Mechanism of interpersonal sensorimotor interaction

Atsushi Takagi

A thesis presented for the degree of
Doctor of Philosophy

Department of Bioengineering
Imperial College London

September 2015

This is to certify that the work in this thesis has been carried out at the Department of Bioengineering, Imperial College of Science, Technology and Medicine and has not been previously submitted to any other university or technical institution for a degree or award. The thesis comprises only my original work, except where due acknowledgement is made in the text.

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives license. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the license terms of this work.

Abstract

This thesis uncovers the mechanism of interaction between human partners in continuous contact. During such sensorimotor interaction, I propose that people interpret their partner's desired movement through the forces they feel. The evidence of this *interpersonal goal integration* mechanism comes from three behavioural experiments and computational modelling of their results. Novel dual-robotic interfaces were used in the experiments for two reasons: first, to eliminate the partner's gaze as it was found to influence the behaviour of interacting pairs of subjects. Second, to systematically and quantitatively measure the effects of sensorimotor interaction on task performance and adaptations in muscular activity, i.e. motor force and joint impedance.

The first experiment revealed that sensorimotor interaction during a tracking task, where partners track a randomly moving target with their cursor, enabled both the worse and better partners in the pair to improve at tracking. This mutually beneficial effect of interaction, regardless of the partner's skill, cannot be explained by existing theories in psychology and neuroscience, but the discovered *interpersonal goal integration* mechanism does capture this behaviour; when this interaction mechanism was implemented in a robot partner, people interacting with this robot behaved similarly to those during interpersonal interaction.

Adaptations in the muscular activity were observed in the second experiment, where pairs were rigidly or weakly coupled during a tracking task. The better partners increased their motor force and joint impedance, whereas the worse partners decreased both. The interaction mechanism also predicted these changes in the motor force.

Finally, an experiment where four partners interacting together during a tracking task showed that the best partner's performance was not hindered by many worse partners, which was also predicted by the *interpersonal goal integration* mechanism.

The *interpersonal goal integration* mechanism is the first robotic algorithm to exhibit human-like behaviour during sensorimotor interaction. It has applications where continuous contact with a human partner is critical, such as the transport of large objects or in the co-operative control of lower-limb exoskeletons. It offers a new approach for robotic rehabilitation through continuous, physical interaction which exploits the benefits of simultaneous action observation and physical training. Finally, the interaction mechanism may be used as an investigative and therapeutic tool for autism spectrum disorder (ASD). Individuals affected by ASD commonly lack the ability to extrapolate someone's goal by observing their actions alone. Similarly, they may lack the ability to estimate a partner's desired movement from observing haptics forces alone, thus the *interpersonal goal integration* mechanism may shed light on the specific differences between people with ASD and healthy individuals.

Acknowledgements

I would like to thank my supervisor Etienne who has been the driving force behind this thesis. The discussions we had were critical to the success of the work presented here. Your encouragement and support enabled me to pursue some crazy projects that somehow worked out in the end!

I would also like to thank Cheng, Saad and Wassif for the awesome years before and during this PhD. Latymer Court FTW!

To the people at University of Tokyo: Abe-san fūfu, Hirashima-san, Kobayashi-san, Nozaki-sensei, Ojio-san, Sasaki-kun, Shinya-san fūfu, Takiyama-san and Takuji-kun; I had a wonderful time, and look forward to going back for the epic food!

To the past and current members of the Human Robotics Group: Alejandro, Ale, Alfredo, Andrew, Arash, Asif, Carlo, Consuelo, Cous-cous, Donald, Elisabeth, Flavia, Francesca, Francesco, Franck, Ganesh, Ildar, Julius, Kelly, Luca, Martina, Matjaz, Mike, Moritz, Nat, Niek, Pascal, Paul, Paulo, Sarah, Siva, Sharah, Shou-han, Vale and Wayne; thanks for the amazing time at Imperial and in London, I'm truly grateful to have worked alongside you!

Arigatō to my relatives for the great times in Japan: Akari-chan, Bā-bā, Fumiyo-no-ojī-chan, Hajime-chan, Hitomi-chan, Hitoshi-san, Jicchan, Mā-chan, Mariko-no-obā-chan, Miwa-san, Nacchan and Nori-chan.

Finally, I would like to thank my family for their financial help and emotional support which fuelled my passion for research. Thank you, Okā-san and Otō-san, for teaching me how to walk and explore this world, and to Daisuke and Kasia for the much needed relaxation in Hawai'i!

This thesis is dedicated to you.

Contents

1 Introduction **8**

2 Social bias during interactions **11**

2.1 Solo force reproduction 12

2.2 Exchanging forces in pairs 14

2.2.1 Interaction separated by a curtain 14

2.2.2 Face-to-face interaction 16

2.2.3 Face-to-face interaction with the opposite sex 17

2.3 Discussion 19

3 Dyadic reaching **20**

3.1 Dyadic reaching 22

3.2 Dyadic reaching model 28

3.3 Discussion 33

4 Sensorimotor interaction mechanism **34**

4.1 Dyadic tracking 35

4.2 Proposed mechanisms of interaction 37

4.3 Dyadic tracking model 38

4.4 Interpersonal goal integration 44

4.4.1 Building an internal model of the partner 44

4.4.2 Human-robot validation experiment 49

4.5 Discussion 51

5 Connection strength on interaction **52**

5.1 Connection strength experiment 54

5.2 Connection strength simulation 58
5.3 Discussion 62

6 Multi-human interaction **63**
6.1 Multi-human experiment 64
6.2 Why do more partners help? 68
6.2.1 Interpersonal goal integration for multiple partners 70
6.2.2 Other partners' performance modulates improvement 71
6.3 Discussion 75

7 Conclusion **76**

List of Figures

1.1 Examples of interpersonal sensorimotor interaction 8
1.2 Experimental setups used 10

2.1 Individual force reproduction 13
2.2 Force exchange protocol 15
2.3 Force reproduction in pairs 18

3.1 Dyadic reaching protocol 23
3.2 Dyadic reaching torque profiles 1 26
3.2 Dyadic reaching torque profiles 2 27
3.3 Interaction torque profile templates from simulations 30
3.4 Interpreting dyadic reaching movements 32

4.1 Compliant interaction experiment 36
4.2 Proposed mechanisms of interaction 36
4.3 Parameter range for sensitivity analysis 43
4.4 Simulation of proposed mechanisms 43
4.5 Forming an internal model of the partner 44
4.6 Identification of partner's goal 47
4.7 Simulation of interpersonal goal integration model 48
4.8 Human-robot validation experiment of interaction model 50
4.9 Joint impedance and tracking skill are not related 50

5.1 Connection strength experimental protocol 55
5.2 Changes in improvement with rigid, medium and compliant interaction 56
5.3 Changes in effort with rigid, medium and compliant interaction 57
5.4 Prediction from naïve interaction model and co-activity 60

5.5 Prediction from modified interpersonal goal integration model 61

6.1 Multi-human sensorimotor interaction experiment 65
6.2 Results from multi-human interaction 67
6.3 Summation of multiple trajectories does not explain improvement 69
6.4 Prediction of interaction model for multi-human interaction 70
6.5 Other partners' skill, not their average, affects improvement 72
6.6 Improvement loop during interaction 74

1 Introduction

We are social animals; we learn, we play, we bond, we help, we fight and we console each other. Interpersonal interaction is critical to our society, which is why verbal communication or interpersonal sensory interaction between humans has been thoroughly investigated in the past. In comparison, motor interaction between partners has been studied little due to technological limitations, as this type of physical interaction requires sensors and robotic devices to investigate. However, it is a vital aspect of interpersonal interaction as the examples in Figure 1.1 demonstrate; the skills we gain as infants, such as learning to walk or ride a bicycle, require the physical assistance from a parent. We need support to regain motor functions that may be lost due to injuries or stroke. We can also perform incredible feats of synchrony, from carrying large objects together to dancing the tango. These contact tasks allow a simultaneous exchange of information, as both partners feel the forces from one another. In contrast, verbal and gestural communication is serial. Thus, sensorimotor interaction enables rapid and complex coordination, but our understanding of the mechanism behind these interactions is missing (Sebanz *et al.* 2006).



Figure 1.1 | Examples of interpersonal sensorimotor interaction in daily life.

Prior to investigating interpersonal sensorimotor interaction, the cognitive influences inherent in interpersonal interaction need to be controlled. Eye-contact, situational context,

the distance separating partners and past relationships have all been documented to affect the behaviours of interacting partners (Argyle & Dean 1965; Little 1965; Mehrabian 1968; Evans & Howard 1973; Bailenson *et al.* 2001; Kenny & Acitelli 2001). Any change in behaviour arising from cognitive influences may be misinterpreted as a function of sensorimotor interaction, hence eliminating cognitive influences is imperative to isolating and understanding motor coordination between partners. So what are the necessary conditions to eliminate cognitive influences from sensorimotor interaction? This is the first question I address in this thesis, treated in Chapter 2.

Although studies on interpersonal sensorimotor interaction are scarce, some groups have examined it. The first study was conducted by Basdogan *et al.* who examined how two subjects controlled a virtual ring to pass it through a wire (Basdogan *et al.* 2000). They tested pairs in two groups; one was given only visual feedback and the other received both visual and haptic feedback of the partner's actions through a robotic interface; the combined feedback reduced the time taken to complete the task, and instilled a greater feeling of coordination than with vision alone. This suggests that the force plays an important role in coordinating actions. Another pioneering study provided clues to the role the forces played; Reed *et al.* tested pairs who rotated a crank as fast as possible to different target positions; they found partners specialising into roles of always accelerating/decelerating during a reaching movement, which enabled pairs to reach targets faster than individuals who used a crank with half the mass (Reed *et al.* 2006). They proposed that these roles were negotiated through the forces, which might explain how subjects in Basdogan's study completed the task faster. Unfortunately, Reed *et al.* show only a few samples of the force profile from specialised partners. So how do roles emerge, and what purpose do they serve? I will provide experimental evidence, computational modelling and simulations to clarify these queries in Chapter 3.

Besides improving the *speed* of completing a task, a recent study from Prof Burdet's group revealed a surprising benefit in the *control* of a task during interpersonal sensorimotor interaction (Ganesh *et al.* 2014). In this study, pairs of subjects were physically connected by a compliant spring and tracked the same, randomly moving target. The worse partner at the tracking task improved when connected to a better partner, but the better partner also improved whilst interacting with a worse partner! How did partners coordinate their actions to improve from each other, through appropriate roles or another mechanism? How do force and impedance adapt in each partner for better performance? Do the benefits of coordination generalise to larger groups?

My thesis addresses these questions by conducting systematic behavioural experiments, which employ dedicated robotic interfaces to investigate interpersonal sensorimotor interaction; two of these interfaces are shown in Figure 1.2. While the interface used by Reed *et al.* (2006) enabled rigidly coupled interaction in one dimension, the TVINS interface

(Figure 1.2a) facilitates any computer-controlled interaction strategies during planar, two-dimensional arm movements. The Hi5 interface (Figure 1.2b) is reduced to one-dimensional wrist flexion/extension movements, which enables me to systematically study the coordinated behaviours in position, torque and muscle activity (Melendez-Calderon *et al.* 2011). These interfaces provided the means to quantitatively measure the metrics of interaction, such as task performance and adaptations in muscular activity. Furthermore, these devices can manipulate the compliance of the coupling between partners, making it rigid or weak to understand its effect on sensorimotor interaction.

The results from these experiments are modelled to understand and interpret the motor behaviours of the partners. Human neuromechanical control is modelled as an optimal motion planner that executes actions to minimise a cost function, normally the distance between the hand and a target. These simulated subjects receive only visual information of the target position. However, a mechanism of interaction can be tested by providing the simulated subject with additional sensory information, such as a haptic measurement of the interaction force. These simulations yield predictions that can be compared against the experimental data. Models of interaction can provide valuable insights into the underlying adaptations in behaviour, and can be used to interpret the features observed in the experiments. This is the approach I use in Chapters 4 to 6 to understand the mechanism of interpersonal sensorimotor interaction.

This work examines human motor control in the *macroscropic* scale, oversimplifying the physiology of the central nervous system. These simplifications are made to understand the underlying principle of sensorimotor interaction. Thus, the physiological mechanism behind the models I propose here is beyond the scope of this thesis.

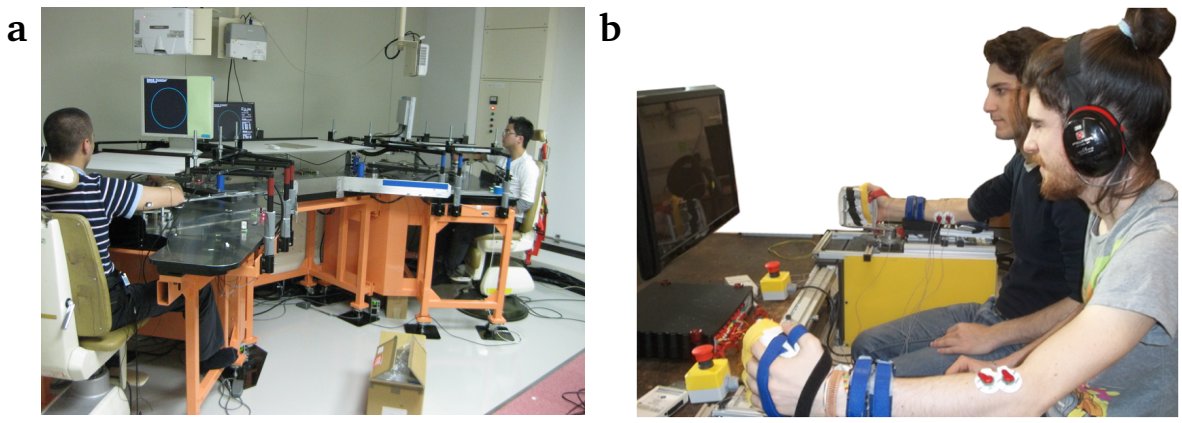


Figure 1.2 | Experimental setups to investigate interpersonal sensorimotor interaction. **a**, Twin planar manipulandums (TVINS) to study interaction in two-dimensional planar arm movements. **b**, Hi5 dual-wrist interface enabling position, force and electromyography readings from wrist flexion/extension (Melendez-Calderon *et al.* 2011).

2
Social bias during interactions

Before undertaking any studies in sensorimotor interaction between humans, it is imperative to understand and control for cognitive influences during interpersonal interaction. Eye-contact, situational context, distance between partners and past relationships have been documented to affect interaction between humans (Argyle & Dean 1965; Mehrabian 1968; Bailenson *et al.* 2001; Evans & Howard 1973; Little 1965; Kenny & Acitelli 2001), but how these factors affect physical interaction between humans is relatively unknown. A previous study conducted by Shergill *et al.* (2003) demonstrated that partners told to exchange the same force would produce increasingly larger forces after every exchange (Shergill *et al.* 2003); this study hypothesised that the escalation of the exchanged force was due to the attenuation of self-produced sensory percepts; just as tickling yourself is difficult, you predict the sensory consequences of your actions and attenuate these sensory stimuli (Blakemore *et al.* 1998). This result is troubling as such behaviour must be mitigated to investigate physical interaction between humans in a controlled environment.

I first tested subjects who knowingly exchanged the same force with a human partner, but with a curtain separating the pair. The behaviour of partners in this curtain condition was compared with individual subjects who reproduced a force from a robot. The act of knowingly exchanging forces with another human did not affect the behaviour. However, when the same experiment was conducted with subjects who sat face-to-face opposite one another, the partners modulated the exchanged force differently to those who were separated by a curtain. Furthermore, men and women modulated their force greater with a partner of the same sex. Altogether, these results show that social cues, not the physiological mechanism of sensory cancellation, affect sensorimotor interaction between humans; moreover, the behavioural influences from social cues can be subdued by blocking the partner's gaze. Thus, obstructing the partner's gaze is necessary to investigate interpersonal sensorimotor interaction.

2.1 Solo force reproduction

First, I examined the accuracy of individuals in sensing and reproducing a force on the right index finger. In *passive reproduction*, the robot sent a force selected randomly between 0.5-10 N for 5 seconds to the finger, which was sandwiched between the robot lever and a mould; this subject reproduced the force by pushing the same lever using the same index finger (see top panel of Figure 2.1a). I also tested *active reproduction* where subjects actively countered the force from the robotic lever such that their finger maintained a fixed position, as depicted in the bottom panel of Figure 2.1a. 4 male and 4 female subjects were tested in both passive and active reproduction modes. All subjects gave their informed consent and all the experiments in this chapter were conducted according to the principles in the Declaration of Helsinki and approved by the Imperial College Research Ethics Committee.

Figure 2.1b shows the force reproduced by subjects as a function of the force from the robot; the diagonal corresponds to perfect reproduction of the force. In active reproduction, the force was accurately reproduced but slightly overestimated for small forces. In passive reproduction, small forces were overestimated but large ones underestimated. Active reproduction may have been more accurate as Golgi tendons could have contributed to sensing the force, whereas in passive reproduction only the mechanoreceptors could be relied upon. The compliance of the lever may have also caused subjects to adopt a quasi-mixture of force and position control, which may have also altered these results (Buchanan & Lloyd 1995) in comparison to previous studies that examined individual force reproduction (Shergill *et al.* 2003; Walsh *et al.* 2011). To avoid unnecessary interaction during sensing, I chose the passive reproduction mode for pairs. Thus, misestimation of the force from the partner should be compared with respect to subjects in the passive reproduction mode, or the *solo condition*.

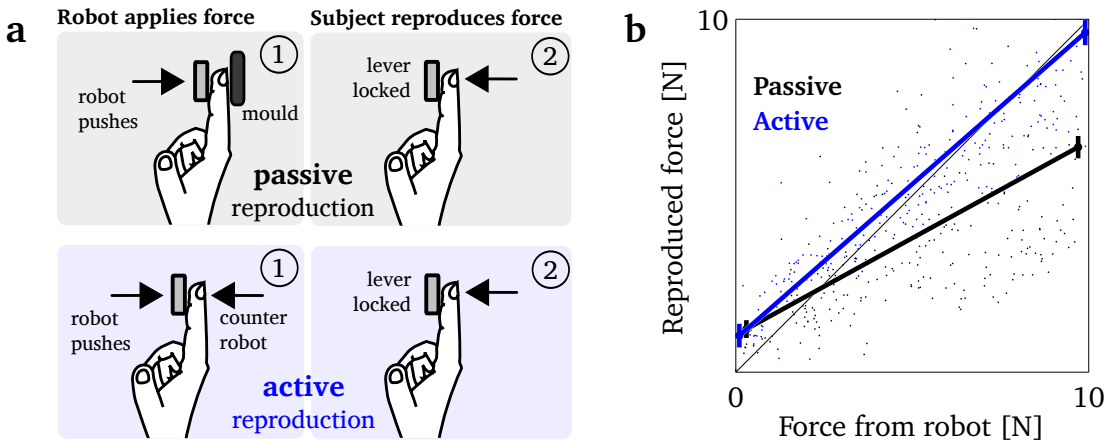


Figure 2.1 | Individual force reproduction. **a**, Protocol force sensing and reproduction using the index finger. In passive reproduction, the finger is sandwiched between a mould and the robot lever to sense the force. In active reproduction, the force from the robotic lever is countered to keep the finger at a fixed position during sensing. The same finger is used to reproduce the force. **b**, Reproduced force versus force from the robot, with linear fits and 95% confidence intervals; subjects were accurate during active sensing but overestimated small forces. In passive reproduction, small forces were overestimated and large ones underestimated.

2.2 Exchanging forces in pairs

How do pairs exchange a force in comparison to individuals? To answer this question, subjects were recruited in pairs and were asked to sense a force applied on the right index finger passively, and then to reproduce this force accurately on to their partner using the same finger (Figure 2.2a). The experiment was composed of 15 trials where each trial was initiated by a robot that picked, in random order, an initial force level of 2, 4, 6, 8 or 10 N, which it applied to the finger of a randomly selected partner. This force was exchanged 6 times between the pair at every trial. Prior to the start of the experiment subjects were introduced to one another to implicitly know the sex of their partner; importantly the partners were strangers to one another to avoid perceptual bias which may have affected the interaction (Little 1965; Evans & Howard 1973; Kenny & Acitelli 2001). Subjects were explicitly told that they would exchange forces with their human partner, and were instructed to "reproduce the same force they had just perceived"; all other instructions came from a computer monitor. Verbal communication was prohibited between the subjects and no feedback concerning the accuracy of the exchange was given to subjects throughout the experiment.

