 ms. The improved performance developed quickly (within 20 trials) once the dyad began working together. Many participants reported a perception of interference from the other participant; few reported cooperation. The explanation cannot be sharing the load because we doubled the crank's rotational inertia in the dyad condition. Rather, we found that the symmetry of the task was spontaneously broken, with the participants taking on different roles.

3.4 Social Facilitation

It has been suggested that Social Facilitation could explain why two people are faster than one on the same task. Social Facilitation research has a long history [29], [30], [31] showing that merely having someone visually present and watching causes a person to perform well on tasks that they are proficient in. However, Social Facilitation cannot explain the results from our experiment because the experimenter was present in all cases. Also, each participant was constantly aware that their performance would be evident to the other participant, since the handle was visibly moving, regardless of whether the trials were individual or dyadic.

3.5 Specialized Forces

To analyze the forces present during a dyadic target acquisition task, it is convenient to transform the forces from members A (F_A) and B (F_B) into a "net force" and a "difference force." The net force $(F_{\mathrm{net}} = F_A + F_B)$ is the sum of the members' forces and is the task-relevant force that accelerates the crank. The difference force $(F_{\mathrm{diff}} = F_A - F_B)$ is a measure of the disagreement of the members and has no physical effect on crank motion. The difference force is a measure of the exerted force that does not go into accelerating the crank. As will be discussed below, the difference force is a possible channel of communication between the participants.

We identified some qualitatively different modes of task sharing in our participant pool. An "active/inert" dyad is shown in Fig. 2a. Early in a trial, member A provides a force toward the target while member B's passive inertia creates a counterproductive force. Late in the trial, member A provides the deceleration force while member B is again passive.

Fig. 2b shows a very different pattern, which we denote "specialized." Member C pushes toward the target early in the trial to accelerate the crank while member D either passively or actively resists the motion. Member D then

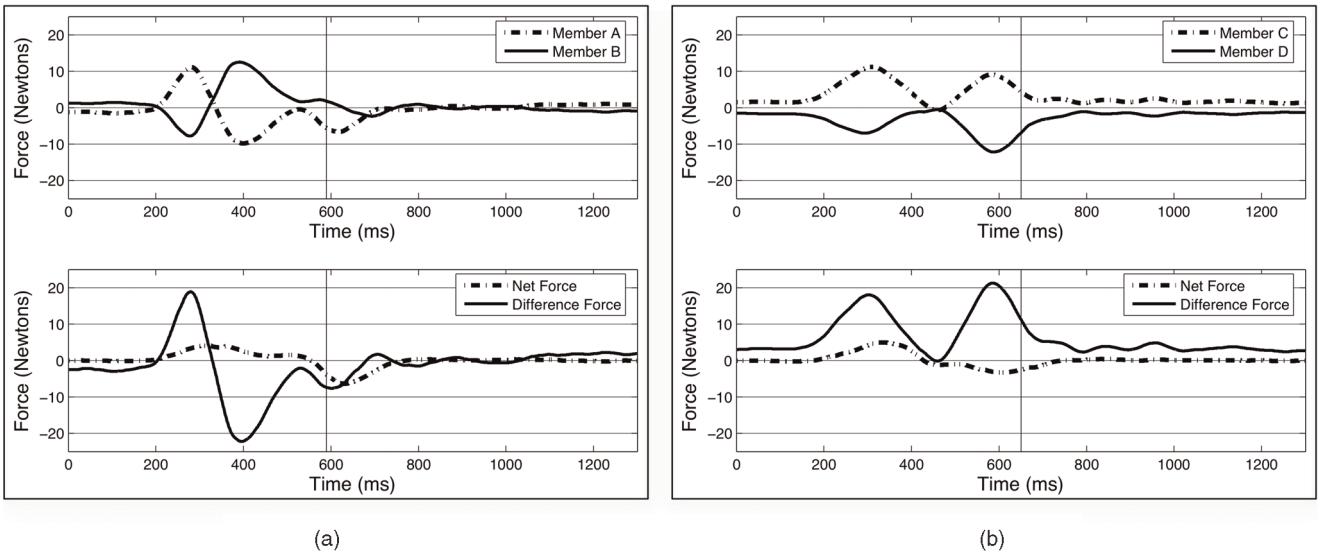

Fig. 2. **Dyad members' force profiles**. Solely through haptic interaction, most dyads developed a new strategy to complete this task by specializing their force production. These graphs show the forces in a single trial for (a) an "active/inert" dyad and (b) a "specialized" dyad. The upper graphs show the forces produced by each member of the dyad when working together. Forces are recorded as positive when applied toward the target. The lower graphs show the sum and difference of their forces: the net force (the sum) contributes to crank motion while the difference force is expended in opposition to one another. In the specialized dyad, the difference force is always the same sign: member C accelerates the crank forward while D is pulled along, and then C continues to push forward (fails to reverse) while D brakes. The vertical line in each plot represents the completion time.

pushes away from the target late in the trial to decelerate the crank while member C continues pushing toward the target. Member C primarily contributes during the acceleration phase and member D primarily contributes during the deceleration phase. This specialization is clearly revealed by inspecting the difference force, which remains always the same sign.

We defined some relevant quantities to help measure and test the suspected specialization. The acceleration phase is defined to begin when the net force is first larger than 0.5 N and ends when the net force changes sign. The deceleration phase begins at that moment and ends when the crank is in the target and the absolute value of the net force drops to and stays under 0.5 N. In Fig. 2b, the acceleration phase lasts from 190 to 430 ms. The deceleration phase lasts from 430 to 740 ms. All times are measured from when the target first appears.

To find the contribution of each participant, we integrated the force of each member over each phase and divided by the integrated net force for that phase as shown by

$$C_{\text{member,phase}} = \frac{\int F_{\text{member,phase}}}{\int F_{\text{net,phase}}}, \tag{1}$$

where $phase$ is either acceleration or deceleration, $member$ is either individual, C is the contribution of each member during each phase, F is the force of each member during each phase, and F_{net} is the net force.

The result provides four fractional contributions: each member's contribution during the acceleration and deceleration phases. The two members' contributions for each phase necessarily sum to one. A negative contribution indicates that the member was accelerating during a deceleration phase, or decelerating during an acceleration phase, even if only due to passive inertia. A contribution greater than one indicates that a member had to compensate for the negative contribution of the other dyad member.

Specialization can occur in terms of phase (acceleration versus deceleration) or direction (left versus right). Left/right (L/R) specialization would be seen as one member predominantly contributing to the left going forces while the other member would predominantly contribute to the right going forces. The difference between phase and direction specialization is not identifiable in a single trial; Fig. 2b could show either type. The trend is apparent when we examine multiple trials.

Fig. 3 shows four clusters of points, each point representing one trial. Clusters S_1 and S_2 are for the two members of an acceleration/deceleration (A/D) specialized dyad and clusters N_1 and N_2 are for the two members of a nonspecialized dyad. The nonspecialized dyad is an example of an active/inert dyad since one participant has a contribution near 1 for both phases of motion. For every point in one of the clusters, there is a point mirrored around (0.5, 0.5), which represents the other member of the dyad. As an example, a dot at (0, 1) represents a trial in which member S_1 contributed 100 percent of the acceleration and 0 percent of the deceleration. In this case, the other member, S_2, necessarily contributed 0 percent of the acceleration and 100 percent of the deceleration.