2.2.1 Interaction separated by a curtain

8 male and 8 female pairs were tested by seating them side-by-side with a curtain separating the pair (see Figure 2.2b); subjects knew they exchanged forces with one another but could not see their partner. The data was fit with a linear mixed-effects model of the form

$$F = \beta_{0i} + \beta_{1i} F_{\mathrm{s}} + \varepsilon_i \tag{2.1}$$

where F is the reproduced force, F_s is the sensed force, β_{0i} is the offset, β_{1i} the slope and ε_i the unexplained variance for each subject i. Figure 2.2c shows the results from male and female pairs in the curtain condition. To see if any difference occurred between the curtain and solo conditions, I tested the alternative model of

$$F = \beta_{0i} + \beta_{1i} F_{\mathrm{s}} + \beta_{2i} V + \beta_{3i} \left(F_{\mathrm{s}} \cdot V\right) + \varepsilon_i \tag{2.2}$$

where V is a variable that describes either the gender or the condition of the experiment. A likelihood ratio test between the models of eqs. (2.1) and (2.2) revealed that males and females behaved similarly to one another in the curtain condition (p=0.70, dashed green and orange traces in Figure 2.2c). It also showed that subjects from the curtain condition behaved similarly to solo subjects (p=0.51, black trace of Figure 2.2c). Thus, the act of knowingly

exchanging forces with a human partner did not bias the behaviour of subjects.

As subjects in the curtain and solo conditions reproduced the force similarly, they were used as a single control population of 40 subjects to predict the force reproduced by an unbiased subject. Force misestimation was judged relative to this control such that positive force misestimate corresponded to overestimation, and negative force misestimate to underestimation of the force from the partner. As the force from the partner was not perfectly reproduced, multiple exchanges led to a convergence in the exchanged force; when the robot initiated the exchange with 2, 4, 6, 8 or 10 N, the exchanged force eventually converged to 3N due to the misestimation of the partner's force in unbiased subjects (see Figure 2.3a for unbiased pair). This convergence is visualised by clustering the data in two dimensions, where the horizontal axis corresponds to the mean force from the partner and the vertical axis corresponds to the mean force misestimate. This allows us to visualise the behaviour of each partner, or how they misestimated the force, as a function of the forces they received from the partner. The

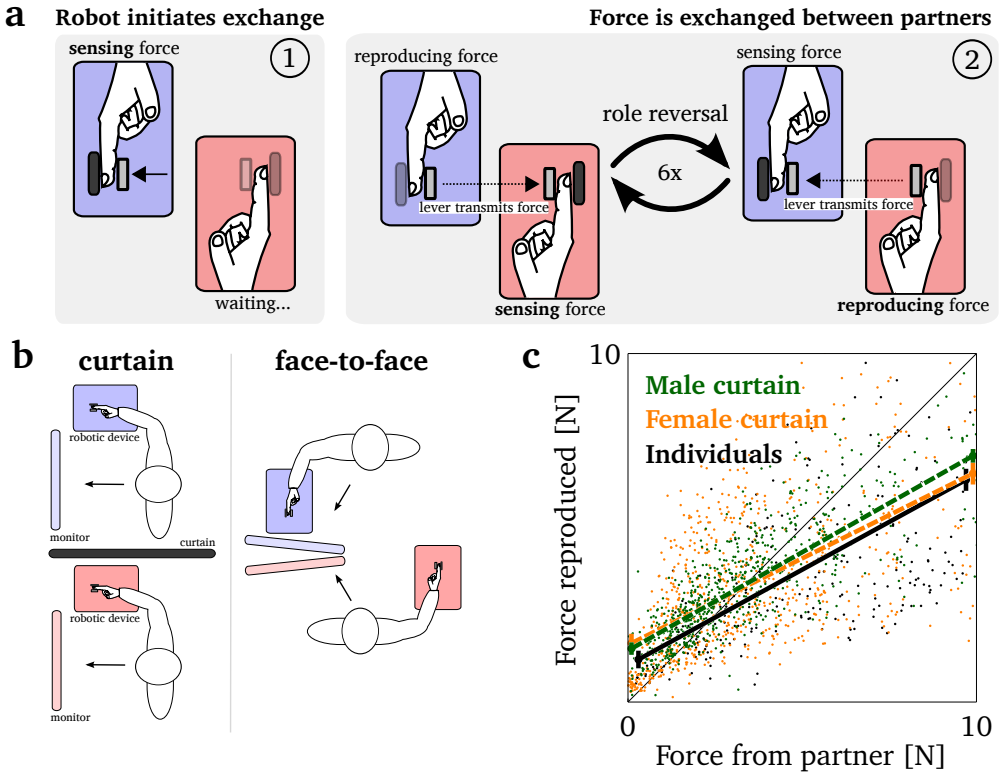


Figure 2.2 | Experimental protocol and setup to exchange forces. **a**, Top-down view of the experimental setup for force sensing and reproduction using the index finger. First, a constant force from the robot is sent to the index finger of a randomly selected partner; then, this person sends the force to the other partner using the robotic lever. This process is repeated for 6 exchanges. **b**, Two experimental conditions where partners were separated by a curtain and sat side-by-side, or sat opposite one another face-to-face. **c**, Force reproduced as a function of the force from the partner in the curtain condition for male and female pairs, and the solo condition.

cluster in the top panel of Figure 2.3b signifies unbiased behaviour, which is what we would expect from partners who were not biased by the interaction.

The data was clustered using a Gaussian mixture distribution, and the number of clusters was determined by minimising the Akaike information criterion. Outliers were visually excluded from clusters. One-way between groups ANOVA using post hoc Games-Howell was conducted to compare differences in all conditions between the horizontal (force from partner) and vertical (force misestimate) means of the centroid of the clusters.

2.2.2 Face-to-face interaction

If knowingly interacting with a human partner does not bias your behaviour, what caused the escalating behaviour in Shergill's study (Shergill *et al.* 2003)? In their study, the partners exchanged forces whilst sitting face-to-face; studies have shown that a person's gaze can alter your behaviour (Argyle & Dean 1965; Mehrabian 1968). To test if the partner's gaze would have an effect on interaction behaviour, an additional 8 male and 8 female pairs were recruited and exchanged forces in the face-to-face condition (Figure 2.2b). Subjects sat opposite one another such that the partner was visible at all times during the experiment; the distance between the partners was kept the same as in the curtain condition. Would the partner's gaze affect the behaviour of subjects in the same task?

Indeed, the behaviour in both males and females changed in response to the partner's gaze. In female pairs, the cluster of unbiased behaviour disappeared altogether; instead, another cluster of inhibitory females emerged, who consistently underestimated the force from the partner (orange cluster in middle panel of Figure 2.3b). Figure 2.3a shows a representative inhibited female pair; forces were consistently underestimated such that the exchanged force rapidly decreased to 1 N. This inhibited behaviour may have occurred to avoid conflict with the partner (Maccoby 1998) as overestimating the force from the partner may be perceived as aggression. Only one female subject did not underestimate the force from the partner; this pair (2 orange crosses in the middle panel of Figure 2.3b) were clear outliers from the cluster, but it is evident from the force misestimate that the other partner was inhibited by the interaction. Out of 16 female subjects, 15 were inhibited by face-to-face interaction.

Male pairs who interacted face-to-face were also inhibited by the partner's gaze, but a significant proportion were unfazed by it; the centroid of the inhibited male cluster (n=8) was similar to the inhibited female cluster (p_{horz}=1.0, p_{vert}=.99) and the unbiased male cluster (n=6) was similar to the unbiased subjects of the curtain condition (p_{horz}=1.0, p_{vert}=1.0). One male pair showed escalating behaviour which belonged to neither inhibited nor unbiased clusters (two green crosses in middle panel of Figure 2.3b). Both partners in the escalating pair consistently overestimated the force from one another; the force increased after every

exchange (Figure 2.3a). This escalating behaviour, which was not observed in females, may be related to rough play which occurs more readily in males (Goldstein, 2003). The emergence of the inhibited and escalating behaviour in males, and the overwhelming majority of female pairs who were inhibited by the face-to-face interaction, suggests that the partner's gaze affects males and females in different ways. In comparison to males, females were overwhelmingly inhibited by the partner's gaze (15/16=94% of females and 8/16=50% of males); therefore, the subject's sex may determine the effect of the partner's gaze.

2.2.3 Face-to-face interaction with the opposite sex

If the effect of the partner's gaze is ingrained in the subject's gender, one would expect more inhibited females than inhibited males when pairs of mixed-sex interact face-to-face. To test this hypothesis, 16 mixed-sex pairs were tested in the face-to-face condition to see if more females than males would exhibit inhibited behaviour.

Contrary to this hypothesis, males and females in mixed-sex pairs behaved similarly. Inhibitory and unbiased behaviours were observed as two clusters (see bottom panel of Figure 2.3b), and males and females were equally likely to be inhibited (10 males and 9 females) or unbiased (5 males and 6 females) by the interaction. Furthermore, the amount of inhibition in both males and females in mixed-sex pairs was less than their counterparts in same-sex pairs; inhibited males in male pairs underestimated the partner's force by 0.43 ± 0.12 N ($p_{vert}=.013$) and inhibited females in female pairs by 0.48 ± 0.12 N ($p_{vert}=.0037$) more than their mixed-sex counterparts. As a result, the average exchanged force between the partners was smaller in same-sex inhibited pairs than mixed-sex ones (0.70 ± 0.15 N, 0.68 ± 0.13 N for male and female inhibited pairs respectively). The centroid of the unbiased cluster of mixed-sex pairs was similar to the unbiased cluster in the control and in same-sex male pairs ($p_{horz}=.23$, $p_{vert}=.96$ and $p_{horz}=.71$, $p_{vert}=.99$ respectively). These results suggest that inhibited mixed-sex partners modulate the force differently depending on the partner's sex.

One male subject showed escalating behaviour (green cross in bottom panel of Figure 2.3b), but his female partner returned the force relatively accurately like an unbiased partner; the behaviour of this mixed-sex pair is different from the escalating same-sex male pair, where both partners were escalating.

Although both males and females were more inhibited by a partner of the same sex, the likelihood of becoming inhibited remained similar in males, but changed with the partner's sex in females; females were more likely to display inhibited behaviour when interacting with another female partner (15/16=94% of the population) than with a male partner (9/16=56%). On the other hand, males were inhibited, unbiased and escalated in approximately the same proportions when interacting with a male (n=8:6:2) or female partner (n=10:5:1). In an-

other study on personal distance, females were more sensitive to a partner's gaze and changed their behaviour more readily than males (Bailenson *et al.* 2001); this suggests that females are selectively more sensitive to the contextual cue of the partner's sex in comparison to males (Little 1965), which may also relate to the greater ability in females to empathise with others (Baron-Cohen 2002).

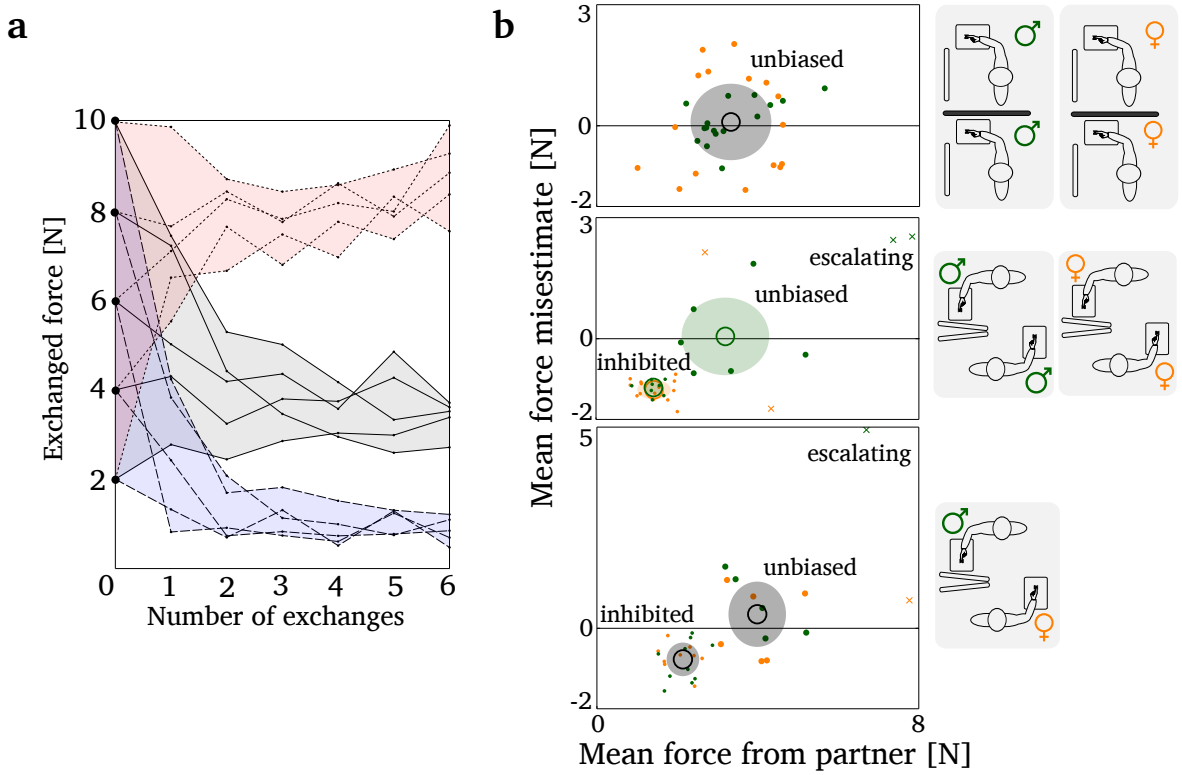


Figure 2.3 | Force reproduction in pairs. **a**, The exchanged force as a function of the number of exchanges from 3 representative pairs who displayed unbiased, inhibited and escalating behaviour. The robot initialised the exchange with 2, 4, 6, 8 and 10N. If the exchanged force converged to roughly 3N, pairs were unbiased by the interaction and behaved similar to individuals. Inhibited pairs underestimated the force from the partner so the exchanged force rapidly decreased. Escalating pairs overestimated the force so the exchanged force increased. **b**, Points are from each partner, circle is the mean and shaded circle is the standard error of the centroid. *Top panel*: Male and female same-sex pairs from the curtain condition, who had similar centroids to each other and the solo condition. *Middle panel*: Male and female pairs from the face-to-face condition. Two outlier pairs, one male and one female, were removed from the clustering. All female and some male partners (n=9) displayed inhibited behaviour which was characterised by a cluster with little variability where small forces were exchanged. In males, another cluster of unbiased subjects (n=6) was observed with a centroid that was not significantly different from the control. *Bottom panel*: Mixed-sex pairs from the face-to-face condition. Inhibited and unbiased behaviours were observed in this condition, with equal populations of male and females in both. However, the inhibited partners were less inhibited than their counterparts in same-sex pairs.

2.3 Discussion

The study of Shergill *et al.* (2003) painted a picture of force escalation during sensorimotor interaction where the attenuation of sensory percepts from self-generated actions caused subjects to overestimate the force from the partner. However, a distinctly different picture emerges from the experiments I conducted. Although knowingly interacting with another human partner did not cause subjects to behave differently to individuals who did the task alone, the partner's gaze caused significant changes in behaviour.

When males and females interacted with a partner of the same sex, some subjects displayed inhibited behaviour, which was characterised by underestimation of the partner's force; this was not observed when partners were separated by a curtain. Furthermore, the partner's sex modulated this inhibited behaviour such that both males and females were more inhibited by a partner of the same, rather than the opposite, sex. Moreover, females' interaction behaviour was more sensitive to the partner's sex than males'; the likelihood of females being inhibited by a female partner was higher than with a male partner. In contrast, males were inhibited, unbiased and escalated in equal proportions, regardless of the partner's sex. Escalating behaviour was observed in only one male pair who interacted face-to-face, a comparatively smaller number than in Shergill's study. However, we did observe face-to-face pairs who behaved similarly to individuals as their study did.

Precisely why sitting face-to-face changes a partner's behaviour is still unclear. Third-person neuroscience has documented the effects of gaze on a person's actions (Bailenson *et al.* 2001), but first- and second-person neuroscience approaches are necessary to reveal changes in the mental state and the mental representation of the partner, which may explain why face-to-face tit-for-tat exchanges in force is different to individual force estimation (Vogeley & Fink 2003; Schilbach *et al.* 2013).

Regardless of the social mechanism involved, these results show that social cues modulate interaction behaviour, and that blocking the partner's gaze is a necessary condition to investigating interpersonal sensorimotor interaction.

3
Dyadic reaching

In the previous Chapter 2, I demonstrated that a curtain is necessary to eliminate social bias during sensorimotor interaction. Using this controlled experimental environment I can examine joint interaction scenarios, like carrying a table together. The pioneering study of (Reed & Peshkin 2008) examined dyads rotating a single crank to different angles. They found signs of dyads specialising into roles where one always accelerated and the other decelerated during the movement. Unfortunately, we do not know if these roles remained persistent or how they emerged in the task. Their protocol of doing movement after movement consecutively, and as fast as possible, is not a sufficiently controlled testing environment; this is because partners applied interaction forces (the force between the partners) of varying magnitude at the beginning of each new movement, which may alter the dyadic reaching movement (Reed & Peshkin 2008). For this reason, I decided to revisit their dyadic reaching task with a controlled experimental protocol to examine one movement at a time, with and without the presence of interaction forces at the start of the movement.

The results of my dyadic reaching experiment indicate that dyadic behaviour is highly consistent; there is no trial-by-trial adaptation of the interaction torque profiles, which suggests little to no negotiation of roles between partners, which was purported to arise during the trials by Reed *et al.*. However, the same dyad occasionally behaved differently with and without interaction forces at the beginning of the movement; the direction of the initial forces was of importance, as some pairs did not switch symmetrically when alternating roles between pushing towards the target and pulling away before jointly reaching. Furthermore, between dyads there were significant differences in behaviour. How do we explain the differences in behaviour with and without forces at the start of the movement, the large interaction forces observed even at the end of the movement and the large variability between the dyads?

I use a simple model of two humans planning and executing their movements independently, who feel the effects of the interaction force. The interaction torque profiles are pre-

dicted by connecting partners with different motion plans, which is composed of the goal position and the state cost function. The redundancy of the cursor, which was the average position of the partners, allowed partners to reach to different goals. The goal position was biased by which direction subjects were applying the force before the movement; subjects pushing towards the target reached for a goal that overshot it, and those who pulled away from the target undershot it. The state cost function appeared to be ingrained in each person as the timing and sequence of the torque profiles did not change during the entire experiment.

These simulations imply that dyadic reaching movements are explained by a mechanism of *co-activity*, where both partners execute movements independently and ignore the interaction forces (Jarrassé *et al.* 2012). In effect, dyadic reaching is not a suitable task to examine interaction between partners. Interaction implies adaptation that occurs over time whilst experiencing the partner's movements. The consistency of the interaction force profiles within dyads gives no indication of such adaptation, and no form of effort optimisation either. What is needed is an interaction task that specifically demands in adaptations in both partners, which will be explored in the next Chapter 4.

3.1 Dyadic reaching

16 subjects, grouped in 4 male and 4 female pairs, were recruited to make point-to-point reaching movements in one dimension using the Hi5 dual robotic wrist interface (Figure 1.2b); all subjects gave their informed consent and the experiment was conducted according to the principles in the Declaration of Helsinki and approved by the Imperial College Research Ethics Committee. The movement started with slight extension of the wrist and ended with medium flexion (-10 deg to 30 deg with a movement amplitude of 40 deg). Subjects were instructed to wait for a countdown of 3 seconds before initiating the reaching movement. This countdown was introduced to synchronise the start of each movement, otherwise the results may be affected by different reaction times of the partners.

The experiment was composed of trials in solo, coupled and push-pull blocks. In solo blocks, partners did the reaching task alone (see 'solo' panel Figure 3.1). Movement times were constrained between 900-1200 ms by a message on a monitor that indicated if the movement was "too fast" or "too slow". In coupled blocks, partners were connected through a rigid virtual connection of stiffness $0.4\,\mathrm{Nm\,deg^{-1}}$ (assuming the wrist is 10 cm from the pivot, this is $\approx 2300\,\mathrm{Nm^{-1}}$), and were given visual feedback of the position-averaged cursor of both partners (see 'coupled' panel Figure 3.1). The movement time was not constrained in the coupled blocks in case it disturbed the movement during interaction.

To examine the effect of interaction forces at the beginning of the movement, I also tested pairs in push-pull blocks. Here, partners were rigidly connected as in the coupled block but were required to produce 0.7 Nm of interaction torque before starting the movement (see 'push-pull' panel Figure 3.1); the interaction torque was produced by one partner pushing towards and the other pulling away from the target. These roles were reversed in the next push-pull block (after a solo block washout). Like in the coupled blocks, the movement times were not constrained in the push-pull blocks.

The top panel of Figure 3.1 shows the block design of the experiment; subjects first completed a solo block, which served as a training block, then one coupled block followed by a washout solo block. Lastly, two pairs of push-pull and solo blocks were completed. In the second push-pull block, the roles of pushing towards and pulling away from the target were reversed. The roles were constant within each block.

To examine interaction behaviour, I look at the interaction torque as a function of time, or the interaction torque profile. This gives an indication of the forces applied by the dyad during the movement. If changes occur in the profile, it can be observed immediately. I also show the mean interaction torque profile from all trials in each coupled and push-pull block to see the overall behaviour of the dyad. Due to the nature of the rigid connection, the force

applied by each partner will result in a nearly symmetric torque profile.