Dots near the $x = y$ line represent one member's equal contribution during the different phases. The perpendicular distance of a dot from the $x = y$ line is a measure of the degree of specialization. The ellipses show the standard deviation of each cluster of trials along the $x = y$ and $x = 1 - y$ lines. Using these measures, a nonspecialized dyad is distinctly different than the specialized dyad.

To assess the significance of the specialization type, we compared A/D and L/R specialization. L/R specialization

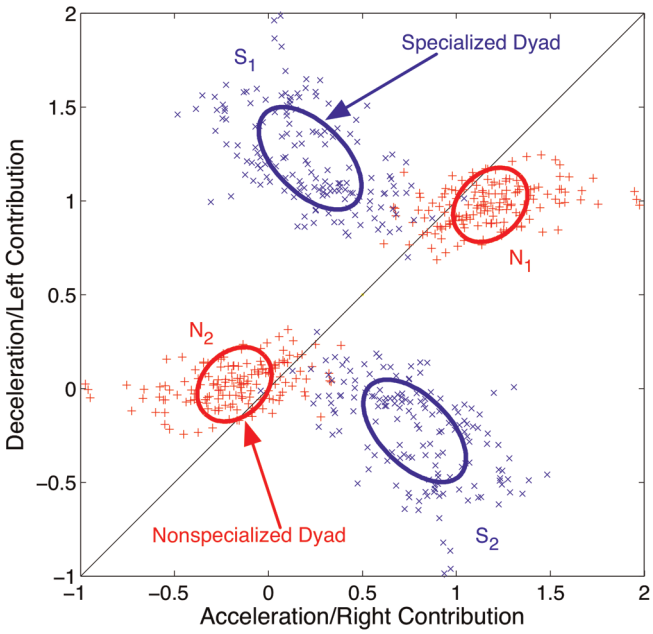

Fig. 3. **Degree of specialization**. A dyad's degree of specialization can be characterized by plotting each member's contribution to different phases of motion. This graph shows superimposed data sets for two different dyads, one of which exhibits significant specialization and one which did not. The contributions from each member of a dyad are necessarily opposite and appear as a dot mirrored about the center of the box (0.5, 0.5). The distance from the $x = y$ line is the degree of specialization. The center of the ellipse is the mean of the distribution, and the major and minor axes show one standard deviation.

was generally not adopted by our dyads even when a right-handed participant worked with a left-handed participant. Fig. 4 compares A/D specialization to L/R specialization. Of the 15 dyads, 11 show significantly more A/D specialization than L/R specialization ($p < 0.05$).

3.6 Discussion of Specialization

When a single person becomes part of a dyad, there is no longer a one-to-one correspondence between dynamics and kinematics due to the redundancy of the coupled limbs. For example, in a dyadic task, one participant can choose not to perform at all, to help with only the deceleration phase, or to use only elbow flexor muscles. Many dyads specialize according to task phase, which is one way they can solve the redundancy problem. We speculate that by specializing, each member is able to focus on a subset of the actions while the other member completes the complimentary actions.

Humans are able to precompute trajectories [8] based upon previous interactions [7]. When a specific event occurs, they are able to quickly perform the prepared action [32]. In the case of two people working together, the deceleration specialist possibly waits for some cue, such as reaching a certain location or velocity, and subsequently begins to decelerate the crank while the accelerator focuses on other aspects of the task.

A single person executing a target acquisition task in our experiment would be expected to use the triphasic burst pattern of muscle activity in which an agonist muscle burst starts the movement and is followed by an antagonist muscle burst to brake the movement and another agonist burst to help hold the limb at the final position [33], [34], [35]. These bursts represent careful planning based on prior knowledge, rather than feedback received during the task [36], [37]. Moreover, these patterns represent optimal movements that best accomplish the task within rather limiting physiological constraints such as the rates at which muscles can be turned on and off [37], [38] and the limited torque generating capacity in different areas of the workspace [39]. Consequently, the limited muscle activation rate becomes less critical if one person is able to initiate the deceleration phase while the other is finishing the acceleration phase.

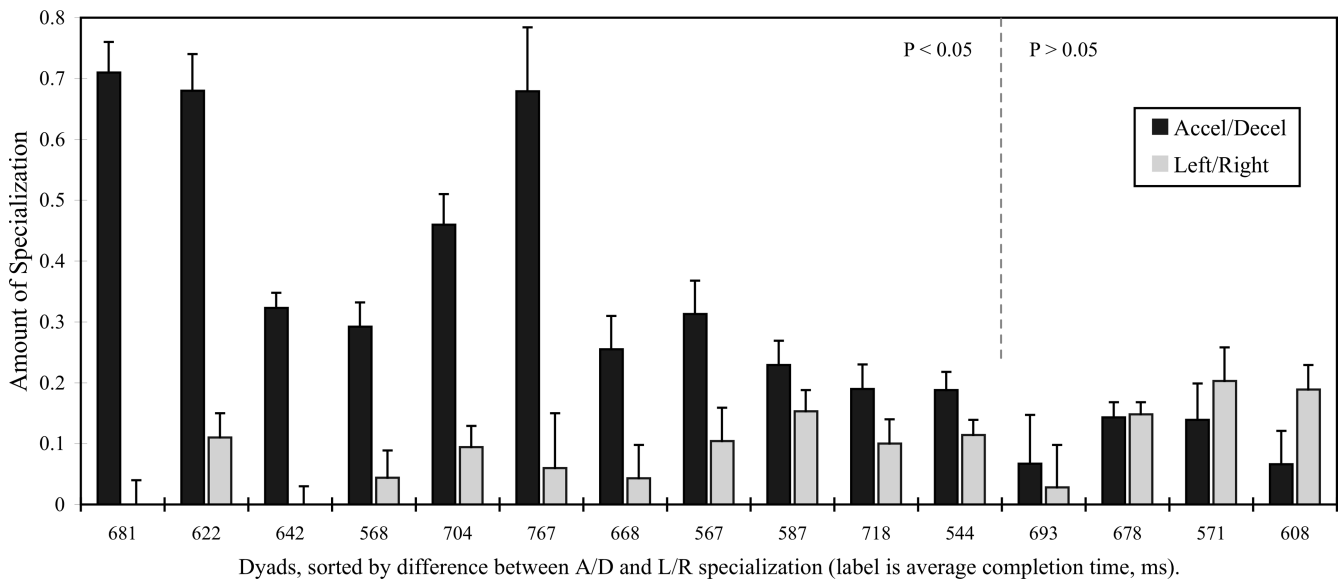

Fig. 4. **Type of specialization**. Each bar shows a dyad's average degree of specialization over all trials, and a 95 percent confidence interval. A value of 0.707 (the distance from the $x = y$ line to the point (1, 0), meaning one member's contribution to acceleration is 100 percent and to deceleration is 0 percent) can be considered "fully specialized." Of the 15 dyads, 11 show significantly ($p < 0.05$) more A/D specialization (black bars) than L/R specialization (gray bars). Many, but not all, dyads show a qualitatively high degree of A/D specialization.