Would there be any difference between the coupled and push-pull blocks due to the interaction torque at the start of the movement? And would partners exhibit symmetrically opposite behaviour when they switched roles in the next push-pull block, or would they adapt to the new condition?

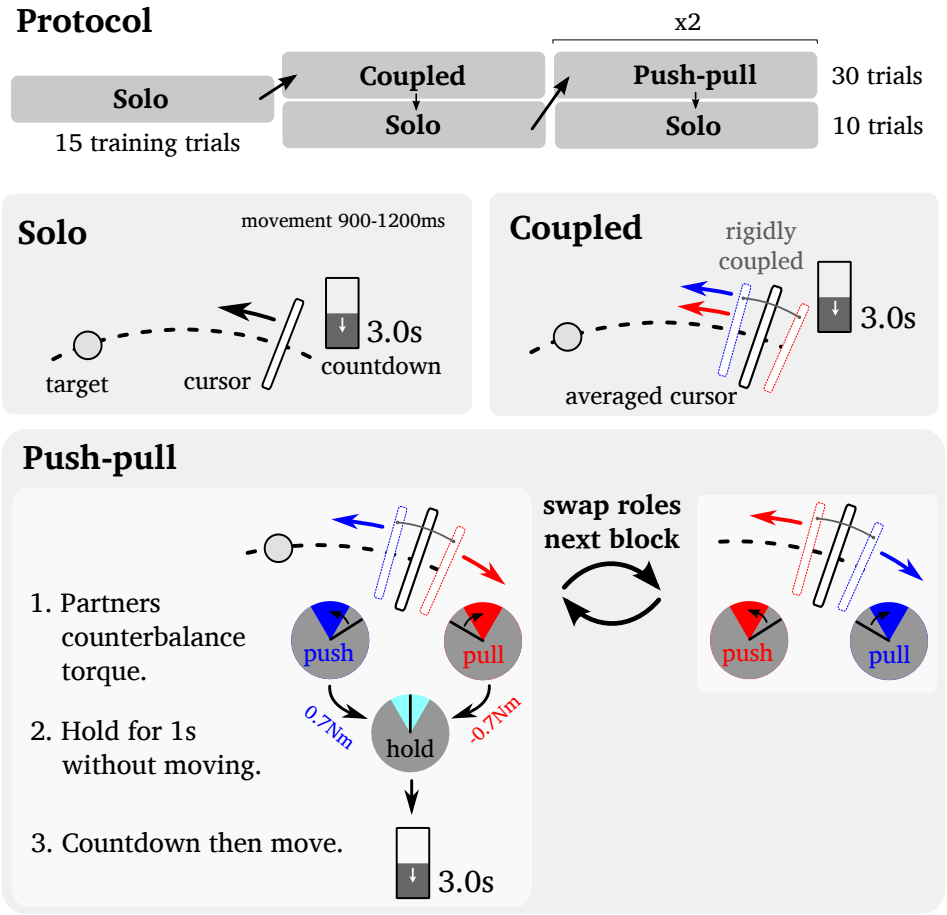


Figure 3.1 | Dyadic reaching protocol. The experiment was split into three types of blocks; in the solo block, partners are disconnected and reach the target alone. Feedback was given to the subject if the movement is faster than 900 ms or slower than 1200 ms. A 3 second continuous countdown bar was used to synchronise the start of the movement of both subjects. In both the coupled and push-pull blocks the partners were rigidly connected by a virtual spring; in the coupled block, partners began the movement with no interaction torque. In push-pull blocks, partners were required to push and pull to counterbalance one another, producing an interaction torque of $\tau_{\text{int}} = 0.7\,\text{Nm}$ whilst keeping the averaged cursor at the initial position θ_0 before initiating the movement. The roles of pushing and pulling away from the target were swapped in the next push-pull block in the experiment.

Results

Figure 3.2 shows the interaction torque profiles from 8 dyads in the coupled and push-pull blocks; the first column corresponds to the coupled block and the second and third columns are from the push-pull blocks. Two vertical lines indicate the peak velocity and the end of the movement (dotted and solid lines respectively). From here onwards, I refer to each dyad by their plot label (for example, Figure 3.2a corresponds to dyad *a*).

Let's first take an overall look at the interaction torques in both coupled and push-pull blocks. Remarkably, the interaction torque profiles are very consistent and deviate little from the mean. Dyads *g* and *h* are different from the others, but only due to high variability as the underlying mean behaviour is similar for each movement within each block. From the perspective of negotiation and role emergence, this makes little sense. Why are the movements so consistent every trial? How could the roles be negotiated so quickly, almost as if they were predetermined before the interaction? The consistency of the interaction torque profiles suggests little to no adaptation in behaviour, and is more indicative of *co-activity* where partners ignore the forces and execute the movement independently (Jarrassé *et al.* 2012).

Now let me examine the behaviour in more detail. Looking at the coupled blocks only (left column), two dyads *d* and *e* had negligible interaction torque. The interaction torque switched direction at the peak velocity of the movement in dyad *a*; this corresponds to a type of specialisation observed in Reed *et al.*'s (2008) study where one partner pushed towards the target at first but pulled away at the end of the movement and the other partner did the opposite. The interaction torque in the remaining five dyads *b, c, f, g* and *h* did not switch signs, corresponding to dyads who accelerated or decelerated during the entire movement; this is another type of specialisation observed in Reed's study. Although the interaction torque never changed direction in these five dyads, the interaction torque profiles are distinct; these dyads cannot be classified as a single "always accelerating/decelerating" group. Instead, the subtle differences in the interaction torques serve as clues to the underlying behaviour of the partners. However, at this stage I cannot interpret the differences in behaviour between the dyads. It is evident that in dyads *a, b, f, g* and *h* a constant interaction torque was observed even after the end of the movement, as reported in Reed's study.

Now let's look at the push-pull blocks (middle and right columns). Dyads *a, b, c* and *d* exhibited symmetry in the interaction torques with relatively symmetrical role reversal. Other dyads *e* and *f* showed signs of asymmetry when switching push-pull roles but partners always applied interaction torques in the same direction during the movement. A change in the direction of the interaction torque was observed in the remaining dyads *g* and *h* such that the partner pushing towards the target at the beginning was pulling away mid-way and vice versa for the other partner. By examining dyad *d*, the difference between the coupled and push-pull

dyad blocks is evident; in the coupled block, negligible interaction torque was observed but as soon as the dyad was forced to produce some interaction torque at the start of the movement in the push-pull blocks, the dyad began to produce significant interaction torque, even at the end of the movement.

This is the furthest I can get with qualitative interpretations alone. I can only speculate how these specialised behaviours may be occurring; extrapolating the change in each partner's behaviour from these interaction torques is an immensely difficult task, which is probably the reason why previous studies looked so sparsely at such data. However, I can try to model the interaction torques using a simple framework to gain some understanding of the underlying behaviours. In the next section, I build a simple controller that models the reaching movement of a single subject and observe what happens when two simulated subjects are connected rigidly to one another. Is it possible to capture the rich behaviours in Figure 3.2 with such a simple framework?

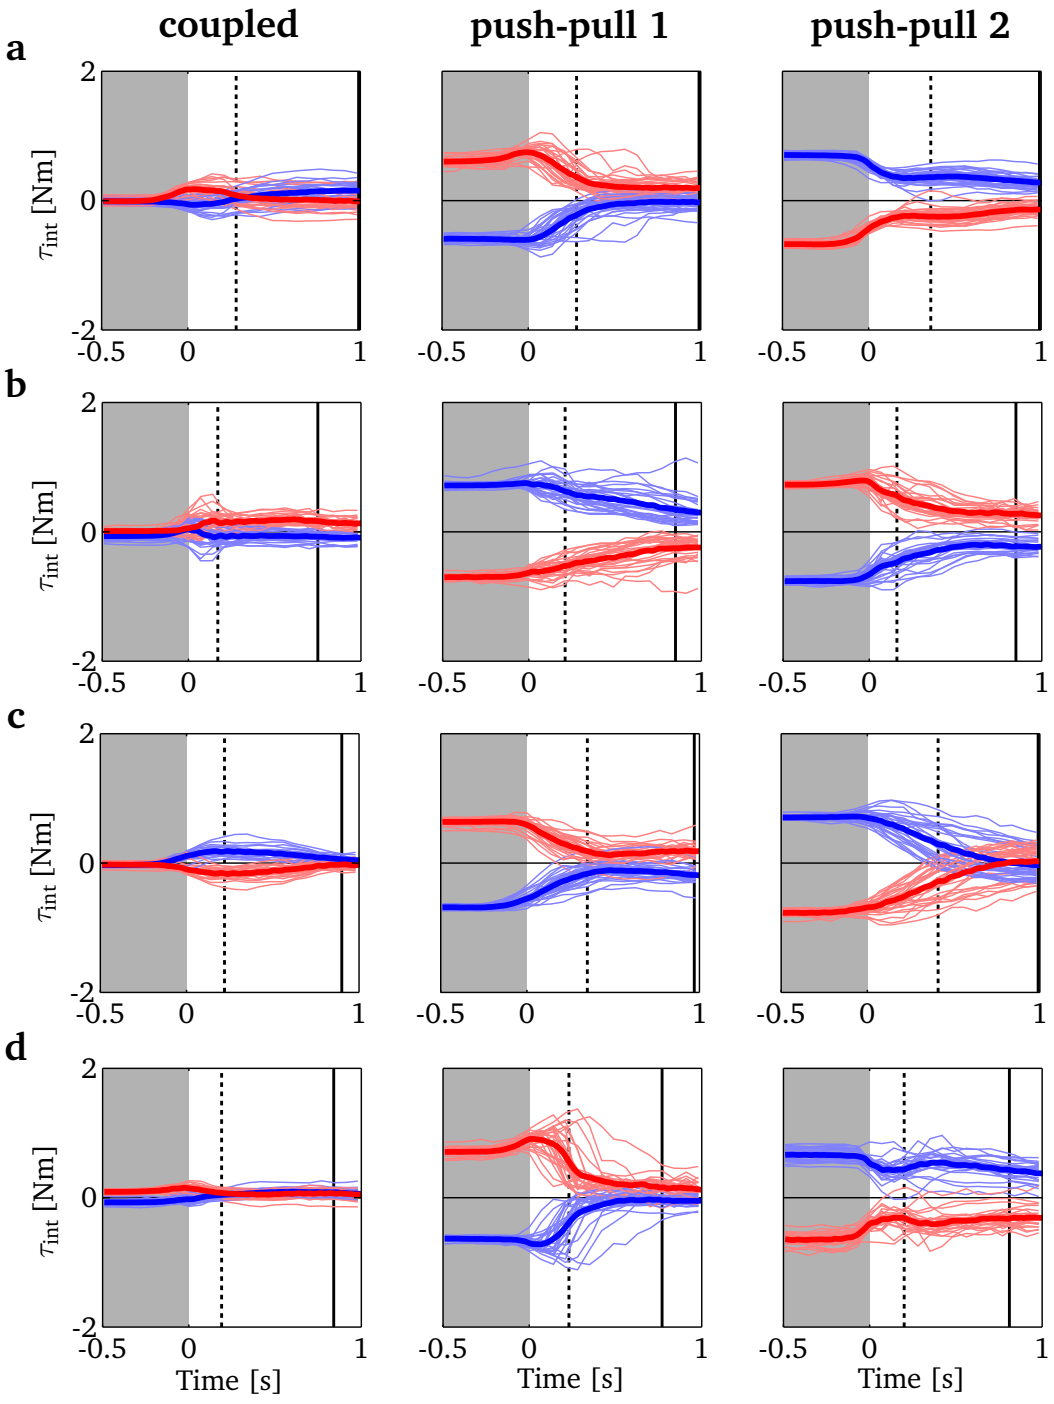


Figure 3.2 (continues next page)

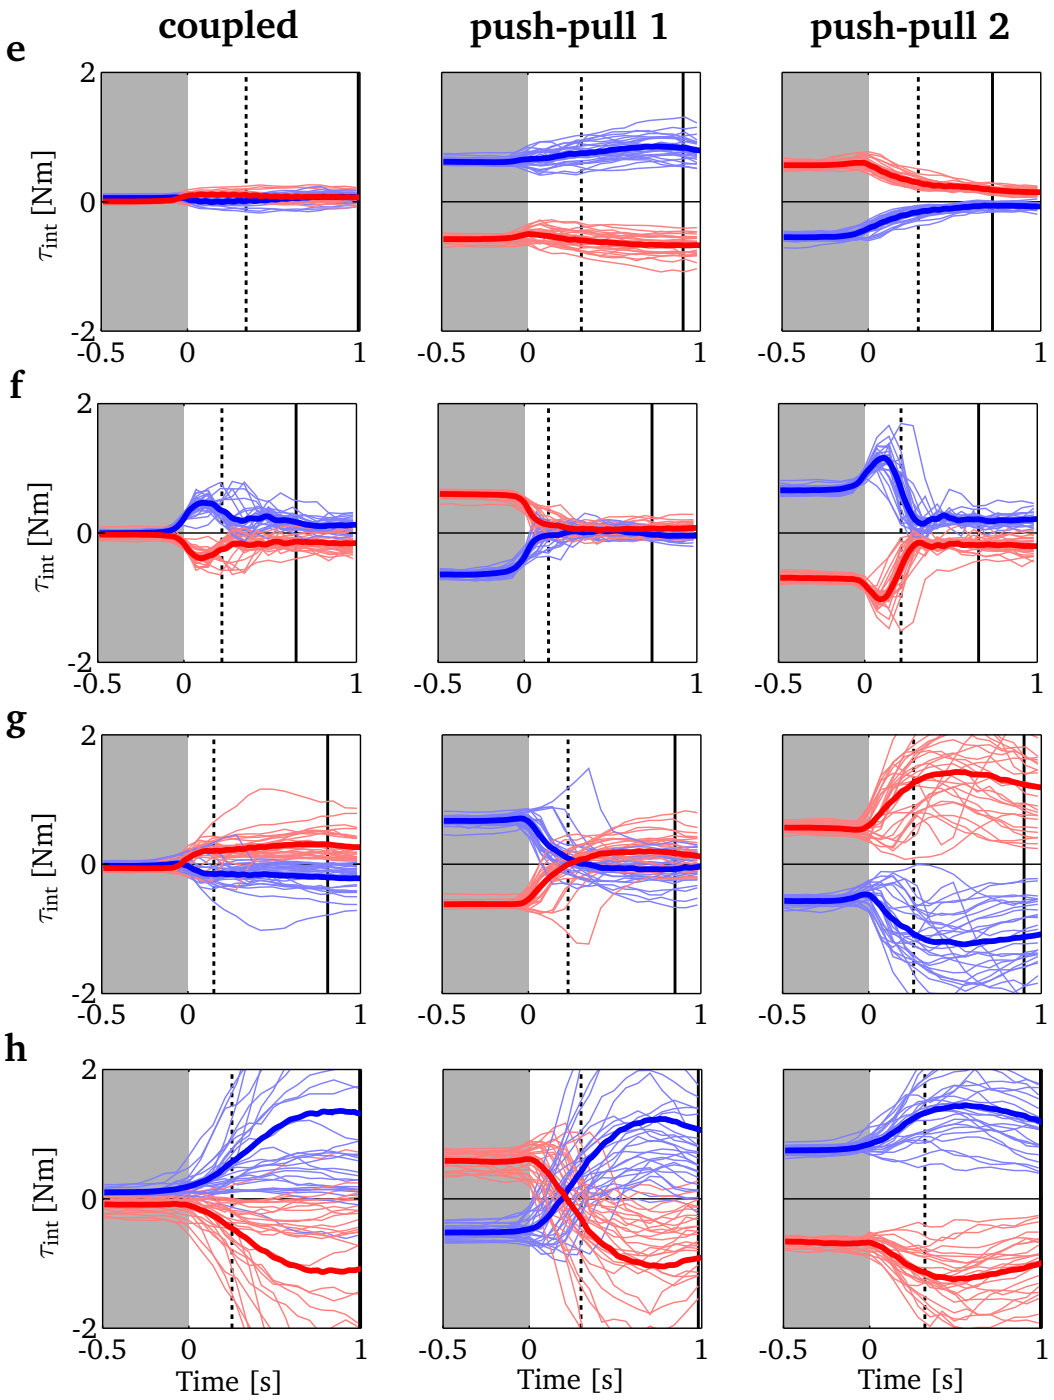


Figure 3.2 | Data from coupled (left column) and push-pull blocks (middle and right columns). Blue and red correspond to partners A and B. Light traces show all trials within each block and the thick trace is the mean of all trials. The gray region is before the start of the movement, which begins at 0s; the interaction torques before the movement are visible in this region in the push-pull blocks.

3.2 Dyadic reaching model

Let's begin with a model of a single subject making a reaching movement in 1 dimension. The state of the human wrist

$$\mathbf{x} = \left[\theta,\, \dot{\theta},\, \tau,\, f\right]^T \tag{3.1}$$

is composed of the angle, angular velocity, torque of the wrist and an auxiliary variable that evolves in discrete time through the state space model

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}u_k + \mathbf{C}\left(\mathbf{x}_k^p - \mathbf{x}_k\right) \tag{3.2}$$

where $\mathbf{x}_k^p$ is the state of the partner, k is the time index,

$$\mathbf{A} \equiv \begin{bmatrix} 1 & dt & 0 & 0 \\ 0 & 1 & dt/I & 0 \\ 0 & 0 & 1 - dt/\tau_u & dt/\tau_u \\ 0 & 0 & 0 & 1 - dt/\tau_u \end{bmatrix},\ \mathbf{B} \equiv \begin{bmatrix} 0 \\ 0 \\ 0 \\ dt/\tau_u \end{bmatrix},\ \mathbf{C} \equiv \begin{bmatrix} 0 & 0 & 0 & 0 \\ K\,dt/I & D\,dt/I & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{3.3}$$

with the moment of inertia of the wrist I, time step dt and the stiffness and damping of the spring K and D that connects partners. The musculoskeletal system of the wrist is modelled using a second-order muscle-like filter with a time constant of $\tau_u = 40\,\text{ms}$, which is why the auxiliary variable, f, is needed (Todorov 2005). The connection matrix $\mathbf{C} \equiv \mathbf{0}$ if the partners are disconnected, as in the solo block. The motor command sent to control the wrist is

$$u_k = -\mathbf{L}\left(\mathbf{x}_k - \mathbf{g}\right) \tag{3.4}$$

where the goal vector $\mathbf{g}$ is

$$\mathbf{g} = [\theta_g, 0, 0, 0]^T \quad \text{during reaching} \tag{3.5}$$

$$\mathbf{g} = [\theta_0, 0, \tau_0, 0]^T \quad \text{before reaching in push-pull block} \tag{3.6}$$

where θ_g and θ_0 are the goal and initial angles of the movement, respectively. τ_0 is the interaction torque imposed at the start of the movement in the push-pull blocks. The controller gain $\mathbf{L}$ is computed by minimising the cost

$$E\left[\sum_{k=0}^{\infty} \mathbf{x}_k^T \mathbf{Q}\, \mathbf{x}_k + u_k^T R u_k\right] \tag{3.7}$$

where the matrix $\mathbf{Q}$ is positive semi-definite and R is positive. This infinite horizon linear quadratic regulator (LQR) is found by solving the Riccati equation (Stengel 2012)

$$\mathbf{S} = \mathbf{Q} + \mathbf{A}^T \left[\mathbf{S} - \mathbf{SB}(R + \mathbf{B}^T\mathbf{SB})^{-1}\mathbf{B}^T\mathbf{S}\right] \mathbf{A} \tag{3.8}$$

and the iterative solution of eq. (3.8) yields the optimal gain

$$\mathbf{L} = (R + \mathbf{B}^T\mathbf{SB})^{-1}\mathbf{B}^T\mathbf{SA}\,. \tag{3.9}$$

A pair of the human-like controllers above were simulated alongside one another, only interacting through the connection matrix $\mathbf{C}$. In other words, each partner was only affected by the interaction forces alone, and did not plan a joint movement with another partner, yielding *co-activity* (Jarrassé *et al.* 2012).

I assume that subjects have perfect knowledge of the noiseless environment, as the purpose of this exercise is to describe complex behaviours arising from human-like controllers. For all simulations, the values of the parameters were set to

$$K = 0.4\,\text{Nm}\,\text{deg}^{-1},\ D = 0.05\,\text{Nms}\,\text{deg}^{-1},\ R = 10^{-2},\ dt = 0.01\,\text{s}\,,\ I = 0.1\,\text{kgm}^2. \tag{3.10}$$

K is from the experiment and damping D is added to model the natural damping properties of the wrist. The inertia I was set such that the torques in the controller matched the torque-normalised muscle activity during the solo blocks (peak acceleration/deceleration between 0.5-1 Nm). R was set to match the speed of the movement observed in the data. The state cost and goal were set to

$$\mathbf{Q} = \text{diag}\,[1, 0.04, 0, 0]\ , \quad \theta_g = 30\,\text{deg} \equiv \frac{30\pi}{180}\,\text{rad} \tag{3.11}$$

to match the velocity profile of movements (approximately bell-shaped but with a longer tail at the end) and the goal was set to the same angle as in the experiment.