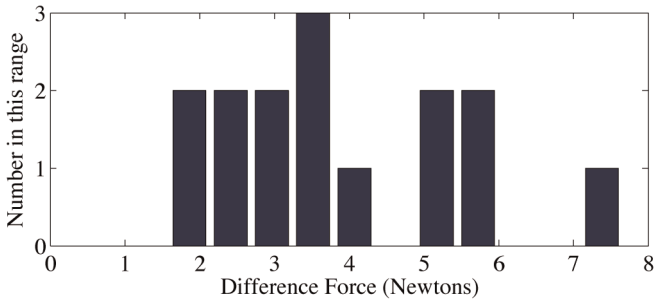

Fig. 5. **Dyadic contraction**. This histogram shows the average difference force once the dyad has reached the target and stayed there for 700 ms. In order for a dyad to complete the task, no force is required when the handle is in the target, yet many participants continued to maintain a significantly large opposition force, possibly for stability reasons similar to cocontraction in one person. Bin size is 0.55 N.

This specialized action may bring the dyad closer to the time-optimal bang-bang control strategy.

A similar specialization strategy has been observed in a single individual performing a bimanual task. Reinkensmeyer et al. [40] show that an individual holding both ends of a pencil with one finger from each hand will provide acceleration forces with one hand and deceleration forces with the other hand. This result might be taken as a bimanual (single person) model for our observed two-person A/D specialization. For a single individual, the inward force allows the pencil to accelerate and decelerate while being rigidly held and not dropped. The tight neural coupling between two arms in an individual allows a person to coordinate the actions of each arm [41]. In our dyadic task there is, of course, no direct neural coupling between the individuals, so the developed strategy must have occurred by haptic communication, since the only communication between participants is physical. Experiment two (Section 4) will further examine the development of specialization with an individual interacting with a consistent specialized partner.

3.7 Dyadic Contraction

We observed another dyadic effect as well. At the end of each trial, the members bring the crank to rest as they wait for the next target to appear. If a participant is working alone, there is no way to apply any force to the handle without also causing the crank to accelerate. If there are two members, however, each member can apply a force even though the crank is not moving, as long as the forces are equal and opposite. This generates a difference force with no net force, hence no acceleration.

We found that dyad members generally exert a significant force in opposition to one another (averaging 3.9 N). Fig. 5 shows the distribution of the average difference force for each pair at the end of each trial. The average standard deviation of all dyads was 3.1 N. Of the 15 dyads, 13 show a dyadic contraction force statistically significantly larger than 2 N ($p < 0.05$).

Dyadic contraction could serve to increase stiffness for the dyad in the same way that muscle cocontraction increases arm stiffness for an individual by cocontracting both the agonist and antagonist muscles to account for perturbations [42], [43]. Since the difference force at the end of the trials resembles cocontraction in individuals, we call the dyadic version "dyadic contraction." Dyadic contraction is a strategy similar to those used in parallel robotics and in human bimanual control [40], [44], [45]. It is possible that an individual could use an opposing force from each arm ("bimanual contraction") to replace some amount of cocontraction in each arm while achieving a similar stabilizing effect.

Dyadic contraction could also serve as a simple message between partners that they are in fact working together. By applying a small force, both participants feel an impedance that helps them determine what is happening on the other side of the curtain. Choi et al. [46] and Chib et al. [47] show that a person may attempt to maintain an optimum haptic force in order to explore a surface. Similarly, with two people working together, each member may be trying to maintain a certain force in order to learn about the partner across the curtain. The third experiment in this report (Section 5) will further examine dyadic contraction with a robotic partner.

4 Experiment Two: Human-Robot Specialization

The roles found in human-human interaction (Section 3.5) suggest that humans have an ability to cooperate and work together by specializing their force inputs. The second experiment, described here, attempts to replicate the interaction of two people on the same simple task by replacing one person with a robotic partner. The robotic partner is designed to mimic one partner's role in the specialization we found in human-human physical interaction. We expected participants to work with the robotic partner in a similar way as they did with a human partner.

It is possible that participants who know they are working with a robot will behave differently than if they believe they are working with a human. We therefore involved a confederate in half of the experiments. Each confederate stood in as if he was the participant's partner, but the confederate did not work with the participant, rather the robotic partner did. When queried later, most of the participants believed they were working with a human partner and were surprised when we explained that the confederate had never worked on the task with them. We expected that a robotic physical partner, thought to be a human partner, would elicit the same response from a participant as working with an actual human partner.

4.1 Robotic Partner

The motor under the table generated simulated torques while the participant was unaware of the origin of the forces. The robotic partner was composed of two parts: an active force production and a simulated inertia. The first part mimics the behavior of a specialized partner who has taken on the role of accelerating the crank. During the individual trials, we captured the acceleration part of a participant's own force profile and later used it as the robot's force profile. This acceleration force trajectory was multiplied by 2.1, which is the average amount an individual increased his/her force by when he/she becomes part of a dyad [28]. This

modified force trajectory becomes a typical force trajectory that could be found in a specialized member of a dyad.

We used a recorded version of the participant's forces to account for variations in forces and completion times among participants, so that any differences can be attributed to the participant working with a similar partner and not because the robotic partner is faster than the participant. The robotic partner playback is activated at the same time as the target is shown to the participant. Note that the recorded forces also contain the average reaction time from the individual, so the initial forces from the participant and the robot occur at a similar time.

The force trajectory for the robotic partner $(RP(t))$ is summarized as

$$RP(t) = 2.1 \times \sum_{i=1}^{100} \frac{\underline{f_i}(t)}{100}, \tag{2}$$

where $f_i(t)$ is a vector containing the forces for individual trials, $i = 1$ to 100 (for each trial), t is the time since the target was shown ($t = 1$ to 1,300), and $\underline{f_i}(t)$ only allows positive values of the force as defined by

$$\underline{f_i}(t) = \begin{cases} f_i(t), & : \quad f_i(t) > 0, \\ 0, & : \quad f_i(t) \leq 0. \end{cases} \tag{3}$$

The second part of the robotic partner simulates the inertia of an arm holding the handle, similar to that of a human partner. Four different confederates assisted throughout the experiments and an average of their arm inertias was used. The inertia (I) of each of their arms was calculated by grabbing the crank with the same grip participants used in the experiments. The motor applied a torque (τ) and we measured the angular acceleration (α) over a frequency range from 1 to 35 Hz. The inertia was then found from $I = \frac{\tau}{\alpha}$. The calculated average inertia used was $0.24\ \mathrm{kg} \cdot \mathrm{m}^2$. The simulated acceleration force was increased to compensate for the inertia of the simulated arm. The control loop for the robotic partner as well as measurements of the forces, acceleration, and position ran at 1 kHz on a computer running QNX.

4.2 Participants

Twenty-two students (seven male; one left-handed), age 18-24, from Northwestern University's Psychology participant pool participated with informed consent. None of the participants were involved in prior experiments.

4.3 Experimental Procedure

The experimental apparatus and procedure were similar to the human-human experiments (Section 3). However, a robotic partner replaced one of the participants and a confederate pretended to be the human participant's partner.

In half the trials, we employed a confederate who stood across from the participant during the experiment but did not physically interact with the participant. In this Human-Robot-Confederate (HRC) group, an experimental run started with the individual (I) or the confederate (C) performing a block of trials individually. Then, the other person completed a block of trials individually. Next, the individual worked with the robotic partner $(\mathrm{I_m})$, which the participants believed to be the human partner, who was actually waiting patiently behind the curtain. This sequence was performed twice, so six participants performed $(\mathrm{I, C, I_m, I, C, I_m})$ and five participants $(\mathrm{C, I, I_m, C, I, I_m})$. Presentation order made no significant difference.

A confederate was not present in the other half of the experiments. In this Human-Robot (HR) group, an experimental run started with the individual performing a block of trials alone followed by a block of trials in which the individual worked with the robotic partner. The participants knew there was not a human on the other side of the curtain. This sequence was performed twice, so 11 participants performed $(\mathrm{I, I_m, I, I_m})$.

4.4 Results

The 11 participants who worked with a robot in the absence of a confederate partner (HR) knew they were working with a nonhuman agent. Of the 11 participants, 10 who worked with a robot in the presence of a confederate (HRC) said they thought they were working with a person. Those 10 were actually quite surprised when we explained that they never worked with the other person. One of them had some doubts but was not sure either way, and this participant's results were not significantly different than the others. At a conscious level, the participants with a confederate present believed they were working with a human partner.

Within each participant, we computed the difference between the individual completion time and the completion time when working with a partner. Compared to an individual working alone: two humans (HH) were 54.5 ms faster (8.5 percent increase), a human working with a robotic partner with a confederate present (HRC) was 5.8 ms faster (0.9 percent increase), and a human working with a robotic partner without a confederate present (HR) was 24.8 ms slower (3.9 percent decrease). The HH group improved by 48.8 ms more than the HRC group $(t(11, 15) = 3.02, p < 0.01)$. The HRC group improved by 30.6 ms more than the HR group but not a statistically significant amount $(t(11, 15) = 1.52, p = 0.14)$. Fig. 6 summarizes the completion time results. The individuals composing the HH group are from experiment one and had an average individual completion time of 695.7 ms. The HR and HRC groups had an average individual completion time of 742.7 ms. The change in performance, not the absolute performance time, is the pertinent factor.

When working with the robotic partner, the participants are given an easy and natural way to specialize. The participant is completely responsible for all the force during deceleration, whereas the participant is free to choose their force during the acceleration phase. The motor applies enough torque to accelerate both the crank and the participant's arm, yet none of the participants developed specialization with the robotic partner. The human-robot force profiles shown in Fig. 7 convey a very different strategy than two specialized people working together (Fig. 2b). The human-robot force profiles show a strategy that is remarkably similar to an individual performing the task. There was not a statistically significant difference between individuals working alone and with the robotic partner during the acceleration portion of the task. Even

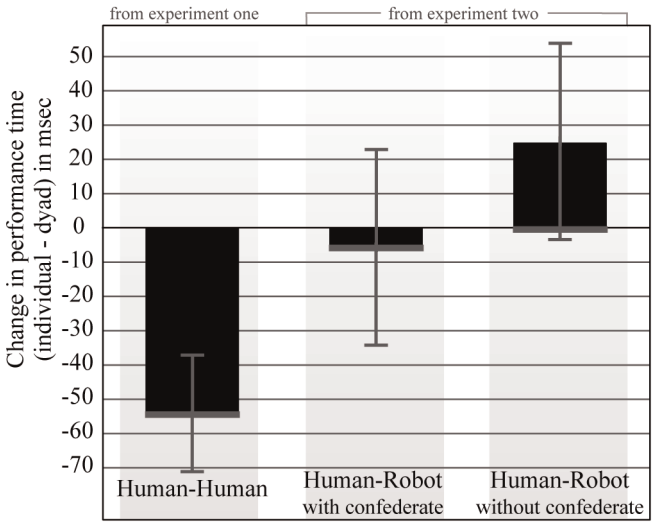

Fig. 6. **Change in completion times**. We programmed a robotic partner to simulate an acceleration specialized partner and compared the improvement time. Participants in the HRC group (middle) knew they were working with a robotic agent while participants in the HR group (right) believed they were working with a human but actually were not. Each dyad time is compared to the participant's individual time (normalized to 0), so a negative value means the two people (or human-robot pair) worked faster together than the individual. Two humans working together (left) were significantly faster than an individual working alone. When working with the robotic partner, participants with a confederate present are faster than participants without a confederate present. Error bars represent 95 percent confidence intervals.

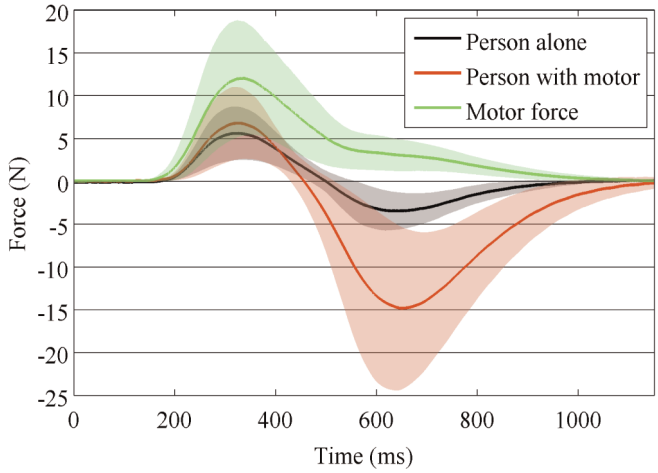

Fig. 7. **Average force profile for individuals working with the robotic partner**. The participants did not significantly change their force when the motor was working with them during the acceleration phase. The participants did significantly change their force during deceleration since the robotic partner was not aiding them during this phase. Semitransparent regions correspond to the standard deviation over all 22 participants.

though the participants believed they were working with a human partner, physically they were not acting as though they were.

4.5 Discussion

When the participants knew they were not working with a person, the robotic partner seemed to hinder their performance relative to working alone. When the participants believed they were working with a human partner, performance was similar to when the participants worked alone. Although the difference between the HR and HRC groups is not statistically significant, the participants appeared to treat a perceived human partner differently than a robot, which suggests that the perceived origin of forces can affect the way in which a person physically interacts. Quite evidently, we were not successfully able to fully mimic a human partner because the performance speedup that occurred was not captured when working with a robot, even when the participant believed he/she was working with a person.

4.5.1 Adaptability to Forces

When working alone, each participant knows what the result of their action will be, so each person can accurately predict the outcome of their actions. When working with a human partner, the outcome is less predictable since a partner's action is unknown. Sebanz et al. [48] show that a person will internally represent the actions of a person nearby when working on a complementary action. Presumably, haptic interactions also allow two people working together to depend on their partner to complete the complementary action of specialization. Thus, we expected that a person would also learn to depend on the robotic partner to complete the complementary action of specialization, but they did not.

Scheidt et al. [49] show that people can learn and adapt to unpredictable forces very quickly. The robotic partner does not learn and is consistent trial to trial, so the forces are more predictable than human forces. Thus, it is surprising to us that the participants did not learn to work with a predictable robotic partner in the same way as they did when working with an unpredictable human partner.

One possibility, however, is that it is the unpredictability of human action that leads partners to develop the observed specialization. Many participants opined after the human-human experiment that their partner was a hindrance, despite the fact that they actually performed faster with a partner. Perhaps specialization in this task serves to minimize the unpredictability of the partner's action by dividing the task into phases that are completed with more autonomy.

4.5.2 Haptic Turing Test

We may understand our robotic partner (in the presence of a nonworking human confederate "partner") as an attempt to pass a Haptic Turing Test. The Turing Test [50] was proposed for verbal communication between a human and a nonhuman agent. A nonhuman agent "passes" the test if it is able to simulate a human sufficiently well to deceive a human interlocutor.