Using this model, I ran simulations of a coupled block where two rigidly connected partners made joint reaching movements. These simulations, results of which are shown in Figure 3.3, predicted three key features in the interaction torque profiles; on the shortest timescale, before the peak velocity of the movement, a difference in the state cost in the position term $\mathbf{Q}_{1,1}$ led to a faster planned movement in one partner, which resulted in a short bump in the interaction torque that decayed rapidly (see Figure 3.3a). On a medium timescale, if one partner had a larger weighting on the velocity term in the state cost $\mathbf{Q}_{2,2}$ (i.e. mostly brake), a flat hump in the interaction torque was observed (Figure 3.3b). Lastly, on the largest timescale, a disagreement in the goal θ_g resulted in an undecaying, constant

interaction torque during and after the movement (Figure 3.3c).

I can use the interaction torque profiles from the simulations as templates to interpret the data in Figure 3.2. By visually examining similarities between these templates and the data, it becomes possible to analyse the cause(s) of the unique interaction torque profiles in the data. Note that $\mathbf{Q}$ and R are related such that changing one is equivalent to changing the other, thus only one need be considered as an independent variable.

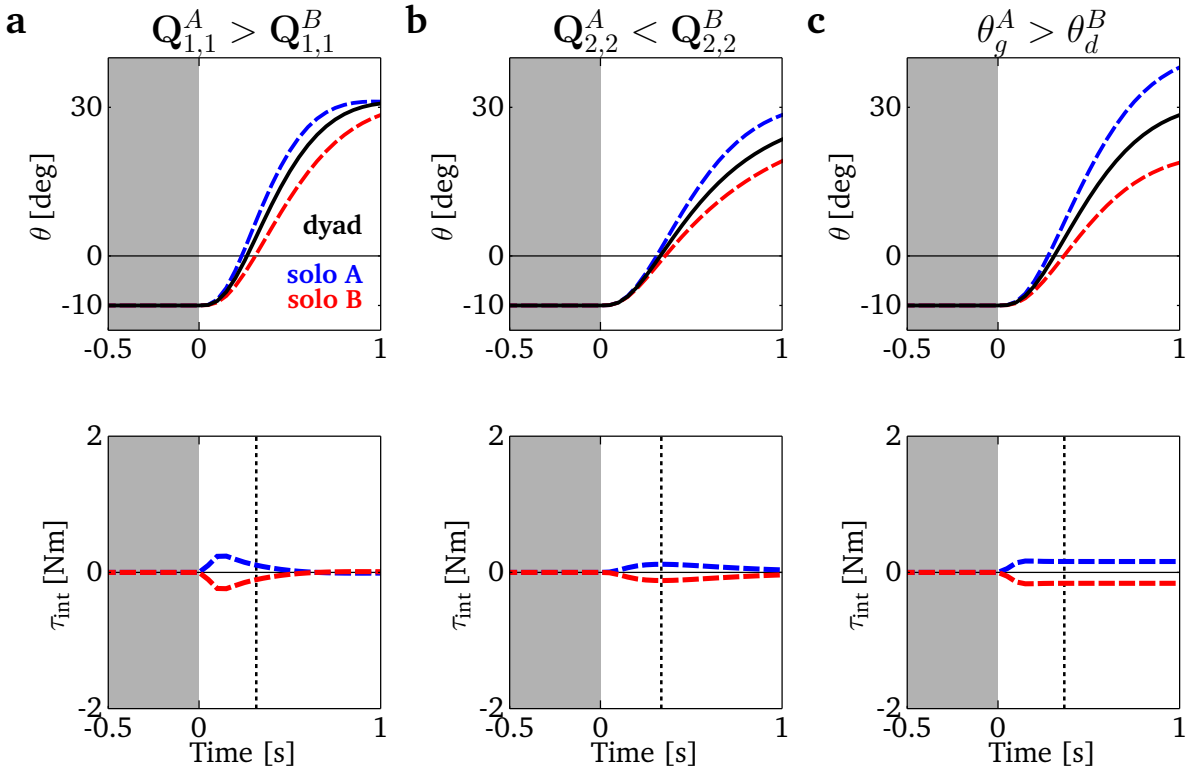


Figure 3.3 | Interaction torque profile templates from simulations of dyadic reaching. The reaching model is implemented with and without connection to observe the predicted interaction torque with differently planned movements. **a**, The position term of the state cost of partner A is larger in comparison to partner B ($\mathbf{Q}^A_{1,1} = 4$, $\mathbf{Q}^B_{1,1} = 1$), resulting in a faster movement during solo reaching (lighter dashed blue and red traces). When these two partners are connected, the faster moving partner A caused a short bump in the interaction torque which rapidly decayed in the first half of the movement. **b**, Partner B changes their movement plan with higher relative weighting on reducing the velocity, i.e. braking ($\mathbf{Q}^A_{2,2} = 0.04$, $\mathbf{Q}^B_{2,2} = 0.4$). This results in a flat hump in the interaction torque which decays over the whole movement. **c**, Due to the redundancy in the cursor angle in coupled and push-pull blocks, the partners can have different desired targets yet still fulfil the task as long as the averaged cursor angle is 30 deg at the end of the movement. Such disagreement in the desired target ($\theta^A_g = \theta_g + 10$ deg and $\theta^B_g = \theta_g - 10$ deg) results in a constant, undecaying interaction torque over the whole movement.

Model-based interpretation of interaction

The easiest trait to pick out from the data is the undecaying, constant interaction torque that remains at the end of the movement (Figure 3.3c). Recall that dyads *a, b, f, g* and *h* showed this feature, which became more prominent in the push-pull block. This implies that, due to the redundancy of the cursor position in connected blocks, partners attempted to reach towards different goals than the original target. Furthermore, the partner's role in the push-pull block determined whether their goal overshot or undershot the target. Partners who pushed towards the target modified their goal to overshoot the target, and partners who pulled away tried to undershoot the target. Dyads *g* and *h* were an exception to this rule, so why did these two dyads behave differently to the others?

These two dyads may have behaved differently due to one partner consistently reaching for a goal that overshot the target. Figure 3.4a shows the data from one of these dyads, and Figure 3.4b is the simulation of this dyad. In this case, partner B (red) consistently reached towards a goal that overshot the target by $5\,\text{deg}$. In other dyads, the simulations show that partners' state cost depended on the direction they were pushing/pulling at the beginning of the movement. Figure 3.4c is an example of such a dyad; the short "bump" in the interaction torque is a sign that partner A (blue trace) had a larger position term cost than partner B (red trace). However, simulating the dyadic movement using the same state costs for partners A and B in all connection blocks (coupled and both push-pull) results in a poor prediction of the interaction torque in the first push-pull block when partner A pulled away and partner B pushed towards the target (Figure 3.4d). The only way to reproduce the data is to reduce the position term cost of partner A, implying that pulling away from the target reduced their position cost.

The simple model of two reaching controllers has provided valuable insight into the motion plans of the partners and how it was responsible for the interaction torque profiles observed in the data. These simulations also indicate one more thing: that subjects' state cost and goal did not change *within* a block, i.e. there was no adaptation in the motion plan within a block. Once again, this implies a mechanism of co-actitivy in the dyadic reaching task (Jarrassé *et al.* 2012).

A disadvantage of the rigid connection is the inherent confound in the analysis and interpretation of the torque profiles; how can I distinguish between one person increasing their position cost, and the other partner reducing it? Clearly, the rigidly connected dyadic reaching task is not a good one to examine interpersonal sensorimotor interaction.

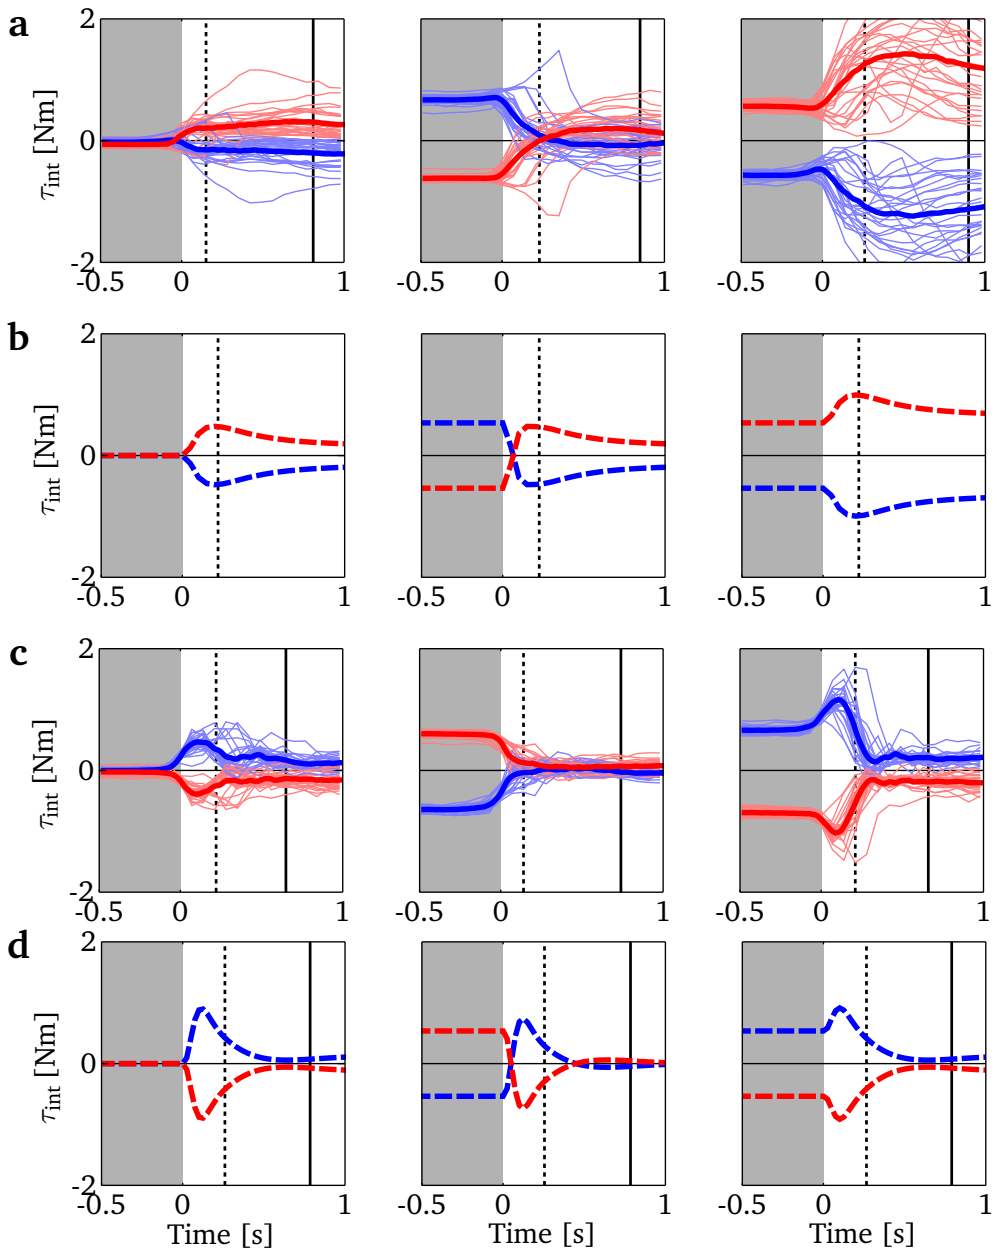


Figure 3.4 | Interpreting dyadic reaching in 2 dyads using the simple model. **a & b**, Data and simulations from a dyad where partner B (red) consistently aimed for a goal that overshot the target. The goal was modified to $\theta_g^A = 25\,\text{deg}$, $\theta_g^B = 35\,\text{deg}$. **c & d**, Data and simulations from a dyad where partner A (blue) had a stronger desire to reach their goal ($\mathbf{Q}^{\text{A}}_{1,1} = 30$). Using the same cost functions, the second push-pull block matches well with the data. However, when partner A is pulling away from the target (middle column), their torque does not become positive halfway through the movement as the simulations predict. This implies that the partner A's desire to reach the goal must have reduced whilst pulling away from the target.

3.3 Discussion

Rigidly coupled partners exhibited many behaviours during reaching movements. The interaction forces observed during these movements are difficult to describe and analyse qualitatively, which may explain why previous studies avoided or showed few of the interaction force profiles (Basdogan *et al.* 2000; Reed & Peshkin 2008; Groten *et al.* 2009; Stefanov *et al.* 2009; van der Wel *et al.* 2011; Madan *et al.* 2015).

To explain the behaviours, I formulated a simulation framework of two humans independently planning and executing a movement, under the influence of the interaction forces. These simulations revealed that partners with different motion plans explained the differences in behaviour. The motion plan, composed of the state cost $\mathbf{Q}$ and goal θ_g, did not evolve as a function of the trials but was determined by the initial conditions of the movement; for instance, the presence of interaction forces at the beginning of the movement biased the goal such that pushing towards the target caused subjects to overshoot.

The framework's ability to predict the interaction torque profiles, and the consistency of movements within each block, are indicative of the mechanism of *co-activity*, where each partner plans their movement independently and are influenced by the interaction forces, but mostly ignore them (Jarrassé *et al.* 2012). Clearly, dyadic reaching is not a suitable task to examine interaction behaviour as no evolution in the behaviour was observed after 30 movements were done together. Furthermore, the rigid connection makes it difficult to distinguish the actions of each partner. What is needed is a task where each partner's actions can be understood, and the behaviour evolves over time. The next Chapter 4 explores such a task.

4
Sensorimotor interaction mechanism

The dyadic reaching task is not a suitable task to investigate interaction between partners. Burdet's group designed a novel experimental paradigm where partners continuously interacted through a *compliant* virtual spring whilst tracking the same target together (Ganesh *et al.* 2014). The spring was compliant such that subjects could feel the interaction forces, and could easily overcome and ignore them if they were a hindrance; they could not relax entirely as the force from the spring was not strong enough to drive subjects to the target. This dyadic tracking experiment revealed the surprising result that interpersonal interaction was mutually beneficial to tracking performance; the better partner was able to improve even with a worse partner!

I propose four mechanisms to explain the mutual improvement in performance. *Co-activity* suggests that the improvement is a by-product of the connection dynamics alone. *Follow the better* proposes that subjects follow the best performing partner at every time instance. *Multisensory integration* suggests the mutual improvement comes from subjects taking a statistically optimal combination of the goal and partner trajectories. Finally, *goal integration* proposes that partners directly shared their goals. Simulations of each mechanism reveal that mutual improvement is predicted by the goal integration model only, but how could partners share their goals directly without access to the partner's vision?

I propose that subjects build an *internal model of the partner* to extract the partner's goal. The partner's goal, which is extrapolated from this internal model, is then combined optimally with their own visual estimate of the goal. This *interpersonal goal integration model* was converted into a computational algorithm to use as a robotic partner in the first human-robot experiment of its kind. Humans interacting with this robot exhibited similar improvements as observed in interpersonal sensorimotor interaction. The physiological basis of this model may lie in the mirror neuron system to understand the goal of the partner's actions.

4.1 Dyadic tracking

This section summarises the results of the experiment in (Ganesh *et al.* 2014); note that the experiments and data are not my own, but the analysis and interpretations are.

The twin visuomotor and haptic interface system (TVINS) shown in Figure 4.1a consists of two 2 degree-of-freedom (DOF) manipulandums located opposite each other. The arm is hidden from view by a table and the position of the hand is displayed as a cursor on a computer monitor placed in front. Each subject can only see their own cursor and target. Subjects were asked to track a smooth but randomly moving target using their cursor. Each trial lasted 60 seconds and the full experiment was separated into four sessions with 10, 20, 20 and 10 trials each respectively. In half of these trials, hereby referred to as *connected trials*, the two subjects interacted physically through a virtual spring with connection stiffness $120\,\mathrm{Nm}^{-1}$ and damping $7\,\mathrm{Nsm}^{-1}$. In *solo trials*, subjects were disconnected and tracked the target alone.

The tracking error, or the mean distance between the target and a subject's cursor, was used as a metric to determine performance in the task. To assess performance improvement due to interaction, the difference between the error in a connected trial and the error in the preceding solo trial was examined; if physical interaction is helpful, the error should be smaller in the connected trial. This performance improvement was plotted as a function of the difference in skill between the partners, which was measured through the relative difference in error in the solo trial such that positive difference in skill implied connection to a better partner, and negative to a worse partner. By having both connected and solo trials, the effect of physical interaction on the tracking performance could be observed, depending on how better or worse the partner was.

Figure 4.1b shows the improvement in performance as a function of the difference in skill. When connected to a better partner, subjects improved monotonically. What is surprising is that the better partner also improved from physical interaction with a worse partner, making continuous interpersonal interaction mutually beneficial. What mechanism could explain this surprising result?

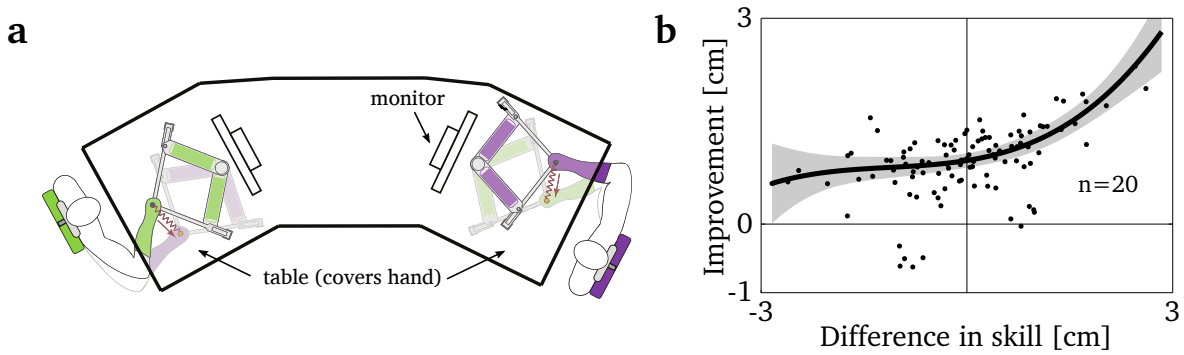


Figure 4.1 | Compliant interaction during dyadic tracking. **a**, Experimental setup for physical interaction in a dyad. For each subject, a 2 DOF manipulandum and a monitor were provided. Subjects had to track a randomly moving target on the monitor using their cursor/arm. In half the trials, the two subjects interacted through a compliant virtual spring, and in the other half they tracked the target alone. **b**, Performance improvement from sensorimotor interaction as a function of the difference in skill. The worse partner improves monotonically with a better partner. Surprisingly, even the better partner improves thanks to a worse partner. Continuous interpersonal sensorimotor interaction is mutually beneficial for both partners.

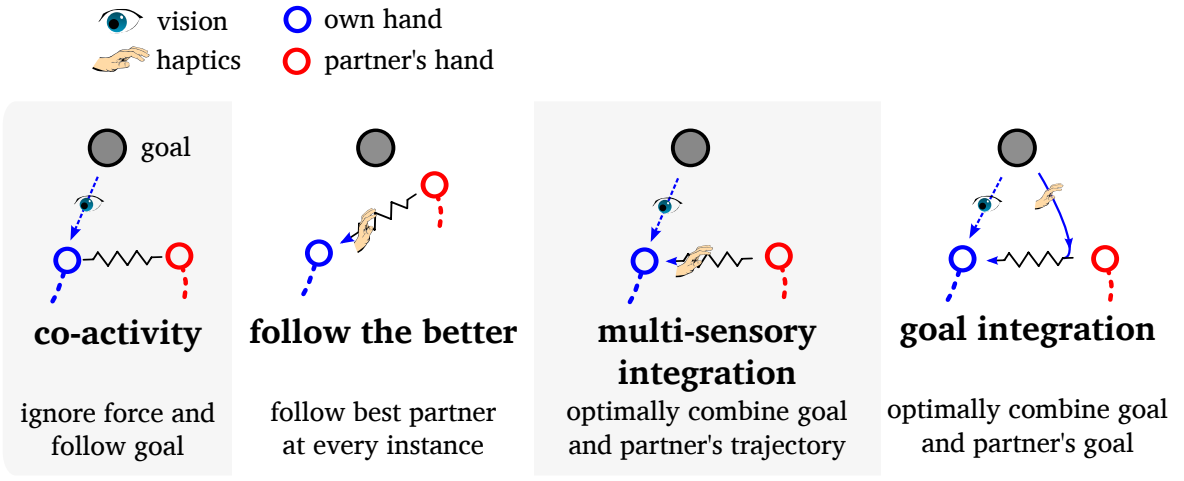


Figure 4.2 | Schematic of proposed mechanisms to explain the results of sensorimotor interaction. Four different mechanisms are proposed to explain the mutual improvement in performance during a dyadic tracking task. These are: *co-activity, follow the better, multi-sensory integration* and *interpersonal goal integration*. In *co-activity*, subjects are affected by the dynamics of the spring connecting them to the partner, but use only their visual information to track their goal and ignore the forces. *Follow the better* proposes that subjects measure the partner's tracking error and switch to following the partner, instead of the goal, if they are better. *Multi-sensory integration* suggests that subjects follow the statistically optimal combination of the goal and the partners' trajectories. Finally, *goal integration* proposes that subjects followed the optimal combination of own goal and the partner's goal.