In one sense, our robotic partner passed the Haptic Turing Test, since human participants believed that they were working with another person. On the other hand, our human participants did not act as they would with a human partner. Evidently, there is subtlety in human-human physical interaction that we were not only unable to capture in our robotic simulation, but that humans were unaware of at a conscious level.

4.5.3 Social Facilitation

The only difference between the HRC and HR groups is the participant's perception of where the force came from, either from the confederate or from the motor. Participants seemed to perform faster when they thought they were working with a human. Social Facilitation, as discussed in Section 3.4, cannot explain this performance increment since the theory relies only upon the visual presence of another person, which was not changing during these experiments.

Interacting physically, however, adds an additional mode of communication, which may cause the participants to discern the haptic presence of their partner. Similar to how Social Facilitation visually causes people to perform faster, haptic presence may also cause people to perform faster by allowing each person to perceive the actions of the other. The HRC group, believing in a human haptic presence, improved by 30.6 ms more than did the HR group, which had an evident lack of human haptic presence. However, if it exists, haptic presence cannot account for the entire performance benefit of dyads since two humans improved significantly more than either of the HR groups did.

5 Experiment Three: Human-Robot Perturbation Rejection

Experiments one and two dealt with interaction throughout the entire trial, where each trial can last more than a second. These repeated interactions are long enough to allow two people to specialize their forces and to develop dyadic contraction. We suspected that two people might similarly be able to asymmetrically divide a task, such as rejecting perturbation forces, over a smaller time period. Other studies [11], [51], [52] have shown that individuals can adjust the impedance of their arm to fit the task at hand, but there have been no studies showing whether two people can cooperatively adjust their joint impedance to better complete a task. Whereas experiment two concentrated on the roles during the motion, experiment three was only concerned with the roles after the motion had ended and was designed to measure how well dyads can jointly reject disturbances. Does the observed dyadic contraction force discussed in Section 3.7 serve to stiffen the participant's arms analogously to cocontraction in an individual?

5.1 Participants

Twenty-two (10 male; eight left-handed) participants, age 18-24, from Northwestern University's Psychology participant pool participated with informed consent. None of the participants were involved in prior experiments. Three dyads were male-male, four were female-male, and four were female-female. One dyad was composed of two left-handed individuals, six dyads were composed of one left- and one right-handed individual, and four dyads were composed of two right-handed individuals.

5.2 Experimental Protocol

We used the same experimental apparatus as the two previous experiments. The procedure was similar to the human-human experiments used in Section 3 except that the motor applied additional forces during the task.

We used the motor to apply perturbations to the individuals and dyads at the end of some trials. The perturbations started after the participants were inside the target for 700 ms and lasted for 100 ms. We measured the resulting displacement at the moment the perturbation force ended. This displacement relates to the impedance of the system composed of the crank and either one or two individuals holding the handles. Similar to the dyad case in experiment one, we physically doubled the inertia, and in this experiment, we also doubled the perturbation force. The force pulse was 10 N for individuals and 20 N for dyads. Since we doubled the inertia and perturbation force, the displacements due to a perturbation would be exactly the same if each individual performed exactly the same alone as in a dyad. Doubling the perturbation force and inertia normalized the results across individuals and dyads.

During individual trials, we also used the motor to simulate dyadic contraction near the targets. We programmed the motor to push toward the center of the workspace with a constant force when the crank was near the targets. Since we were unconcerned with the intervening motion, the simulated dyadic contraction force would linearly transition from side to side as the participants moved the crank. The transition occurred far from the targets. The constant force near the handles was one of three values, either 0, 5, or 10 N. The constant simulated dyadic contraction force was the same for 48 continuous trials. The individual blocks (154 trials) consisted of the following sequence:

$$\begin{array}{rl} \text{10 trials} & \text{Warm up (no force or perturbations)} \\ \text{48 trials each} & \left\{\begin{array}{l}\text{No dyadic contraction force,}\\ \text{5 N dyadic contraction force,}\\ \text{10 N dyadic contraction force.}\end{array}\right. \\ \text{with 8 perturbations} & \\ \text{(random order)} & \end{array}$$

The entire experiment consisted of the following sequence:

$$\begin{array}{rl} \text{154 trials each} & \left\{\begin{array}{l}\text{A alone with simulated}\\ \quad\text{dyadic contraction force,}\\ \text{B alone with simulated}\\ \quad\text{dyadic contraction force,}\\ \text{A and B together with no simulated}\\ \quad\text{dyadic contraction force,}\end{array}\right. \\ \text{(random order)} & \end{array}$$

Repeat same sequence once.

The perturbation force randomly varied between the dominant and nondominant sides of the workspace and occurred in both clockwise (CW) and counterclockwise (CCW) directions. Fig. 8 shows the possible arrangements and perturbation types of the participants during the experiment. Each block of 154 trials contained 24 perturbations so one in six trials, randomly spaced, included a perturbation. In the individual case, there were 12 combinations since there were four perturbation types and three simulated dyadic contraction forces. Each block of 154 was repeated, so each perturbation/force combination was tested four times per 30-minute experiment. In the same-handed dyad case, there were four perturbation types and no simulated contraction force so there were four possible combinations. Each perturbation type was tested 12 times per 30-minute experiment. The table at the bottom of Fig. 8 summarizes the total number of perturbations for each experiment.

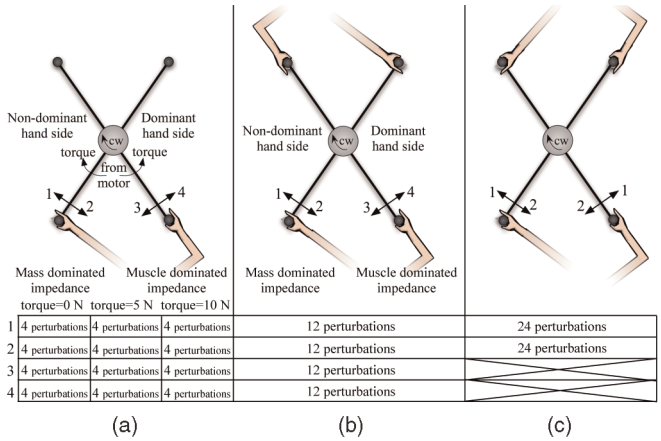

1	4 perturbations	4 perturbations	4 perturbations	12 perturbations	24 perturbations
2	4 perturbations	4 perturbations	4 perturbations	12 perturbations	24 perturbations
3	4 perturbations	4 perturbations	4 perturbations	12 perturbations	
4	4 perturbations	4 perturbations	4 perturbations	12 perturbations	

Fig. 8. **Participant configuration during experiment three**. (a) Participant alone. (b) Same-handed Dyads. (c) Different-Handed Dyads. The arrows indicate the direction and location of each perturbation type. (a) and (b) have four different perturbation types. (c) has two perturbation types since both members' arms are always in an opposite configuration. Listed below the configurations are the total number of perturbations per experiment.

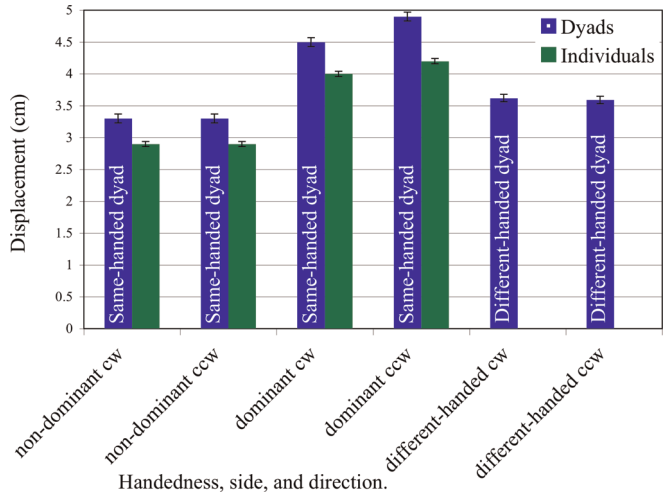

Fig. 9. **Average resulting displacement from a force perturbation**. The same-handed dyads show a similar pattern as individuals, but the dyads had a larger displacement in every case. Unlike specialization, the same-handed dyads were not able to improve their performance by joining forces. The different-handed dyads were able to use the best of each participant's ability for each perturbation type to create a consistently good joint effort. Error bars represent 95 percent confidence intervals.

As seen in Figs. 8b and 8c, there are two ways in which a dyad could be formed. Same-handed dyads consist of either two right-handed or two left-handed participants. Different-handed dyads consist of one left-handed and one right-handed participant. In all previous experiments, this difference has not been statistically significant, but due to the asymmetric perturbation rejection characteristics of the arm configuration, it is relevant in this analysis.

5.3 Results

We conducted a one-way analysis of variance (ANOVA) to determine if the displacement due to perturbation was significantly different within and across each of the three groups: individuals, same-handed dyads, and different-handed dyads. When the ANOVA yielded significant results, we used Tukey's honest significant difference test for post hoc analyses. Fig. 9 summarizes the resulting displacements for individuals and dyads.

5.3.1 Individuals

Individuals performed this task with both a simulated dyadic contraction force and an occasional perturbation force applied after they reached the target. We specified the dyadic contraction force, whereas in dyad trials the members cooperatively negotiated a dyadic contraction force. Specifying the simulated dyadic contraction force for individuals allowed us to measure the effect an external force had on rejecting perturbations. At each force level, we found a statistically significant difference in the displacement $(F(2,21) = 27.5, p < 0.05)$. At 0 N, the displacement was 3.9 cm; at 5 N, it was 3.7 cm; and at 10 N, it was 3.4 cm.

Our results match well with the more elaborate force ellipses published in several studies. Perreault et al. [53] and Gomi and Osu [54] examined how the forces and geometry of the arm can change the viscoelastic properties of the arm muscles. Similar to our results, they found that an external force increased the impedance of the arm. The benefit of an external force implies that the dyadic contraction force we observed may benefit dyads in rejecting a disturbance, but as we will discuss in Section 5.3.2, this result does not extend to two humans.

The direction and location of the perturbation affected the displacement in individuals. There was a statistically significant difference between the nondominant side (3.0 cm) and both perturbation directions (CW: 4.0 cm; CCW: 4.2 cm) on the dominant side $(F(3,21) = 207.1, p < 0.01)$. On the nondominant side, the two perturbation directions were not statistically different. The difference between dominant and nondominant sides was largely due to the position of the arm. On the nondominant side, the arm was reaching across the body, so the effective inertia was greater.

On the dominant side the larger displacements occurred when the perturbation was in the same direction that the individual was pushing (i.e., away from the center). Near the targets, the simulated dyadic contraction force was constant. Cocontraction in an individual works by contracting opposing muscles that act like springs. Any deviation from the equilibrium point would cause a restoring force. The simulated dyadic contraction force did not increase like a muscle would, so there was no restoring spring force to resist the perturbation when the perturbation was in the same direction that an individual was pushing.

5.3.2 Dyads

Completing this task as a dyad surprisingly resulted in worse perturbation rejection characteristics of the connected humans. We expected they would be at least as good as individuals. If each member of a dyad performed exactly the same as they did as an individual, the displacement would have been exactly the same for individuals and dyads, but this did not occur. Fig. 9 shows that same-handed dyads had larger displacements than individuals in all four perturbation types. Each perturbation direction was statistically significantly different between same-handed dyads and

individuals ($F(7,25) = 107.5, p < 0.05$). Same-handed dyads exhibited the same characteristics as an individual: the dominant side was less resistant than the nondominant side, which was largely due to the inertia being greater on the nondominant side. The displacement on the dominant side (CW: 4.5 cm; CCW: 4.9 cm) was affected by the perturbation direction ($F(3,4) = 172.1, p < 0.01$) but not on the nondominant side (3.3 cm). Unlike during task movement, same-handed dyads were unable to benefit from working with a partner.

If the same-handed dyad's impedance was simply the sum of both participant's arm impedance, the displacement would have been the same as the individual's. Other mechanisms for rejecting disturbances such as stretch reflexes and cocontraction within each individual would have continued to operate within each member of the dyad. Like the individuals, dyads had a dyadic contraction force (3.8 N average with a standard deviation of 3.0), but dyads were unable to beneficially unite arm impedances when working in a parallel arrangement. Conversely, different-handed dyads do tend toward an average of the individual member's impedance, but this is largely due to the configuration of their arms. The highest impedance position (mass dominated) of one member will be paired with the lowest impedance position from the other member. Excluding arm configuration, dyads were not able to divide the task or improve on a quick interaction as they did during the longer interactions.

During the task, each participant may have been unable to definitively identify the source of the perturbation. Each member may have incorrectly identified the source of the perturbation as the other member, which could have caused that participant to depend less on their partner. This is only speculation but does highlight one of the difficulties of cooperating with another agent in an imperfect environment.

6 Discussion

When two people cooperate to complete a task, their forces combine to produce a new set of actions. The redundant set of actions allowed the dyads to specialize in time such that one member takes on acceleration and the other deceleration. On the target acquisition task in experiment one, we found that dyads were 8.5 percent faster than individuals despite their perception that a partner was a hindrance. This improved performance, negotiation, and specialization all developed solely through haptic communication.

Haptic communication is not beneficial for all dyadic tasks, however. When a brief force perturbs a dyad, the dyad had a 12.2 percent larger displacement than an individual, which demonstrates a lack of beneficial cooperation for perturbation rejection. Presumably, the interaction was too short to allow a haptic negotiation.

By replicating the haptic interactions found in two humans, we replaced one of the members with a robotic partner, which allowed us to perform a Haptic Turing Test. Our Haptic Turing Test revealed that our acceleration specialized robotic partner was able to convince participants that they were working with another human. On the other hand, we were unable to get participants working with a robot to perform as quickly as a human-human pair or to adopt the same specialized strategy, so perhaps we failed a stricter Turing Test. The difference between participants who believed the robot was a human and those who knew they were working with a robot suggests that the perception of what and where the forces come from can affect a person's interactions. Further study is required to determine how fast two people can communicate haptically and also how to induce a human-robot pair to work as well as a human-human pair.

When working cooperatively, each member of a dyad will feel that they are contributing less to the overall task than they actually are since self-generated forces are perceived as weaker than externally generated forces [25]. We speculate that an escalation in force could cause improved performance since each participant may attempt to achieve an equal perceived contribution. Both members of a dyad cannot each simultaneously perceive that they are contributing half of the required force. An attempt to apply equal perceived forces could be recognized by the other member as a desire to move faster, which would cause the performance to escalate. It is unclear, however, whether increased forces cause faster performance or the other way around.