4.2 Proposed mechanisms of interaction

Several candidates could possibly explain the mutually beneficial effects of interpersonal sensorimotor interaction. Figure 4.2 shows a schematic of the proposed mechanisms of sensorimotor interaction, which are outlined below.

Co-activity

The interaction forces from the spring connecting the partners alone may have caused the improvement; the subjects simply tracked their goal using vision, and ignored the forces (Jarrassé *et al.* 2012); the performance improvement is a by-product of the mechanical connection between the partners.

Follow the better

The subjects could actively monitor the haptic forces, and switch to following the partner in instances when the partner's performance was better than their own; this *follow the better* strategy is motivated by previous work in collective decision-making, where the decision of the best performing partner is taken (Laughlin & Ellis 1986; Hastie & Kameda 2005).

Multi-sensory integration

The subject may extract the position of the partner from the interaction force and follow a statistically optimal combination of the goal and the partners' trajectories; such *multi-sensory integration* of visual and haptic information has been observed in humans (Ernst & Banks 2002; Körding & Wolpert 2004) and could explain performance improvement due to the reduction in estimation error of the target's position.

Goal integration

Subjects may have shared their goals directly, which was combined optimally with the visual estimate of the goal. This mechanism is distinctly different from *multi-sensory integration* as it requires one to know the partner's goal, and not their trajectory.

The next Section 4.3 formulates a simulation framework for dyadic tracking to test each mechanism and compare their predictions to the data.

4.3 Dyadic tracking model

Let me describe the control strategy from the point of view of an agent who would be aware of the interaction mechanism: when tracking a randomly moving target, I need to estimate the difference between my hand position h and the position of the goal g I want to follow. The movement of both my hand and a smoothly moving goal can be modelled by a second order system buffeted by unbiased Gaussian white noise in the acceleration. The state vector is

$$\mathbf{x} \equiv \begin{bmatrix} h - g \\ \dot{h} - \dot{g} \\ \ddot{h} - \ddot{g} \end{bmatrix} \tag{4.1}$$

which evolves in discrete time according to

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}u_k + \boldsymbol{\mu}\,, \quad \mathbf{A} \equiv \begin{bmatrix} 1 & dt & dt^2/(2m) \\ 0 & 1 & dt/m \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{B} \equiv \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \tag{4.2}$$

where k is the time index and dt the sampling time, and m is the (point) mass of the hand. $\mathbf{x}$ represents my planned movement towards the target. $\boldsymbol{\mu} \in \mathcal{G}(\mathbf{0}, \mathbf{M})$ is a vector of unbiased Gaussian white noise processes with covariance matrix

$$\mathbf{M} \equiv E\left[\boldsymbol{\mu}\,\boldsymbol{\mu}^T\right] = \sigma_a^2 \begin{bmatrix} dt^3/6 \\ dt^2/2 \\ dt \end{bmatrix} \begin{bmatrix} dt^3/6 & dt^2/2 & dt \end{bmatrix} = \sigma_a^2 \begin{bmatrix} dt^6/36 & dt^5/12 & dt^4/6 \\ dt^5/12 & dt^4/4 & dt^3/2 \\ dt^4/6 & dt^3/2 & dt^2 \end{bmatrix} \tag{4.3}$$

where σ_a^2 is the standard deviation of the white noise in acceleration, enabling me to identify the unmodelled moving goal.

The position difference measured from vision and proprioception is modelled as

$$z_{k'} = \mathbf{C}\mathbf{x}_{k'} + \nu\,, \quad \nu \in \mathcal{G}(0, N) \tag{4.4}$$

with

$$z_{k'} = h_{k'} - g_{k'}\,, \quad \mathbf{C} = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}, \quad N = \sigma^2, \tag{4.5}$$

where $k' \equiv k - \delta$ is the time index with sensory delay δ, and σ^2 is the sensory noise, which incorporates both visual and proprioceptive noise.

The recursive least square estimation of the state $\mathbf{x}$ (minimising $E[||z_{k'} - \hat{z}_{k'}||^2]$) yields (Simon 2006):

$$\begin{aligned}
\hat{\mathbf{x}}_{k'+1} &= \mathbf{A}\hat{\mathbf{x}}_{k'} + \mathbf{B}u_{k'} + \mathbf{K}_{k'}\left[z_{k'} - \mathbf{C}(\mathbf{A}\hat{\mathbf{x}}_{k'} + \mathbf{B}u_{k'})\right] \\
\mathbf{Y}_{k'} &= \mathbf{A}\mathbf{P}_{k'}\mathbf{A}^T + \mathbf{M} \\
\mathbf{K}_{k'} &= \mathbf{Y}_{k'}\mathbf{C}^T(\mathbf{C}\mathbf{Y}_{k'}\mathbf{C}^T + N_{k'})^{-1} \\
\mathbf{P}_{k'+1} &= (\mathbf{I} - \mathbf{K}_{k'}\mathbf{C})\mathbf{Y}_{k'}\,.
\end{aligned} \tag{4.6}$$

For simplicity, I assume that the sensory noise σ^2 is known to the central nervous system through experience. However, the covariance matching method (Mehra 1972) could be used to compute the value of the sensory noise online. In this method, the difference between the theoretical covariance $\mathbf{C}\mathbf{P}\mathbf{C}^T + N$ and the actual covariance is minimised. This yields the adaptation law based on the innovation, $z - \mathbf{C}(\mathbf{A}\hat{\mathbf{x}} + \mathbf{B}u)$,

$$\widehat{N}_{k'} = -\mathbf{C}\mathbf{Y}_{k'}\mathbf{C}^T + \frac{1}{k'}\sum_{i=0}^{k'}\left[z_i - \mathbf{C}(\mathbf{A}\mathbf{x}_i + \mathbf{B}u_i)\right]\left[z_i - \mathbf{C}(\mathbf{A}\mathbf{x}_i + \mathbf{B}u_i)\right]^T. \tag{4.7}$$

I can move my hand towards the target to minimise this estimated difference by sending the motor command (change in acceleration)

$$u_k = -\mathbf{L}\widehat{\mathbf{x}}_k \tag{4.8}$$

with the controller gain $\mathbf{L}$ selected by minimising the cost

$$E\left[\sum_{k=0}^{\infty}\mathbf{x}_k^T\mathbf{Q}\,\mathbf{x}_k + u_k^T R u_k\right] \tag{4.9}$$

where the matrix $\mathbf{Q}$ is positive semi-definite and R is positive. This infinite horizon linear quadratic regulator (LQR) is found by solving the Riccati equation (Stengel 2012)

$$\mathbf{L} = (R + \mathbf{B}^T\mathbf{S}\mathbf{B})^{-1}\mathbf{B}^T\mathbf{S}\mathbf{A}\,, \quad \mathbf{S} = \mathbf{Q} + \mathbf{A}^T\left[\mathbf{S} - \mathbf{S}\mathbf{B}(R + \mathbf{B}^T\mathbf{S}\mathbf{B})^{-1}\mathbf{B}^T\mathbf{S}\right]\mathbf{A}\,. \tag{4.10}$$

Note that eq. (4.8) requires the state estimate at time step k. To send the appropriate motor command, I must compensate for the sensory delay δ by iterating for δ steps using the forward model

$$\begin{aligned}
\widehat{\mathbf{x}}_k &= \mathbf{A}^{\delta-1}\left(\mathbf{A}\widehat{\mathbf{x}}_{k'} + \mathbf{B}u_{k'}\right) \\
&= \mathbf{A}^{\delta-1}\left(\mathbf{A} - \mathbf{B}\mathbf{L}\right)\widehat{\mathbf{x}}_{k'}\,.
\end{aligned} \tag{4.11}$$

During interaction trials, I experience the elastic force

$$F_k \equiv -K\left(h_k - \widetilde{h}_k\right) + D\left(\dot{h}_k - \dot{\widetilde{h}}_k\right) \tag{4.12}$$

where K and D are the connection stiffness and damping respectively. Hence, the state equation of the tracking process then yields

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\left(u_k + F_k\right) + \boldsymbol{\mu}\,, \quad \boldsymbol{\mu} \in \mathcal{G}(\mathbf{0}, \mathbf{M})\,. \tag{4.13}$$

The forward model compensation of the delay in eq. (4.11) is modified to

$$\widehat{\mathbf{x}}_k = \mathbf{A}^{\delta-1}\left[\left(\mathbf{A} - \mathbf{B}\mathbf{L}\right)\widehat{\mathbf{x}}_{k'} + \mathbf{B}F_{k'}\right] \tag{4.14}$$

where the last known elastic force $F_{k'}$ is used.

Simulations of proposed mechanisms

I simulated the proposed mechanisms in Section 4.2 using this simulation framework. Recall that in *co-activity*, subjects neglect the haptic force and track their visual goal independently of the force they feel. *Co-activity* is simulated using partners who have only visual information of the goal, as in the the solo trials, but are affected by the spring force F (i.e. the dynamics of the connection). The observed behaviours would then be the by-product of the mechanical connection between two individuals performing their own task independently. Simulations reveal that this mechanism predicts linear improvement in the worse partner and a linear decrease in performance for the better partner (see Figure 4.4).

Next, *follow the better* proposes that the worse partner simply follows the better partner. Here, I assume that the ability to haptically track a goal trajectory is equivalent to tracking a visual goal with the hand (i.e. that visual and haptic sensory noise are equivalent). The elastic force F corresponds to the connection with the partner, from which subjects can estimate their partner's hand position

$$\widetilde{h}_k \approx F_k/K + h_k \tag{4.15}$$

where the stiffness K is identified by the subject through the history of interaction. These haptic forces are monitored to calculate the partner's tracking error and my error

$$e_k = h_k - g_k\,, \quad \widetilde{e}_k = \widetilde{h}_k - g_k \tag{4.16}$$

at every time instance. I will switch to following my partner if their error is lower. The sensory

observation is modified to

$$z_{k'} = \begin{cases} h_{k'} - g_{k'} & \text{if } |e_{k'}| \leq |\widetilde{e}_{k'}| \\ h_{k'} - \widetilde{h}_{k'} & \text{otherwise.} \end{cases} \tag{4.17}$$

This *follow the better* mechanism is also insufficient as partners never improve from interaction (see Figure 4.4). This is not surprising as the simulated subjects are trying to haptically follow a partner with the same ability as visually tracking a goal; the partner's trajectory is already deviated from the goal, so partners would be better off following their goal at all times. I also simulated another version of the follow the better, "follow the better by relaxing"; in this version, the partner who is worse at that time instance relaxes by setting their motor command $u_k = 0$. This *follow the better by relaxing* mechanism resulted in the same improvement behaviour as *co-activity*. The assumption that visual and haptic tracking have similar noise may not be reasonable, but even with perfect haptic tracking the better partner will not improve. The different forms of the follow the better mechanism are not sufficient to explain the results of human-human interaction.

In *multi-sensory integration*, the haptic forces are assumed to provide another measurement of the goal; subjects follow the statistically optimal combination of this haptic goal, or the partner's position, and their visual goal. This is described by modifying the measurement eq. (4.5) with

$$\mathbf{z}_{k'} \equiv \begin{bmatrix} h_{k'} - g_{k'} \\ h_{k'} - \widetilde{h}_{k'} \end{bmatrix}, \mathbf{C} \equiv \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \mathbf{N} \equiv \begin{bmatrix} \sigma^2 & 0 \\ 0 & \widetilde{\sigma}^2 \end{bmatrix}, \tag{4.18}$$

where σ^2 and $\widetilde{\sigma}^2$ are the sensory noises from each respective partner. As the simulation results in Figure 4.4 shows, multi-sensory integration does not predict mutually beneficial effects on performance. Interestingly, *follow the better* and *multi-sensory integration* are two sides of the same coin; the only difference between these two mechanisms is the weighting between the goal and the partner's trajectory. As the simulations show, no weighting between the visual estimate of the goal trajectory and the partner's trajectory will yield a more accurate estimate of the goal consistently.

In *goal integration*, subjects could extract another sensory measurement of the goal from the partner and combine it with the visual estimate. This can be simulated by changing eq. (4.5) to

$$\mathbf{z}_{k'} \equiv \begin{bmatrix} h_{k'} - g_{k'} \\ h_{k'} - \widetilde{g}_{k'} \end{bmatrix}, \mathbf{C} \equiv \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \mathbf{N} \equiv \begin{bmatrix} \sigma^2 & 0 \\ 0 & \widetilde{\sigma}^2 \end{bmatrix}, \tag{4.19}$$

which will yield, through eq. (4.6), the stochastic optimal integration of the partner's goal $\widetilde{g}$ with own visual estimate of the goal g. Figure 4.4 shows the simulation from the goal integration model, which predicts mutual performance improvement in both the better and

worse partners, as in the data.

Sensitivity of simulation results to the parameters

The parameters used in the simulations were $R = 10^{-5}$, $\lambda = 10^{-5}$ and $m = 1\,\text{kg}$. The sensory noise determines the tracking error of a partner; to compare with the experimental data directly, partners were simulated with sensory noise corresponding to the same tracking errors as in the data.

I use the sum of squared error

$$\text{SSE} = \sum_{i=0}^{n} (\mathbf{y}_i - \hat{\mathbf{y}}_i)^2 \tag{4.20}$$

as the metric for predictive power, where n is the number of data points, $\mathbf{y}$ is a vector of data points and $\hat{\mathbf{y}}$ is a vector of predictions of a mechanism.

There are two free parameters in the model, which are the magnitude of the Gaussian noise in the target's jerk σ_a^2 and the state cost multiplier q. Low values of $\sigma_a^2 < 1$ result in unstable tracking as the Kalman filter cannot capture the motion of the target sufficiently. To determine a reasonable upper bound for σ_a^2, the spectrum of the trajectory in the data can be compared with the simulation output. As Figure 4.3a shows, larger values of σ_a^2 increase the spectral content of the trajectory at $\approx 1\,\text{Hz}$, resulting in less and less human-like trajectories. Interestingly, the predictive power of each mechanism increases with larger σ_a^2, i.e. higher frequency movements result in greater performance improvement for all mechanisms, with the exception of the *interpersonal goal integration* model which is insensitive to this change (see Figure 4.3b).

The other free parameter q affects the controller's strength, and affects the state cost matrix through

$$\mathbf{Q} = q\,\text{diag}(1, 0.1, 0) \tag{4.21}$$

where the ratio between the position, velocity and acceleration costs was fixed because varying their balance did not significantly affect the simulation predictions, e.g. $\mathbf{Q} \equiv \text{diag}(1, 0.1, 0)$ and $\mathbf{Q} \equiv \text{diag}(1, 10, 0)$ yielded similar results for all proposed mechanisms.

q was varied in powers of 10 between 10^{-1}-10^5. Too low q results in motion where the connection force F overpowers the controller's strength, resulting in instability; too high q corresponds to superhuman power in maintaining a desired trajectory, corresponding to arm impedance which is beyond the realm of human strength. Figure 4.3b reveals that the *follow the better* and *multi-sensory integration* mechanisms were best with large q, and the *co-activity* mechanism has a minimum when $q = 1$. The *interpersonal goal integration* model was, once again, insensitive to changes in q. Not only is this model robust to changes in σ_a^2 and q, but it

has the highest predictive power (lowest SSE) over the whole parameter space.

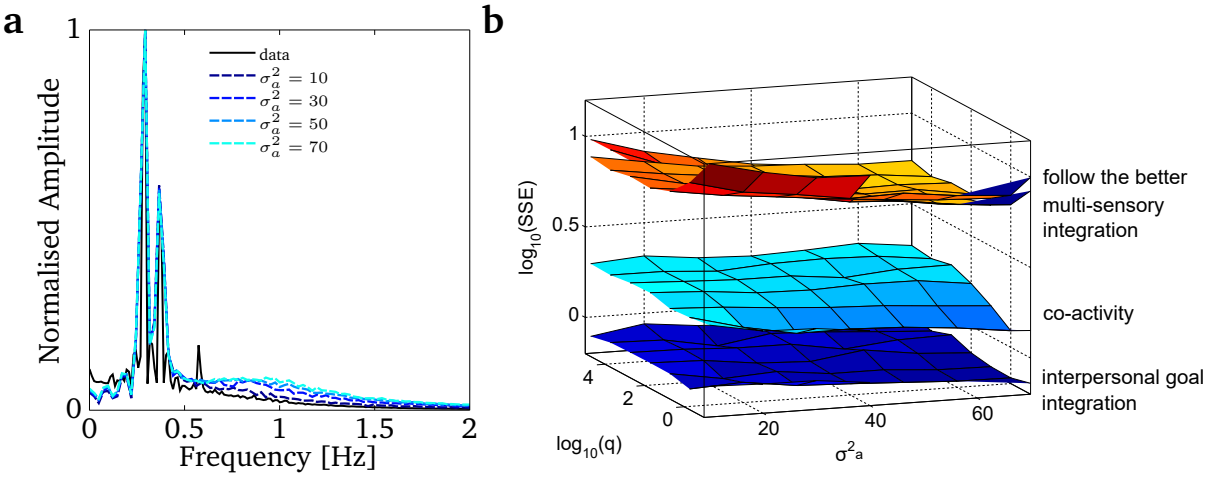


Figure 4.3 (a) Deciding the range for σ_a^2 through the trajectory's spectrum. (b) The predictive power, measured by SSE, of the four proposed mechanisms of interaction as a function of σ_a^2 and q.

However, how could subjects identify the partner's goal through the interaction forces alone? What is the underlying mechanism behind the *goal integration* model?

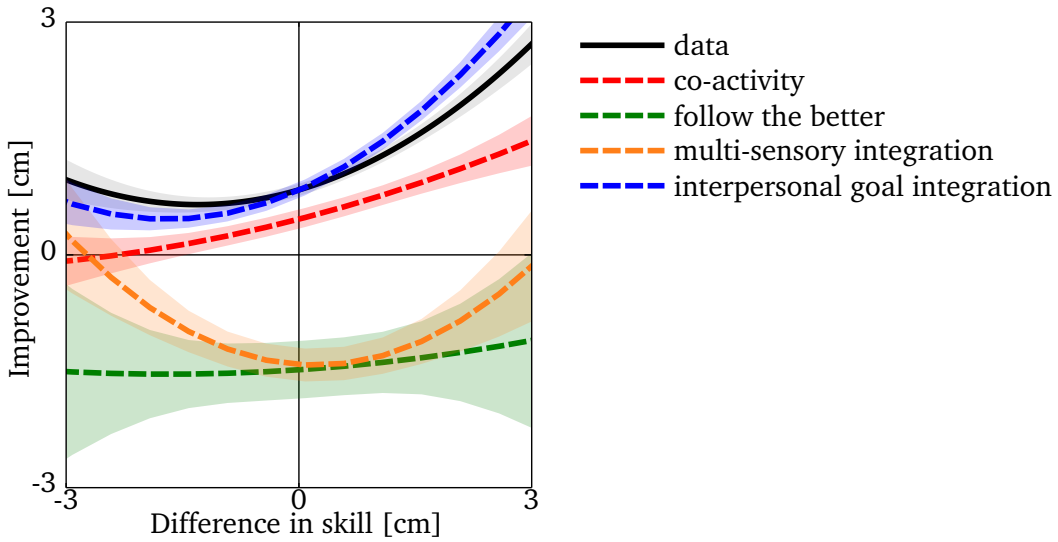


Figure 4.4 | Best predicted performance improvement from simulations of different mechanisms. The parameters from the sensitivity analysis that yielded the best prediction were selected for each mechanism, and their predictions are shown here. The *follow the better* and *multi-sensory integration* mechanisms both predict negative improvement which does not match the data. *Co-activity* predicts linear improvement for the worse and linear deterioration in the better partner's performance. On the other hand, the *interpersonal goal integration* model predicts performance improvement in both the better and worse partners, corresponding to the data.

4.4 Interpersonal goal integration

I showed using simulations that only the goal integration model yields the mutual improvement during sensorimotor interaction. However, the mechanism to extract the partner's goal is not clear. Here, I propose that subjects build an internal model of the partner from the visual estimate of the goal, the proprioceptive estimate of the arm and the haptic measurement of the external (interaction) force.

This hypothesis, of building an internal model of the partner, is an extension of well-established motor learning principles, where subjects learned to counteract external forces from a state-dependent force field by learning the abstract relationship between the state of the arm and the forces (Shadmehr & Mussa-Ivaldi 1994; Conditt *et al.* 1997). Humans can also adapt to loads with different dynamic properties (Bock 1990; Dingwell *et al.* 2002; Dingwell *et al.* 2004). The *internal model* is a forward model of the system dynamics, i.e. what is expected to happen as a consequence of a motor force. In continuous interpersonal sensorimotor interaction, I propose that subjects form an internal model of the partner's actions using vision, proprioception and haptics.

4.4.1 Building an internal model of the partner

At this point, one might ask why an internal model of the partner is necessary. Why not build a look-up table between the interaction forces and the goal? Imagine being connected to a partner who is tracking the same goal as you. One can create some relationship between the forces and the goal, but this look-up table would instantly become worthless if your own or the partner's movement changed. This form of rote memorisation between the interaction forces and the goal is not likely to be used, as humans have been shown to form internal models of force fields based on the kinematic state of the arm (Conditt *et al.* 1997).

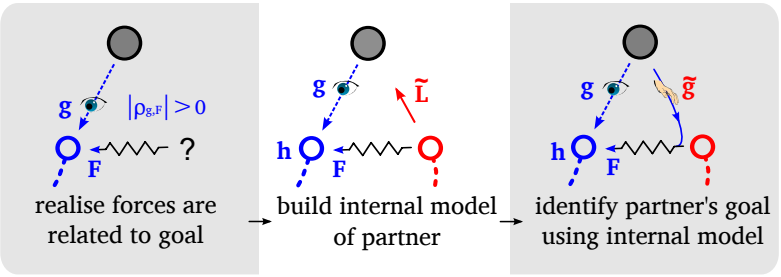


Figure 4.5 | Identifying the partner's goal through an internal model of the partner. Subjects first realise that the interaction forces and the goal are significantly correlated. Subjects then form an internal model of the partner assuming that they are following the same goal. The visual estimate of the goal, proprioceptive estimate of own hand and the haptic estimate of the interaction force are used to estimate the partner's state and their control policy.

Vision, proprioception and haptics will be crucial to forming an internal model of the partner. Figure 4.5 shows a schematic of how such an internal model can be identified; first, subjects feel external forces that are significantly correlated with the visual estimate of the goal. The central nervous system (CNS) will create an internal model to predict these forces, and in doing so it must create a model of the partner who follows the same goal to fully model the interaction forces.

The CNS considers a system consisting of two partners following a target, connected by a spring of stiffness K and damping D. The state of this system, in one dimension, is defined as

$$\boldsymbol{\xi} = \left[\mathbf{h},\, \widetilde{\mathbf{h}},\, \widetilde{\mathbf{g}},\, F\right]^T \tag{4.22}$$

where $\mathbf{h}$ is the state of own arm (position, velocity and acceleration), $\widetilde{\mathbf{h}}$ is the state of the partner's arm, $\mathbf{g}$ is the goal state and F is the interaction force. Considering the similarity between the partner and my trajectories (correlation of $\rho = 0.6$ from the data), it would make sense to use pre-existing neural networks of my tracking task to infer the partner's actions. My CNS models the partners' motion as if they are tracking my goal, and I identify the inverse model of their tracking, i.e. the partner's control gain $\widetilde{\mathbf{L}}$, through their motor command

$$\widetilde{u}_k = -\widetilde{\mathbf{L}}_k \left(\widetilde{\mathbf{h}}_k - \widetilde{\mathbf{g}}_k\right) \tag{4.23}$$

where $\widetilde{\mathbf{L}} = \left[l_p,\, l_v,\, l_a\right]$ is a row-vector. The full state is thus extended to include the partner's control gain

$$\boldsymbol{\xi} = \left[\mathbf{h},\, \widetilde{\mathbf{h}},\, \mathbf{g},\, F,\, \widetilde{\mathbf{L}}\right]^T \tag{4.24}$$

and is described by the nonlinear state equation

$$\boldsymbol{\xi}_{k+1} = \mathbf{f}(\boldsymbol{\xi}_k) \tag{4.25}$$

which evolves according to

$$\begin{bmatrix} \mathbf{h}_{k+1} \\ \widetilde{\mathbf{h}}_{k+1} \\ \widetilde{\mathbf{g}}_{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A} \end{bmatrix} \begin{bmatrix} \mathbf{h}_k \\ \widetilde{\mathbf{h}}_k \\ \widetilde{\mathbf{g}}_k \end{bmatrix} + \begin{bmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{0} & \mathbf{B} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} u_k \\ \widetilde{u}_k \end{bmatrix}$$

$$\begin{aligned} \widetilde{u}_k &= -\widetilde{\mathbf{L}}_k \left(\widetilde{\mathbf{h}}_k - \widetilde{\mathbf{g}}_k\right) \\ F_k &= \begin{bmatrix} K & D & 0 \end{bmatrix} \left(\widetilde{\mathbf{h}}_k - \mathbf{h}_k\right) \\ \widetilde{\mathbf{L}}_{k+1} &= \widetilde{\mathbf{L}}_k \end{aligned} \tag{4.26}$$

where

$$\mathbf{A} \equiv \begin{bmatrix} 1 & dt & dt^2/(2m) \\ 0 & 1 & dt/m \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{B} \equiv \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}. \tag{4.27}$$

The partner's inverse model is identified by allowing the partner's control gain to evolve with the system noise. The state and the partner's inverse model must be estimated simultaneously, and is solved by linearising the system of eq. (4.25) and recursively minimising $E[||z_k - \hat{z}_k||^2]$ through (Ljung 1979)

$$\begin{aligned} \hat{\boldsymbol{\xi}}_{k+1} &= f(\hat{\boldsymbol{\xi}}_k) + \mathbf{K}_k \left[z_k - \mathbf{C}\mathbf{f}(\hat{\boldsymbol{\xi}}_k) \right] \\ \mathbf{Y}_k &= \mathbf{F}_k \mathbf{P}_k \mathbf{F}_k^T + \overline{\mathbf{M}} \\ \mathbf{K}_k &= \mathbf{Y}_k \mathbf{C}^T (\mathbf{C}\mathbf{Y}_k\mathbf{C}^T + \overline{\mathbf{N}})^{-1} \\ \mathbf{P}_{k+1} &= (\mathbf{I} - \mathbf{K}_k\mathbf{C})\mathbf{Y}_k \,. \end{aligned} \tag{4.28}$$

where $\mathbf{F}_k = (\frac{\partial f_i}{\partial \xi_j})\big|_{\hat{\mathbf{x}}_k}$ is the first-order linearisation of the nonlinear function f. This extended Kalman filter (EKF) is given the proprioceptive estimate of hand $\mathbf{h}$, visual estimate of goal $\mathbf{g}$ and haptic estimate of interaction force F through the observation

$$\begin{aligned} \mathbf{z}_k &\equiv \left[h_k,\, \dot{h}_k,\, \ddot{h}_k,\, g_k,\, \dot{g}_k,\, \ddot{g}_k,\, F_k \right]^T \\ \overline{\mathbf{N}} &\equiv \text{diag} \left[\sigma_h^2,\, \sigma_{\dot{h}}^2,\, \sigma_{\ddot{h}}^2,\, \sigma_g^2,\, \sigma_{\dot{g}}^2,\, \sigma_{\ddot{g}}^2,\, \sigma_F^2 \right] \end{aligned} \tag{4.29}$$

where σ_h^2, $\sigma_{\dot{h}}^2$ and $\sigma_{\ddot{h}}^2$ are the measurement noises for own hand position, velocity and acceleration; σ_g^2, $\sigma_{\dot{g}}^2$ and $\sigma_{\ddot{g}}^2$ are the noise for goal position, velocity and acceleration and σ_F^2 for the interaction force respectively. I assume that all sensory measurements have equal noise. The observation matrix $\mathbf{C}$ is linear with feedback given to the appropriate state in the EKF.

Does this internal model of the partner truly yield another measurement of the goal? To demonstrate this, I can make the internal model blind to the goal after the partner's control gains have been reasonably identified; in this scenario, the internal model must estimate the goal based on proprioception, haptics and the estimated partner's control gains alone. The sensory observations during such "blind tracking" are

$$\begin{aligned} \mathbf{z}_k &\equiv \left[h_k,\, \dot{h}_k,\, \ddot{h}_k,\, F_k,\, \hat{l}_p,\, \hat{l}_v,\, \hat{l}_a \right]^T \\ \mathbf{N} &\equiv \text{diag} \left[\sigma_h^2,\, \sigma_{\dot{h}}^2,\, \sigma_{\ddot{h}}^2,\, \sigma_F^2,\, \sigma_{l_p}^2,\, \sigma_{l_v}^2,\, \sigma_{l_a}^2 \right]. \end{aligned} \tag{4.30}$$

Figure 4.6a shows a simulation from two partners where I remove the sensory observation of the goal in one partner from 30 seconds onwards; this blind partner can only estimate the goal using proprioception and haptic feedback. Once the visual observation of the goal is removed,

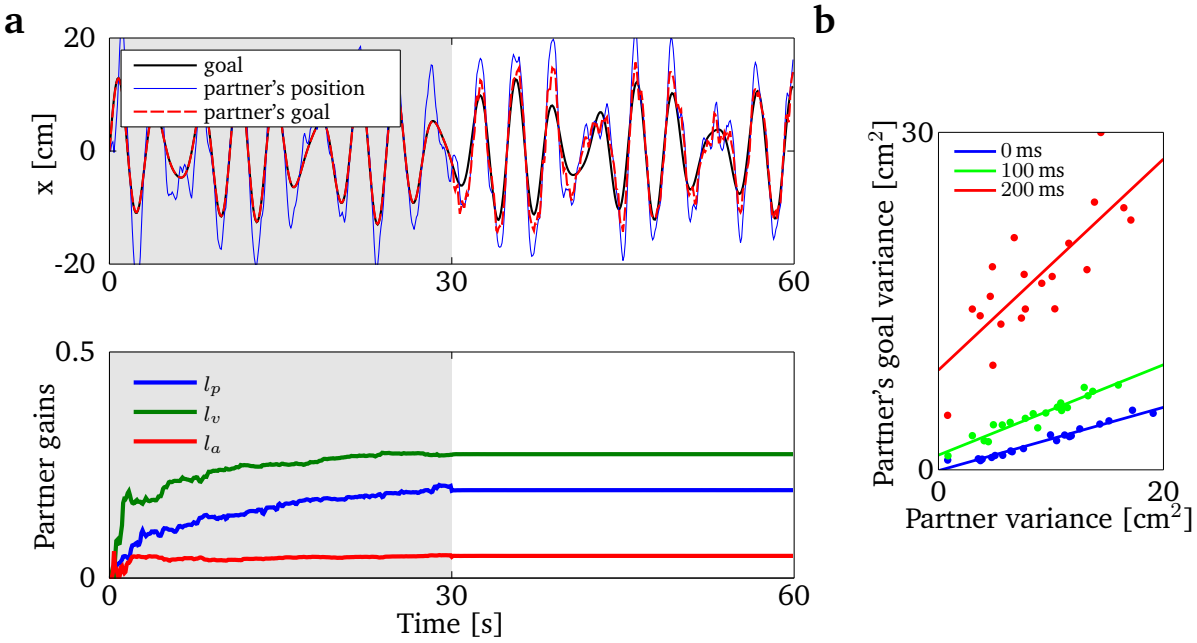


Figure 4.6 | Estimate of the partner's goal from the internal model of a partner. a, After 30s, the internal model of the partner becomes blind to the visual estimate of the goal, and must rely on proprioception, haptics and the partner's gains to estimate the goal. Although the estimate of the goal deteriorates without vision, the internal model has a reasonable estimate of the goal that does not drift many seconds after visual feedback of the goal is removed, but this is only if the partner's behaviour remains consistent over time. **b**, The noise in the partner's goal from the internal model as a function of the variance in the partner's trajectory. The noise in the partner's goal is found to linearly increase as the variance in the partner's trajectory increases, which is logical as any least square approximation of the partner's goal is reliant on the consistency of the partner's behaviour. This relationship between the variances of the partner's goal and trajectory is modulated by the sensory delay in the internal model; as the delay increases, so does the noise in the partner's goal even for the same variance in the partner's trajectory.

the estimate of the goal $\mathbf{g}$ deteriorates in comparison to when vision was available, but the internal model of the partner gives a reasonably good estimate with no drift in the goal's position. Of course, the partner's behaviour must stay constant over time, which may not be a reasonable assumption. The blind estimation of the partner's goal is likely to be accurate only within several hundred milliseconds after visual feedback of the goal is removed.

The noise in the partner's goal is a linear function of the variance in the partner's trajectory, and the greater the sensory delay in the internal model, the steeper the gradient becomes (see Figure 4.6b).

The internal model of the partner above was run alongside the interaction framework described in Section 4.3. Previously defined parameters were kept the same, and the process

noise was set to

$$\overline{\mathbf{M}} = \begin{pmatrix} \mathbf{Q} & & \cdots & & 0 \\ & \mathbf{Q} & & & \\ \vdots & & \mathbf{Q} & & \vdots \\ & & & 1 & \\ 0 & & \cdots & & \gamma\mathbf{I} \end{pmatrix} \tag{4.31}$$

where γ is the learning rate of the partner's control gains, which was set to $\gamma = 10^{-5}$. The sensory delay in the internal model was set to $200\,\text{ms}$.

Figure 4.7 shows the simulated improvement from the full interpersonal goal integration model, internal model of the partner included. As the partner's goal is successfully extracted from the interaction, mutual improvement in performance is observed for both better and worse partners, as in the data. As the interpersonal goal integration model reproduced key features of the human-human interaction data, it was implemented as a robotic partner to interact with a human in real-time.

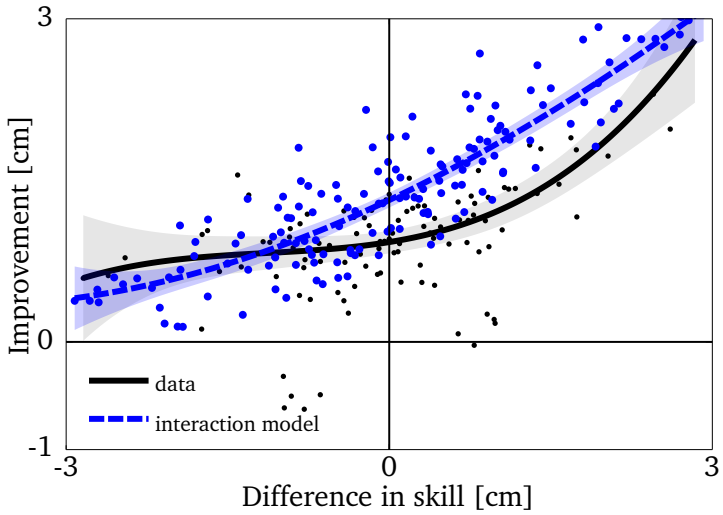


Figure 4.7 | Simulation of the interpersonal goal integration model. 150 simulated trials of the internal model formulation of the *goal integration* model. This *interpersonal goal integration* model predicts the mutual performance improvement of both the better and worse partners, and makes a similar prediction to the goal integration model.

4.4.2 Human-robot validation experiment

The interpersonal goal integration model in Section 4.3 was implemented as a robotic partner and tested in the first human-robot experiment of its kind. The full *interpersonal goal integration* model required two 23 dimensional EKFs running in the background, which was not feasible with the TVINS setup; for this reason, the simple goal integration model without an internal model of the partner was used as it yielded the same improvements as the full model. The robot's visual noise was set to $\sigma = 0.001,\ 0.02$ and 0.07 m which resulted in tracking errors of $1.0,\ 3.5$ and 6.1 cm respectively. Technical limitations also constrained the maximum sensory delay that could be implemented, hence the delay was set to $\delta = 150$ ms.

12 right-handed subjects were recruited in pairs and gave their informed consent to participate in the experiment. The experiment was conducted according to the principles in the Declaration of Helsinki and approved by the ethics committee at Advanced Telecommunication Research Institute. Each subject interacted with their own robot. The same target and trial length as in (Ganesh *et al.* 2014) were used. The experiment consisted of 36 trials where solo and connected trials were experienced one after another (i.e. solo, connected, solo, connected, etc). Subjects interacted with the good, medium then bad robot in the first, middle and last 12 trials respectively.

Figure 4.8a shows the results of the human-robot interaction, from the human partner only. The human partners were able to improve both with a robot partner worse than themselves and with a better robot partner, just like when interacting with a human partner. Humans interacting with a good robot partner appear to improve less than those interacting with a human, but this discrepancy can be explained by the absolute errors of the partners. For the same difference in skill, two partners can both have high or both low errors, but the plot cannot distinguish between these two. The maximum-likelihood estimate predicts greater improvement if both partners have higher errors, even for the same difference in skill, as there is more room to improve. The good robot partner in the human-robot interaction had a low error of 1.0 cm, which was never observed in humans. Hence, the worse humans in the human-robot interaction were significantly better than the worse humans in the human-human interaction, and improved less as a result. Figure 4.8b shows the improvement as a ratio of the solo error; interaction with the robot or human resulted in equivalent improvement.

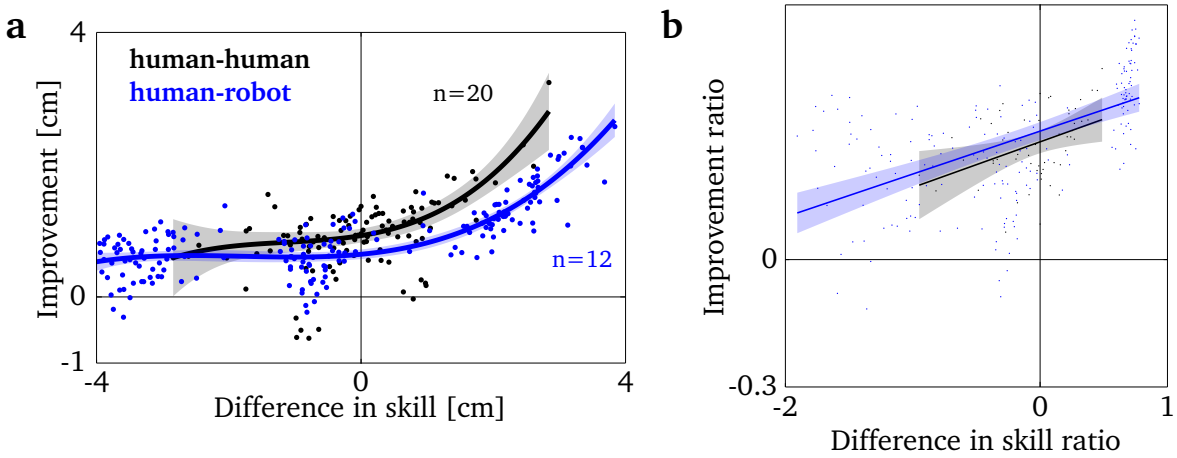


Figure 4.8 | Human-robot experiment to validate the interpersonal goal integration model. **a**, The interpersonal goal integration model was implemented as a robotic partner with three different levels of performance (good, medium and bad). The human-robot data is from the human only. The improvement curve was found to be comparable to the data in the original human-human interaction experiment. However, it seems like a worse human interacting with a better robot improved less than worse humans interacting with a better human. This difference can be explained by the absolute error of the partners; as the good robot had low error, the worse humans in human-robot interaction were much better than those in the human-human pairs. Thus, the human interacting with the robot had less room to improve. By plotting the improvement and difference in skill as a ratio of the solo error (**b**), the improvement between human-robot and human-human interaction are equivalent.

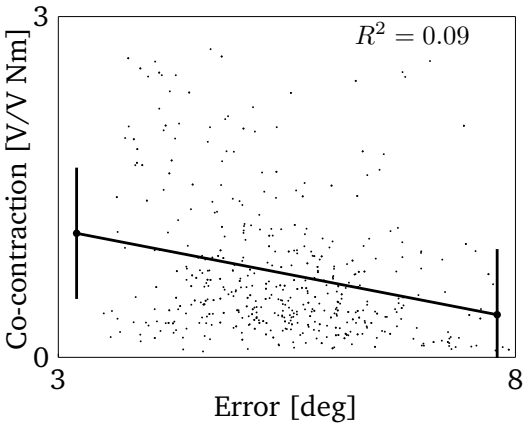


Figure 4.9 | Co-contraction as a function of error in solo trials only, from the experiment in Chapter 5. Torque-normalised co-contraction as a function of error, with error bars showing 1 standard deviation. The relationship between co-contraction and error is weak, with an $R^2 = 0.09$. It is unlikely that subjects infer the performance of the partner through their impedance.

4.5 Discussion

During a dyadic tracking task, partners improved their performance even with a partner worse than themselves. Several mechanisms were examined to explain the mutually beneficial effects of interpersonal sensorimotor interaction. *Co-activity*, where the dynamics of a spring connecting the partners together, could only predict improvement in the worse partner but not the better. *Follow the better* predicted deterioration for partners of all tracking skill. *Multi-sensory integration* similarly predicted deterioration in performance. *The goal integration model*, where partners share their goals, was the only model to reproduce the mutually beneficial effects during interaction. How could the partners share their goals directly? Subjects must form an internal model of the partner using vision, proprioception and haptics to infer the partner's actions, and thereby extract the partner's goal.