The experiments presented in this paper all focused on a one DOF task. Further experiments on a cooperative task in SE(2), such as positioning a desk in a room, will further reveal specialization methods a dyad could exploit using the redundant controls that are available in planar motion. One possible method of specialization could be dividing control of each axis and cooperatively controlling the rotation. This would be akin to A/D specialization. Another method might take the form of coupling all three DOFs with continual haptic interaction through each motion, much like dyadic contraction. The relative location of each member will likely have a significant affect on the type of specialization and resulting actions.

Mirror neurons have been shown to encode actions in an individual visually observing another person performing a task, but the authors are not aware of any studies examining mirror neurons when haptically interacting. If mirror neurons encode the actions of another when physically interacting, the mirror neurons could provide an estimate of a partner's future actions. Such a neurological internal loop would allow dyads to closer achieve the neural coupling found in bimanual interaction and, thus, explain how they are able to achieve a similar type of specialization.

Acknowledgments

This work was supported by the US National Science Foundation under Grant ECS-0433948. The authors would like to thank James Patton, Mitra J. Hartmann, Peter Vishton, Marcia Grabowecky, J. Edward Colgate, Kevin Lynch, and Satoru Suzuki for their help with our research.

References

[1] V. Klingspor, J. Demiris, and M. Kaiser, "Human-Robot Communication and Machine Learning," *Applied Artificial Intelligence J.*, vol. 11, pp. 719-746, 1997.

[2] N. Sebanz, H. Bekkering, and G. Knoblich, "Joint Action: Bodies and Minds Moving Together," *Trends in Cognitive Sciences*, vol. 10, no. 2, pp. 70-76, 2006.

[3] R.C. Schmidt, C. Carello, and M.T. Turvey, "Phase Transitions and Critical Fluctuations in the Visual Coordination of Rhythmic Movements between People," *J. Experimental Psychology: Human Perception and Performance,* vol. 16, no. 2, pp. 227-247, 1990.

[4] G. Pellegrino, L. Fadiga, L. Fogassi, V. Gallese, and G. Rizzolatti, "Motor Facilitation during Action Observation: A Magnetic Stimulations Study," *J. Neurophysiology,* vol. 73, pp. 2608-2611, 1992.

[5] G. Rizzolatti, L. Fogassi, and V. Gallese, "Neurophysiological Mechanisms Underlying the Understanding and Imitation of Action," *Nature Rev. Neuroscience,* vol. 2, no. 9, pp. 661-670, 2001.

[6] J.W. Krakauer, M.F. Ghilardi, and G. Ghez, "Independent Learning of Internal Models for Kinematic and Dynamic Control of Reaching," *Nature Neuroscience,* vol. 2, pp. 1026-1031, 1999.

[7] A.M. Gordon, G. Westling, K.J. Cole, and R.S. Johansson, "Memory Representations Underlying Motor Commands Used during Manipulation of Common and Novel Objects," *J. Neurophysiology,* vol. 69, no. 6, pp. 1789-1796, 1993.

[8] R.S. Johansson, "Sensory Input and Control of Grip," *Proc. Novartis Foundation Symp. Sensory Guidance of Movement,* vol. 218, pp. 45-59, 1998.

[9] Y. Uno, M. Kawato, and R. Suzuki, "Formation and Control of Optimal Trajectory in Human Multijoint Arm Movement. Minimum Torque-Change Model," *Biological Cybernetics,* vol. 61, no. 2, pp. 89-101, 1989.

[10] T. Flash and N. Hogan, "The Coordination of Arm Movements: An Experimentally Confirmed Model," *J. Neuroscience,* vol. 5, no. 7, pp. 1688-1703, 1985.

[11] D.W. Franklin, E. Burdet, R. Osu, M. Kawato, and T.E. Milner, "Functional Significance of Stiffness in Adaptation of Multijoint Arm Movements to Stable and Unstable Dynamics," *Experimental Brain Research,* vol. 151, pp. 145-157, 2003.

[12] R. Shadmehr and F.A. Mussa-Ivaldi, "Adaptive Representation of Dynamics during Learning of a Motor Task," *J. Neuroscience,* vol. 14, no. 5, pp. 3208-3224, 1994.

[13] J.K. Hodgin, W.L. Wooten, D.C. Brogan, and J.F. O'Brien, "Animating Human Athletes," *Proc. ACM SIGCHI Conf. Computer Graphics and Interactive Techniques,* pp. 71-78, 1995.

[14] P.M. Fitts, "The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement," *J. Experimental Psychology,* vol. 47, pp. 381-391, 1954.

[15] M. Li and A. Okamura, "Recognition of Operator Motions for Real-Time Assistance Using Virtual Fixtures," *Proc. 11th Int'l Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS '03),* pp. 125-131, 2003.

[16] B. Corteville, E. Aertbelien, H. Bruyninckx, J.D. Schutter, and H.V. Brussel, "Human-Inspired Robot Assistant for Fast Point-to-Point Movements," *Proc. Int'l Conf. Robotics and Automation (ICRA),* 2007.

[17] S. Gentry, "Dancing Cheek to Cheek: Haptic Communication between Partner Dancers and Swing as a Finite State Machine," PhD thesis, Massachusetts Inst. Technology, 2005.

[18] N. Wegner and D. Zeaman, "Team and Individual Performance on a Motor Learning Task," *J. General Psychology,* vol. 55, pp. 127-142, 1956.

[19] E. Field and D. Harris, "A Comparative Survey of the Utility of Cross-Cockpit Linkages and Autoflight Systems' Backfeed to the Control Inceptors of Commercial Aircraft," *Ergonomics,* vol. 41, no. 10, pp. 1462-1477, 1998.

[20] L. Summers, J. Shannon, T. White, and R. Shiner, "Fly-by-Wire Sidestick Controller Evaluation," *Proc. Aerospace Technology Conf. and Exposition (AeroTech '87),* Soc. Automotive Engineers, SAE Technical Paper 871761, 1987.

[21] A. Karniel, R. Meir, and G.F. Inbar, "Exploiting the Virtue of Redundancy," *Proc. Int'l Joint Conf. Neural Networks (IJCNN),* 1999.

[22] F. Lacquaniti and C. Maioli, "Distributed Control of Limb Position and Force," *Tutorials in Motor Behavior II,* R. GESaJ, ed., pp. 31-54, 1992.

[23] S. Glynn and R. Henning, "Can Teams Outperform Individuals in a Simulated Dynamic Control Task?" *Proc. Ann. Meeting Human Factors and Ergonomics Soc. (HFES '00)* (ISSN: 10711813), vol. 6, pp. 141-144, 2000.

[24] G. Knoblich and J.S. Jordon, "Action Coordination in Groups and Individuals: Learning Anticipatory Control," *J. Experimental Psychology: Learning, Memory, and Cognition,* vol. 29, no. 5, pp. 1006-1016, 2003.

[25] S.S. Shergill, P.M. Bays, C.D. Frith, and D.M. Wolpert, "Two Eyes for an Eye: The Neuroscience of Force Escalation," *Science Magazine,* vol. 301, p. 187, 2003.

[26] I. Elhajj, H. Hummert, N. Xi, and Y. Liu, "Real-Time Haptic Feedback in Internet-Based Telerobotic Operation," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA),* 2000.