This *interpersonal goal integration* model was converted into a computational algorithm and implemented as a robotic partner in a human-robot interaction experiment; humans interacting with this robot improved similarly to interaction with another human partner, yielding a Turing test of sensorimotor control.

The dyadic tracking study in (Ganesh *et al.* 2014) is similar, yet distinct, from a previous study which examined a human-human perceptual decision-making task (Bahrami *et al.* 2010). Bahrami *et al*. proposed a mechanism where the partner's goal and their confidence of it were shared. The authors of the decision-making study mention that the distinction cannot be made between their model and the direct sharing of the partners' goals, as partners may have suboptimally combined their goals. Importantly, the dyadic tracking task *can* make this distinction; mutually improving partners can be explained only if the goals are shared between them. Goal sharing is only possible through an internal model of the partner, which is identified using own sensory information and reliability only. The extracted partner's goal is dependent on own sensory delay and the variability of the partner's trajectory, i.e. the partner's true performance at the task. Therefore, interpersonal physical interaction, unlike verbal communication, enables subjects to "cut through bullshit" to obtain the partner's true goal and true reliability.

One could argue that the mechanical impedance of the arm (or co-contraction) is a measure of the partner's confidence, which is what subjects used to estimate the reliability of the partner's goal. However, an experiment I conduct in the next Chapter 5, with electromyographic measurements of joint co-contraction, shows that there is no discernible relationship between absolute level of co-contraction and tracking error (Figure 4.9). Hence, arm impedance is not used to assess the variability in the partner's goal.

5
Connection strength on interaction

Unlike the dyadic reaching task where subjects did not adapt to their partner, the dyadic *tracking* task caused partners to adapt their behaviour, enabling me to identify the mechanism of interpersonal sensorimotor interaction. Besides changing the task, compliance was added to the physical connection between the partners in order to understand the adaptations in each partner separately. Now that the mechanism is identified, I can examine how changing the compliance of the connection may have affected sensorimotor interaction. Intuitively, a very compliant connection is close to no interaction, i.e. the benefits of sensorimotor interaction would be reduced. On the other hand, a rigid connection may enable partners to understand each other's actions with greater accuracy, but it will not allow the better partner to counter the worse partner's perturbing actions. Is there a trade-off between better understanding of the partner's goal and the ability to counter the partner's motion, or can the CNS find a win-win strategy to gain extra sensory information whilst absorbing the partner's perturbing actions?

In order to examine this question, I conducted a dyadic tracking experiment in 1 dimension with rigid, medium and compliant connections between the partners. The 1DOF system enabled me to study the muscle activity of each partner systematically. The performance, measured by the tracking error, of the worse partner was better with a more rigid connection. Surprisingly, the better partner was little affected by a rigid connection with a worse partner. Measurements of muscular activity revealed that the better partner's ability to maintain their performance relied on increased motor force and higher joint impedance to reject the partner's actions. Furthermore, the worse partner relaxed both their motor force and impedance, helping to reduce their negative impact on the better partner.

I ran simulations of the experiment using the *interpersonal goal integration* model, which was adjusted to include the effects of compliance on the noise in the partner's goal as the connection strength weakened; the simulations predicted the performance improvement, sup-

porting the hypothesis that the strength of the spring connecting interacting partners modulates the quality of the sensory information of the partner's goal. Furthermore, the simulations also predicted the increased motor force of the better partner, and relaxation in the worse partner. It may seem obvious that the worse partner relaxes, but the *co-activity* mechanism does not predict such relaxation. This tendency to relax when the other partner is better should not be mistaken for a leader-follower paradigm, but it is the by-product of optimal fusion between visual and haptic (i.e. the partner's) information of the goal.

5.1 Connection strength experiment

A dyadic tracking experiment was conducted using the Hi5 dual wrist interface (Figure 5.1a) to systematically test differences in performance improvement during interpersonal sensorimotor interaction with rigid ($1500\,\mathrm{Nm}^{-1}$), medium ($150\,\mathrm{Nm}^{-1}$) and compliant ($25\,\mathrm{Nm}^{-1}$) connection strengths. 32 subjects (16 dyads) were recruited and gave their informed consent to participate in the experiment. The experiment was conducted according to the principles in the Declaration of Helsinki and approved by the Imperial College Research Ethics Committee. Subjects were asked to track a moving goal

$$\theta_t = -7.75\sin(0.12\cdot 2\pi t) - 1.64\sin(0.28\cdot 2\pi t) + 9.35\sin(0.37\cdot 2\pi t) - 10.62\sin(0.63\cdot 2\pi t) \quad (5.1)$$

for 40 seconds. Visual position noise, which was kept fixed for each partner during the whole experiment, was added to the target to augment differences in skill between the partners; the target was corrupted with visual noise to manipulate the performance of subjects. The total experiment lasted 45 trials, which were separated into three equal-sized blocks of 15 trials where the connection strength was different in each block (Figure 5.1b). Dyads always experienced the medium connection first to ensure any learning of the task and interaction occurred at the same connection strength. The order of the second and third blocks were randomised between rigid-compliant and compliant-rigid.

Similar to the experiment of (Ganesh *et al.* 2014), two types of trials (solo and connected) were tested. Performance was measured by calculating the tracking error, or the root-mean squared distance between the cursor and target over the whole trial. By examining the tracking error of two subjects in the previous solo trial, I could assess how much each partner improved as a function of the difference in error, or skill, between the partners.

Surface electromyography of the wrist flexor carpi radialis (FCR) and extensor carpi radialis longus (ECRL) were measured to calculate the effort expended during the tracking task. Prior to the experiment, the EMG signals were calibrated by asking the subject to isometrically contract each muscle to produce constant torques of $1, 2, 3$ and $4\,\mathrm{Nm}$. The torque produced by the muscles were linearly regressed with the EMG to calculate torque-normalised muscular activity from each muscle (for more details on the calibration process, see (Melendez-Calderon *et al.* 2015)). Total effort was taken as the mean torque-normalised activity of the FCR and ECRL over a whole trial. Reciprocal activation (RA) was the mean absolute difference in activity between the FCR and ECRL muscles, which corresponds to motor force. Co-contraction (CC) was taken to be the mean minimum overlap in torque-normalised activity between the FCR and ECRL muscles, which corresponds to joint impedance.

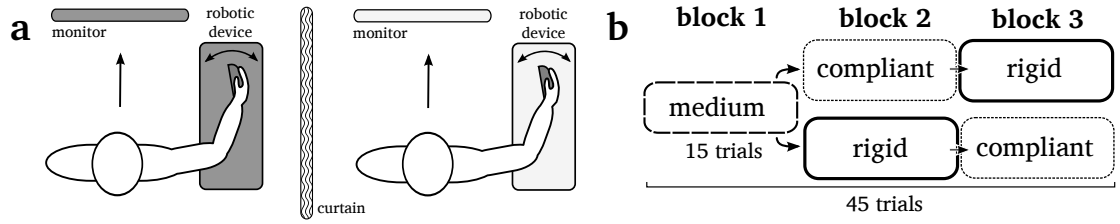


Figure 5.1 | Investigating the connection strength's effects on interpersonal sensorimotor interaction. **a**, Dual-wrist robotic interface setup to test interpersonal interaction in 1 DOF. Partners, separated by a curtain, controlled a cursor by moving their wrist. **b**, The experiment was composed of 3 blocks where the connection strength was different in each block. Each block consisted of 15 trials, populated by equal proportions of solo and connected trials. All dyads experienced the medium block first, then completed either compliant-rigid or rigid-compliant blocks.

Results

The performance improvement was measured by taking the difference in error between a connected trial and the previous solo trial, while the difference in skill was the relative difference in the partners' error in the solo trials. Figure 5.2a shows the performance improvement in all rigid, medium and compliant connection blocks. The worse partner improves more if the connection to the better partner is rigid, which is expected. Surprisingly, the better partner is only slightly hindered by the worse partner. The total improvement of the dyad is shown in Figure 5.2b, revealing that the rigid connection leads to most improvement in both partners.

How was the better partner able to maintain their performance even when rigidly connected with a worse partner? The answer lies in the change in muscular activity in the connected trials. Better partners with a rigid connection expended 50% more effort during interaction and worse partners relaxed (see Figure 5.3a). Approximately half of this total change in effort came from a change in reciprocal activation (motor force) and the other half from co-contraction (joint impedance). With a compliant connection, a gradual linear decrease in extra effort was observed. However, the difference in effort between the connected and solo trials with a compliant spring was small. With a medium strength spring, the worse partner never relaxed during interaction; the better partner exerted more effort with a medium connection but not as much as with the rigid connection.

How does the effort change during the 40 second trial? The mean change in effort was computed at different time instances of the trial. The percentage change is calculated with respect to the mean RA/CC in the previous solo trial. Figure 5.3b shows the change in total mean effort at $1, 5, 20$ and 40 seconds. There are noticeable changes in the effort, the most radical occurring with a rigid connection. Figure 5.3c shows the 'better' and 'worse' partners' reciprocal activation (RA) and co-contraction (CC) as a function of time. In this case, a subject with error 2 degrees less than their partner was considered to be a 'better' partner, and one with an error 2 degrees more than their partner was the 'worse' partner. The first

thing to notice is the simultaneous increase/decrease in RA and CC; clearly, the RA and CC have a close relationship such that an increase in one led to an increase in the other. With the exception of the rigidly connected better partners, the RA and CC change rapidly in the first 10 seconds, but stabilise and do not change during the remaining 30 seconds of the trial. Rigidly connected better partners kept increasing their RA and CC until 20 seconds after the trial began; the worse partners started with the same level of RA and CC at the 2 second mark, but rapidly decreased as the trial progressed. This is an indication that the identification of the partner's control gains and their performance took approximately 10 seconds.

Although the performance improvement appears to be modulated monotonically by the connection strength (Figure 5.2a), the same cannot be said for the change in effort. What does the interpersonal goal integration model predict?

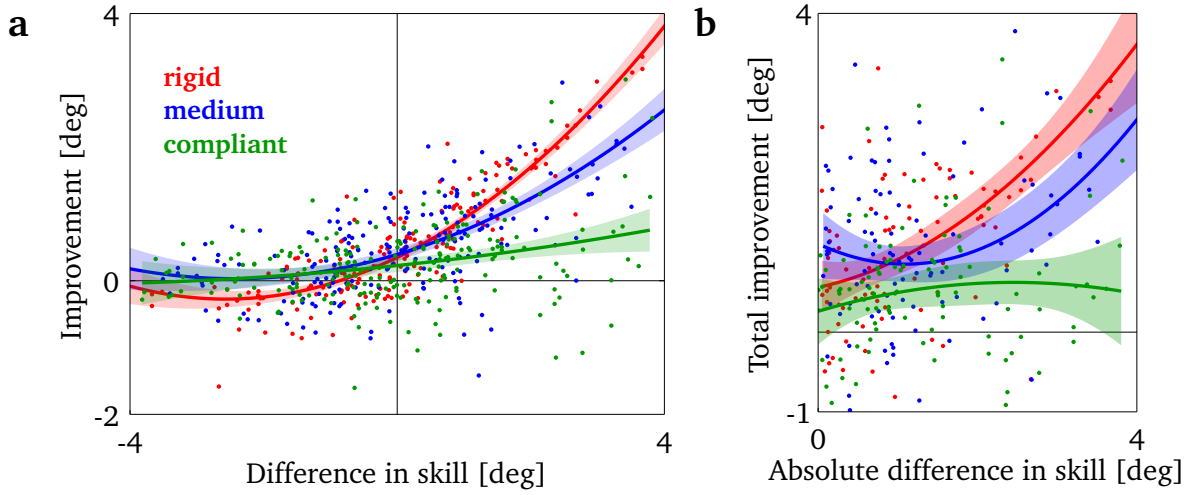


Figure 5.2 | Improvement and effort during rigid, medium and compliant interaction. **a**, Performance improvement as a function of the difference in skill between the partners for rigid, medium and compliant connection strengths. **b**, Total improvement in the dyad as a function of the relative difference in skill between the partners. The dyad as a whole improved the most with a rigid connection.

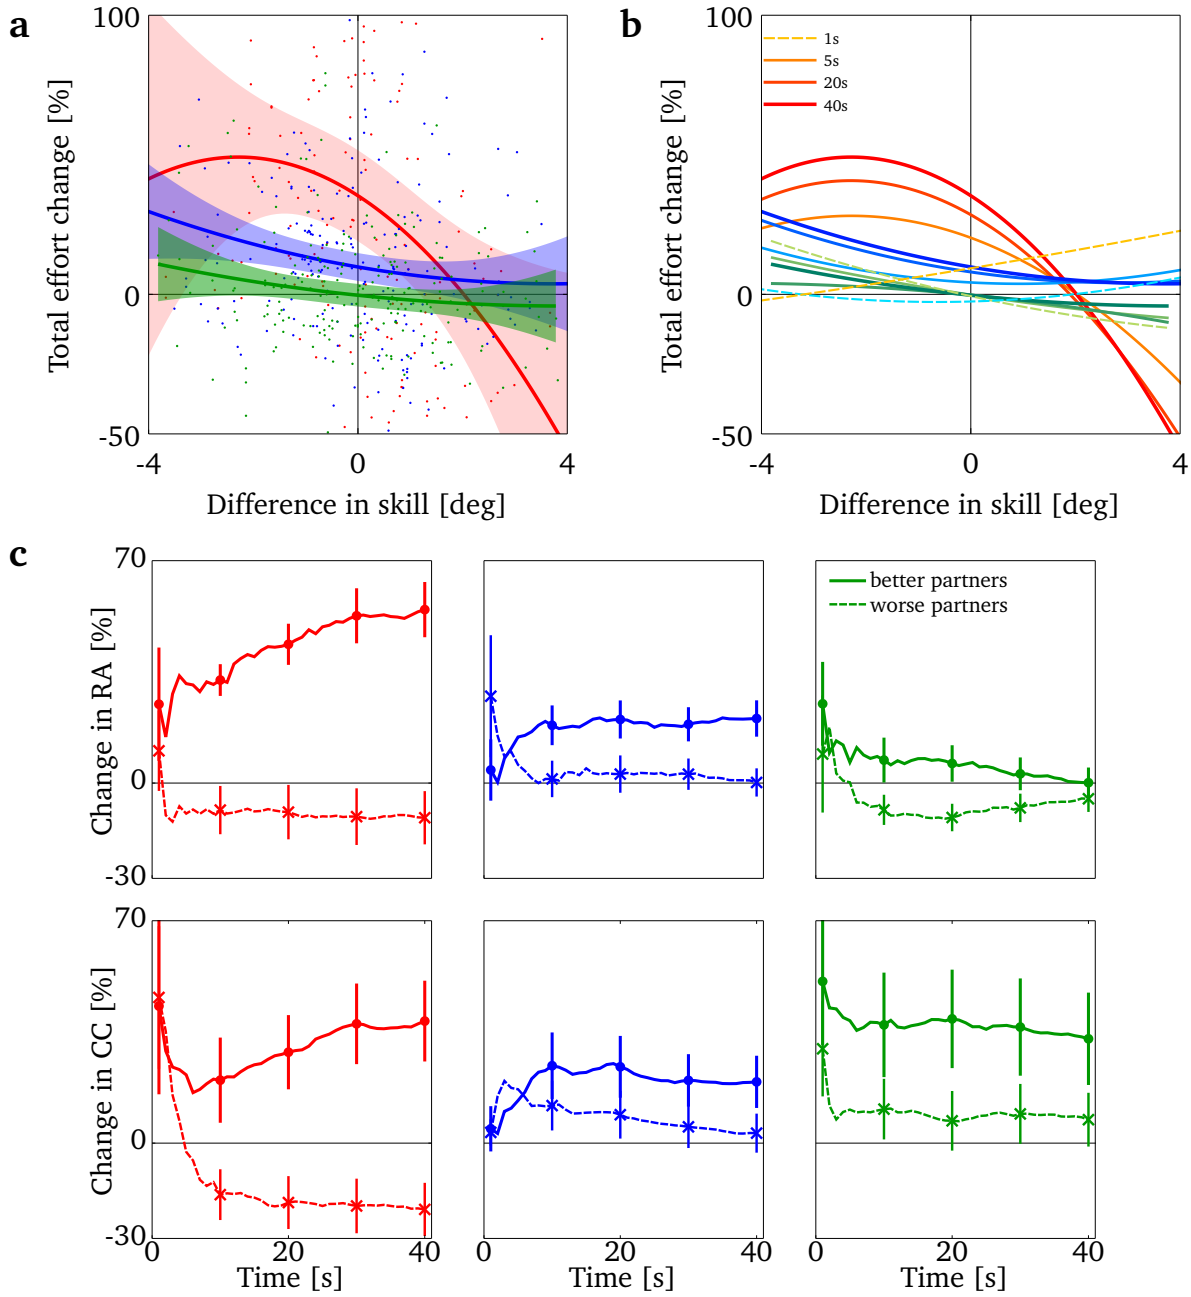


Figure 5.3 | Effort during rigid, medium and compliant interaction. **a**, Change in total effort as a function of the difference in skill. With a rigid connection, worse subjects only relaxed if the difference in skill between the partners was approximately 2 deg. With a medium connection, both better and worse partners spent more effort during interaction. **b**, Effort increase as a function of the difference in skill at different trial times. The second-order polynomial fits do not do the data justice, but one can see the total effort changes greatly as the trial progresses. The next plot **c** demonstrates the quantitative changes that occur. **c**, Change in reciprocal activation (RA) and co-contraction (CC), with respect to the RA and CC in the previous solo trial, as a function of time for better and worse partners (defined as more than 2 deg difference in skill). The better and worse partners have similar RA and CC at the start of the trial. However, radical changes occur as the trial progresses, indicative of subjects identifying the relative error of the partner.

5.2 Connection strength simulation

Let's simulate the results of Section 5.1 by naïvely assuming that the partner's goal can be identified with similar accuracy in all rigid, medium and compliant connections. The interpersonal goal integration model in Section 4.3 was used to simulate two partners in 1 dimension. The parameters were set to

$$dt = 0.005\,\text{s}\,,\ I = 0.002\,\text{kgm}^2,\ \sigma_a^2 = 0.005\,\text{deg}^2,\ \mathbf{Q} = \text{diag}\,[1, 0.005, 0]\ ,\ R = 10^{-10} \qquad (5.2)$$

to match the error ranges observed during the experiment (between 4-8 deg).

Figures 5.4a and 5.4b show the predictions of the performance improvement and change in effort from the *interpersonal goal integration* model assuming that the connection strength does not affect the accuracy in identifying the partner's goal; the modulation of the improvement, which occurred in the data, is not observed and relaxation is predicted in the medium and compliant connection strengths, which does not match the data. The relaxation of the worse partner in the rigid connection may seem an obvious result, but *co-activity*, where partners are only affected by the connection dynamics, does not predict such relaxation (Figure 5.4d); this implies that the spring dynamics alone does not predict the relaxation of the worse partner, and it is only because of the *interpersonal goal integration* model that such relaxation is observed.

The differences between the simulations of the *interpersonal goal integration* model and the data indicate some change in the interaction that cannot be explained by altering the dynamics of the spring alone. So what aspect of interaction changes when the connection strength is modified?

Connection compliance modulates identification of partner's goal

The connection strength directly affects the magnitude of the interaction forces that a dyad experiences. It is more difficult to assess the partner's actions if the forces are weak. To estimate how accurately a subject could sense a goal from the interaction forces, I conducted a haptic tracking experiment; the experimental setup in Figure 5.1a was used but with only one subject sitting on the left interface. 6 subjects were asked to follow a haptic target by feeling the forces from the robotic device. The monitor displayed only the subject's cursor position and not the target. Subjects completed 5 trials per connection strength in the order of rigid, medium then compliant for a total 15 trials. In all trials, subjects received haptic feedback of the target's position through a virtual spring that connected their cursor/wrist to the target. The root-mean squared distance between the target and cursor was taken as the

performance metric of the haptic tracking task.

Figure 5.5a shows the results from the haptic tracking experiment; the horizontal axis corresponds to the additional error that arises from the compliance in the spring. As expected, the weaker the connection strength became, the greater the error in judging the position of the goal was from purely the interaction forces. The additional error from the compliance in the connection strength can be used to modify the interpersonal goal integration model. The predicted performance improvement from the maximum likelihood estimate, assuming that $\sigma = e$, is given by

$$\Delta e_{1,2} = e_{1,2} - \sqrt{\frac{e_1^2\, e_2^2}{e_1^2 + e_2^2}}\,, \tag{5.3}$$

where e_1 and e_2 are the tracking errors of partners 1 and 2 respectively. This improvement is modified to account for the decrease in accuracy in determining the partner's goal due to the connection strength,

$$\Delta e_1 = e_1 - \sqrt{\frac{e_1^2(e_2+e_a)^2}{e_1^2 + (e_2+e_a)^2}}\,, \quad \Delta e_2 = e_2 - \sqrt{\frac{e_2^2(e_1+e_a)^2}{e_2^2 + (e_1+e_a)^2}} \tag{5.4}$$

where e_a is the additional error arising from the compliance in the connection. This additional error can be converted into an equivalent additional noise value $\sigma_{e_a}^2$ and added onto the sensory noise of each partner.