[27] J.P. Hespanha, M. McLaughlin, G.S. Sukhatme, M. Akbarian, R. Garg, and W. Zhu, "Haptic Collaboration over the Internet," *Proc. Fifth Phantom Users Group Workshop (PUG),* 2000.

[28] K. Reed, M. Peshkin, M.J. Hartmann, M. Grabowecky, J. Patton, and P.M. Vishton, "Haptically Linked Dyads: Are Two Motor Control Systems Better than One?" *Psychological Science,* vol. 17, no. 5, pp. 365-366, 2006.

[29] N. Triplett, "The Dynamogenic Factors in Pacemaking and Competition," *Am. J. Psychology,* no. 9, pp. 507-533, 1898.

[30] R.B. Zajonc, "Social Facilitation," *Science Magazine,* vol. 149, pp. 269-274, 1965.

[31] B.H. Schmitt, T. Gilovich, N. Goore, and L. Joseph, "Mere Presence and Social Facilitation: One More Time," *J. Experimental Social Psychology,* vol. 22, pp. 242-248, 1986.

[32] D.A. Rosenbaum, "Human Movement Initiation: Specification of Arm, Direction, and Extent," *J. Experimental Psychology: General,* vol. 109, pp. 444-474, 1980.

[33] M. Hallett, B. Shahani, and R. Young, "EMG Analysis of Stereotyped Voluntary Movements in Man," *J. Neurosurgery Psychiatry,* vol. 38, pp. 1154-1162, 1975.

[34] B. Hannaford and L. Stark, "Roles of the Elements of the Triphasic Control Signal," *Experimental Neurology,* vol. 90, pp. 619-634, 1985.

[35] G.L. Gottlieb, D.M. Corcos, G.C. Agarwal, and M.L. Latash, "Principles Underlying Single Joint Movement Strategies," *Multiple Muscle Systems,* J.M. Winters and S.L.Y. Woo, Springer-Verlag, pp. 236-250, 1990.

[36] G.L. Gottlieb, D.M. Corcos, and G.C. Agarwal, "Organizing Principles for Single-Joint Movements. I. A Speed Insensitive Strategy," *J. Neurophysiology,* vol. 62, pp. 342-357, 1989.

[37] G.L. Gottlieb, C.H. Chen, and D.M. Corcos, "Nonlinear Control of Movement Distance at the Human Elbow," *Experimental Brain Research,* vol. 112, pp. 289-297, 1996.

[38] C.F. Ramos, S.S. Hacisalihzade, and L.W. Stark, "Behavior Space of a Stretch Reflex Model and Its Implications for the Neural Control of Voluntary Movement," *Medical and Biological Eng. and Computing,* vol. 28, pp. 15-23, 1990.

[39] J. Prodoehl, G. Gottlieb, and D. Corcos, "The Neural Control of Single Degree-of-Freedom Elbow Movements: Effect of Starting Joint Position," *Experimental Brain Research,* vol. 153, pp. 7-15, 2003.

[40] D.J. Reinkensmeyer, P.S. Lum, and S.L. Lehman, "Human Control of a Simple Two-Hand Grasp," *Biological Cybernetics,* vol. 67, no. 6, pp. 553-564, 1992.

[41] S.P. Swinnen and N. Wenderoth, "Two Hands, One Brain: Cognitive Neuroscience of Bimanual Skill," *Trends in Cognitive Sciences,* vol. 8, no. 1, pp. 18-25, 2004.

[42] R. Osu, D.W. Franklin, H. Kato, H. Gomi, K. Domen, T. Yoshioka, and M. Kawato, "Short- and Long-Term Changes in Joint Co-Contraction Associated with Motor Learning as Revealed from Surface EMG," *J. Neurophysiology,* vol. 88, pp. 991-1004, 2001.

[43] T.E. Milner, "Adaptation to Destabilizing Dynamics by Means of Muscle Cocontraction," *Experimental Brain Research,* vol. 143, pp. 406-416, 2002.

[44] J.L. Patton and P. Elkins, *Training with a Bimanual-Grasp Beneficially Influences Single Limb Performance.* Soc. Neuroscience, 2002.

[45] V. Chib, J. Patton, K. Lynch, and F. Mussa-Ivaldi, "Haptic Discrimination of Perturbing Fields and Object Boundaries," *Proc. 12th Int'l Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS '04),* vol. 0-7695-2112-6/04, 2004.

[46] S. Choi, L. Walker, H.Z. Tan, S. Crittenden, and R. Reifenberger, "Force Constancy and Its Effect on Haptic Perception of Virtual Surfaces," *ACM Trans. Applied Perception,* vol. 2, no. 2, pp. 89-105, 2005.

[47] V.S. Chib, J.L. Patton, K.M. Lynch, and F.A. Mussa-Ivaldi, "Haptic Identification of Surfaces as Fields of Force," *J. Neurophysiology,* vol. 95, pp. 1068-1077, 2006.

[48] N. Sebanz, G. Knoblich, and W. Prinz, "Representing Others' Actions: Just Like One's Own?" *Cognition,* vol. 88, pp. 11-21, 2003.

[49] R.A. Scheidt, J.B. Dingwell, and F.A. Mussa-Ivaldi, "Learning to Move amid Uncertainty," *J. Neurophysiology,* vol. 86, 2001.

[50] A. Turing, "Computing Machinery and Intelligence," *Mind,* vol. LIX, no. 236, pp. 433-460, 1950.

[51] D.W. Franklin and T.E. Milner, "Adaptive Control of Stiffness to Stabilize Hand Position with Large Loads," *Experimental Brain Research,* vol. 152, pp. 211-220, 2003.
[52] M. Darainy, N. Malfait, P.L. Gribble, F. Towhidkhah, and D.J. Ostry, "Learning to Control Arm Stiffness under Static Conditions," *J. Neurophysiology,* vol. 92, pp. 3344-3350, 2004.
[53] E.J. Perreault, R.F. Kirsch, and P.E. Crago, "Effects of Voluntary Force Generation on the Elastic Components of Endpoint Stiffness," *Experimental Brain Research,* vol. 141, pp. 312-323, 2001.
[54] H. Gomi and R. Osu, "Task-Dependent Viscoelasticity of Human Multijoint Arm and Its Spatial Characteristics for Interaction with Environments," *J. Neuroscience*, vol. 18, no. 21, pp. 8965-8978, 1998.

Kyle B. Reed received the BS degree in mechanical engineering from the University of Tennessee in 2001 and the MS and PhD degrees in mechanical engineering from Northwestern University in 2004 and 2007, respectively. He is currently a postdoctoral fellow in the Laboratory for Computational Sensing and Robotics, Johns Hopkins University. His interests include haptics, human-machine interaction, rehabilitation engineering, medical robotics, and engineering education. He was the recipient of the 2001 US National Science Foundation Graduate Fellowship. He is a member of the IEEE.

Michael A. Peshkin received the PhD degree. He is a professor of mechanical engineering in the Department of Mechanical Engineering, Northwestern University. His research is in robotics and human-machine interaction, and rehabilitation robotics. He also works in novel actuators and sensors. He has cofounded three spin-offs: Mako Surgical (image guided surgery), Cobotics (assistive devices for materials handling), and Kinea Design (gait and balance training robot for rehabilitation after stroke). He holds 15 patents and is a coinventor (with J. Edward Colgate) of Cobots: collaborative robots for direct interaction with humans in a shared task. He is a member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.