Figure 5.5 shows the simulations from this *modified interpersonal goal integration* model. The modification to the accuracy of identifying the partner's goal in eq. (5.4) successfully predicts the performance improvement in the rigid, medium and compliant connection strengths (Figure 5.5b). Furthermore, the changes in effort for all connection strength in the simulations match the data (Figure 5.5c). This suggests that the same mechanism of *interpersonal goal integration* was at play, but the partner's goal was more difficult to identify with a compliant connection.

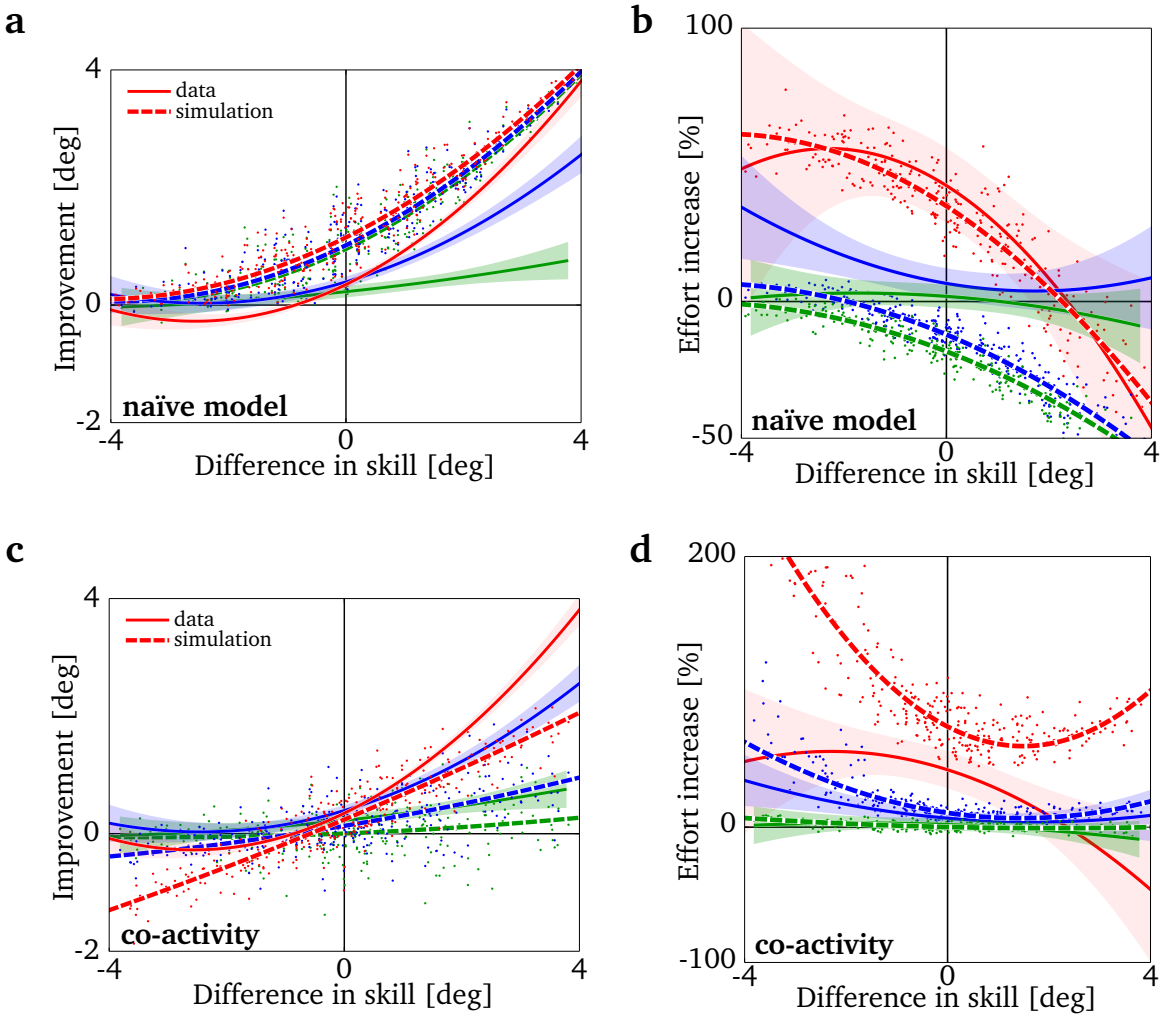


Figure 5.4 | Simulations of naïve interpersonal goal integration model and co-activity with different connection strengths. **a**, The modulation in the performance improvement with connection strength is not observed in the naïve interpersonal goal integration model. **b**, The naïve model also predicts relaxation in the medium and compliant connections in the worse partner, which is inconsistent with the data. **c**, In *co-activity*, the performance improvement monotonically changes with the connection strength and the better partner becomes progressively worse if rigidly connected to a worse partner. **d**, The change in effort predicted by *co-activity* changes monotonically but not as one would expect; the relaxation of the worse partner is not predicted. This is somewhat counter-intuitive as one would expect a change in the gradient of the effort curve such that the worse partner relaxes more and more as the connection strength is increased. However, *co-activity* predicts an increase in effort in both the better and worse partners with a rigid connection.

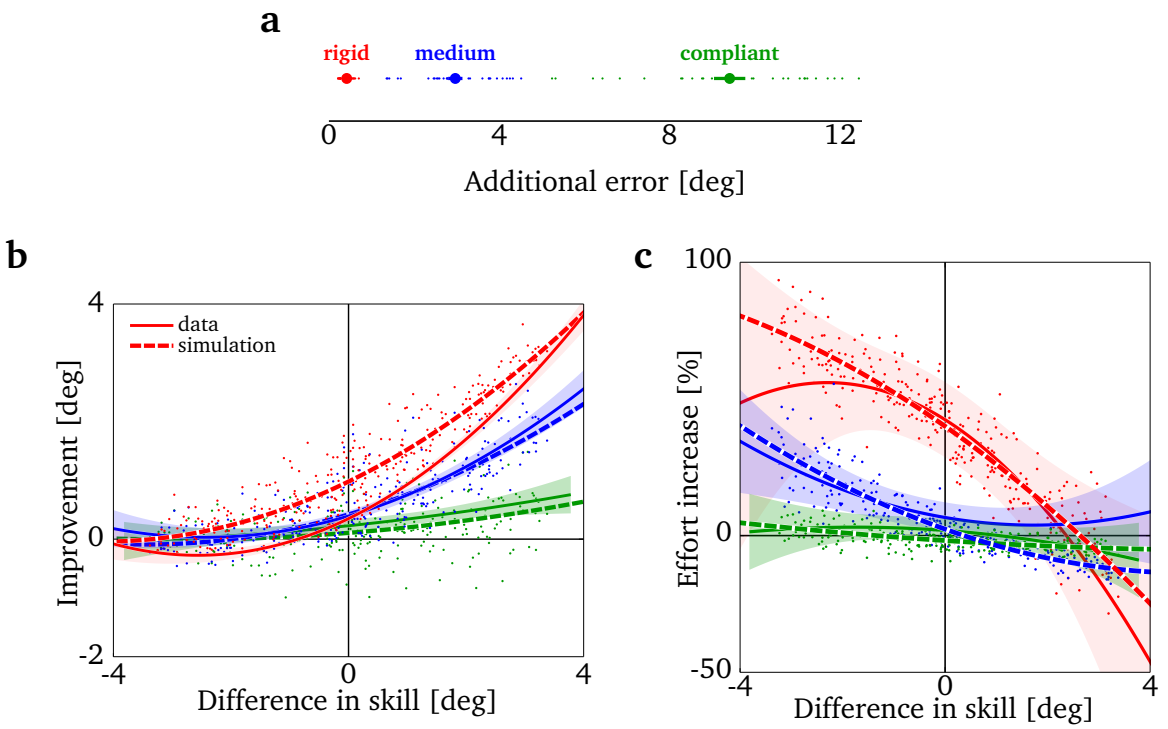


Figure 5.5 | Modified interpersonal goal integration model. **a**, Additional error e_a caused by compliance in connection for rigid, medium and compliant connection strengths. This additional error is used to modify the interpersonal goal integration model to account for the accuracy in determining the partner's goal as a function of connection strength. **b**, Performance improvement prediction from the modified interpersonal goal integration model. The prediction is overly optimistic for the rigid connection but matches the medium and compliant connections well. **c**, Change in effort predicted by the modified interaction model; the rigid and compliant connections are predicted by the model, but the medium connection predicts slightly more relaxation in the worse partner.

5.3 Discussion

The strength of the spring connecting interacting partners had an effect on both the performance of the task and the effort during interaction. The worse partner improved more with a stiffer connection, and because the better partner was not hindered by the interaction, the total improvement of both partners was overall the best with the rigid connection. Analysing the muscular activity, I found that the better partner had significantly greater reciprocal activation and co-contraction (i.e. force and impedance) during interaction, and the worse partner relaxed both RA and CC components of muscle activity. Was this the reason why the better partner could maintain their performance and why the worse partner could improve so much? In fact, the performance and effort arise as a consequence of the interaction mechanism.

The *interpersonal goal integration* model, developed in the earlier Chapter 4, was used to simulate interaction with a rigid, medium and compliant connection. The interaction model was modified to give less accurate measurements of the partner's goal when the connection was compliant. I used a haptic tracking experiment to assess how well subjects could follow a haptic goal with rigid, medium and compliant connection; this enabled me to determine the accuracy in determining the partner's goal as a function of the compliance of the connection, which I used to modify the interaction model. The *modified interpersonal goal integration* model predicted both the changes in performance and the reciprocal activation during interaction for all connection strengths. What this implies is that the change in motor force, such as the relaxation of the worse partner and the increased motor force of the better partner, are not some high-level strategy of deciding to relax and follow a partner; it is a by-product of a strategy whereby the CNS optimally combines visual and haptic information of the goal.

The model is still missing the co-contraction component, or impedance, of the effort. Since changes in the reciprocal activation and co-contraction occurred simultaneously and in roughly the same manner, the co-contraction may be changing according to a simple algorithm where fast changes reciprocal activation would result in increased co-contraction (Franklin *et al.* 2008). This v-shape algorithm could be implemented on top of the existing *interpersonal goal integration* model to adapt the co-contraction when the motor force rapidly switches direction, i.e. when agonist-antagonist muscle pairs rapidly activate asynchronously.

6
Multi-human interaction

The *interpersonal goal integration* model in Chapter 4 predicted that a better partner interacting with a worse partner would either improve or maintain their performance, and should never become worse. An extension of this prediction is that a better partner interacting with two or more partners will never become worse, a prediction that defies intuition.

This prediction was tested by conducting a multi-human tracking experiment, where pairs and groups of 4 tracked the same target to examine sensorimotor interaction with 2, 3 and 4 partners. The partners interacted via a compliant spring. Visual noise was added to the target to manipulate the tracking performance of each partner to create many interaction situations, such as one good partner and three bad ones or four equally bad ones.

This multi-human experiment confirmed the prediction of the interaction model; the best partner was able to maintain the same performance even when connected to 1, 2 or 3 worse partners. The worse partners improved more when the interaction group size was larger. The total improvement of the group, and the expected improvement for each partner, became greater as the group size increased.

There is also evidence to suggest that partners were continuously improving from the interaction. Using a group of 2 partners as an example, partner A would improve thanks to B, but partner B would improve from A as well. The results suggest that cyclical improvement does occur.

6.1 Multi-human experiment

A tracking experiment was conducted using four Phantom 1.5HF interfaces (see Figure 6.1a). Subjects were recruited in pairs (10 pairs, n=20) or in groups of 4 (12 groups, n=48). All subjects gave their informed consent and the experiment was conducted according to the principles in the Declaration of Helsinki and approved by the ethics committee of the Graduate School of Education at the University of Tokyo. Each subject used one of the haptic interface to control a cursor in 2 dimensions, which was displayed on their monitor. All subjects were separated by a curtain and completed 10 training trials at the beginning of the experiment. Pairs experienced solo and connected trials, similar to the protocol used in Chapters 4 and 5, completing a total 80 trials. Groups of 4 were tested to gather interaction data from 3 and 4 partners simultaneously in one experiment; Figure 6.1b shows the protocol used for groups. After the 10 training trials, four permutations of 3 partner interaction were tested and, finally, 4 partner interaction at the end of the experiment for a total 115 trials. Subjects were asked to track the same, pseudo-randomly moving target (in cm) with the motion

$$\begin{aligned} x_t &= 1.6\sin(0.1\cdot 2\pi t_x) + 4\sin(0.3\cdot 2\pi t_x) + 0.8\sin(0.5\cdot 2\pi t_x) + 2.4\sin(0.8\cdot 2\pi t_x) \\ y_t &= 0.8\sin(0.2\cdot 2\pi t_y) + 2.4\sin(0.3\cdot 2\pi t_y) + 4\sin(0.6\cdot 2\pi t_y) + 2.4\sin(0.8\cdot 2\pi t_y) \end{aligned} \tag{6.1}$$

where t_x and t_y are time variables shifted by

$$\begin{aligned} t_x(0) &= [0, 2.63, 3.54, 5, 6.45, 7.36]\,\text{s} \\ t_y(0) &= [0, 1.38, 3.49, 5, 6.5, 8.61]\,\text{s} \end{aligned} \tag{6.2}$$

at the start of the trial to initiate the trajectory at random points to prevent learning of the trajectory. The target was composed of small spots, which were generated around the target with a position and velocity picked from normal distributions with standard deviations of $\sigma_p = 0.5$ cm for the position and from ten uniformly spaced values between $\sigma_v = 0.5$-30 cm/s for the velocity. For each level of visual velocity noise σ_v, the tracking error was calculated for each partner separately during solo trials; this data was linearly regressed to predict a subject's tracking error from just the visual velocity noise on the target (see Figure 6.1c). Figure 6.1b shows the experimental protocol with the 30 solo trials distributed at the beginning, middle and end of the experiment. If learning did occur, the relationship between the tracking error and visual velocity noise should change. However, as Figure 6.1c shows, the relationship between tracking error and visual velocity noise remained fixed throughout the experiment, demonstrating that learning did not take place. Thus, the visual noise on the target is a reliable predictor for the solo tracking error of a subject.

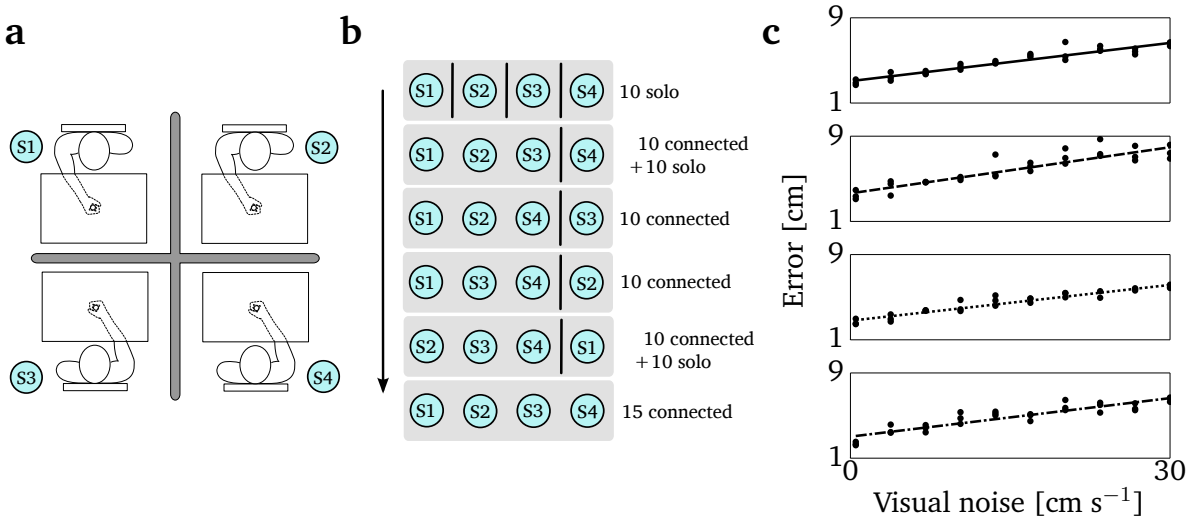


Figure 6.1 | Multi-human sensorimotor interaction. **a**, Experimental setup for physical interaction between two, three and four partners; partners were isolated using a curtain and completed a 15 second tracking task. Each partner held onto the handle of a Phantom 1.5HF manipulandum, which was constrained along a horizontal plane. Subjects were recruited in pairs or in groups of four. **b**, Experimental protocol when testing groups of 4. First, 10 solo trials acted as a training period, then four permutations of groups of 3 were tested. Lastly, 4 person interaction was tested. Solo trials were mixed-in throughout the experiment to ensure no learning of the task occurred. **c**, Tracking error as a function of the visual noise on the target from all four partners in a group of 4; although the disconnected trials are spread uniformly throughout the experiment, no learning of the task was observed. The error can be accurately estimated from the visual noise due to the consistency in the relationship between the two.

Partners were connected by a compliant spring of stiffness of $K = 100\,\mathrm{Nm}^{-1}$. The force experienced at the hand of each subject i was equivalent to the sum of the connection to all other partners,

$$F_i = \sum_{j \neq i}^{N} K\left(p_i - p_j\right) . \tag{6.3}$$

where $N = 2,\ 3,\ 4$ is the total number of partners.

The performance improvement is the difference between the solo error (e_i^{s}) and the connected error (e_i^{c}) for the i^{th} partner. The solo error was not measured, but it was predicted using the regression between the tracking error and visual velocity noise in the solo trials (see Figure 6.1c showing the regression from a sample group),

$$e_i^{\mathrm{s}} = \alpha_i \sigma_{v_i} + \beta_i \tag{6.4}$$

such that the performance improvement was

$$\begin{aligned} \Delta e_i &= e_i^{\mathrm{s}} - e_i^{\mathrm{c}} \\ &= \left(\alpha_i \sigma_{v_i} + \beta_i\right) - e_i^{\mathrm{c}} \end{aligned} \tag{6.5}$$

where σ_{v_i} is the same visual velocity noise as experienced during the connected trial. The mean difference in skill for the i^{th} partner, μ_{d_i}, was calculated as

$$\mu_{d_i} = e_i^{\text{s}} - \frac{1}{N-1}\sum_{j\neq i}^{N} e_j^{\text{s}}\,. \tag{6.6}$$

By plotting the improvement as a function of the mean difference in skill, I can identify how much a partner improved depending on how good the other partners were. Would the addition of worse partners improve performance, as predicted by the *interpersonal goal integration* model, or would it hinder the better partner as intuition leads us to believe?

Results

Figure 6.2a shows the performance improvement as a function of the mean difference in skill for groups of 2 (red), 3 (green) and 4 (blue) partners. Worse partners improved more when connected to more better partners. Incredibly, the best partner in the group was not hindered by more worse partners, but were able to maintain their task performance. The total performance improvement in the whole group was larger as more partners were involved in the interaction (see solid trace Figure 6.2b). The average improvement, or the total improvement divided by the number of partners in the group, was larger for a given mean skill of the group (see dashed trace Figure 6.2b). If two groups had similarly skilled partners in the group, the larger group would lead to greatest improvement in performance.

Does the improvement change as a function of time? Figure 6.2c shows the improvement change as a function of time for groups of 2 only. At first, the improvement is close to the *co-activity* prediction (see Figure 6.3c for *co-activity* predictions) but it moves rapidly away from this prediction, and converges to the final improvement curve. The mean root-mean squared distance between the improvement curve at different times and the *co-activity* prediction was plotted as a function of time (see Figure 6.2d). As the improvement curves rapidly converge and stabilise to a fixed value, exponential curves are fit to the change in improvement Δe of the form

$$\frac{\partial(\Delta e)}{\partial t} = a_1 + a_2 e^{-\lambda t} \tag{6.7}$$

where t is time and λ is the time constant. Figure 6.2d shows the exponential fits and the time constants for groups of 2, 3 and 4; z-tests show that no statistical difference exists between the time constants ($\text{z}_{\lambda_2\lambda_4} = 0.93$ and $\text{z}_{\lambda_3\lambda_4} = 0.95$). This means that identifying the other partners' error is independent of the group size of the interaction, which may indicate that the other partners are treated as a single partner. It also demonstrates the scalability of sensorimotor interaction, in comparison to verbal communication which would take longer to come to a

consensus with more people in the group. The fact that no partners are hindered, and most improve, when the interacting group size is bigger is not intuitive. What could explain this improvement?

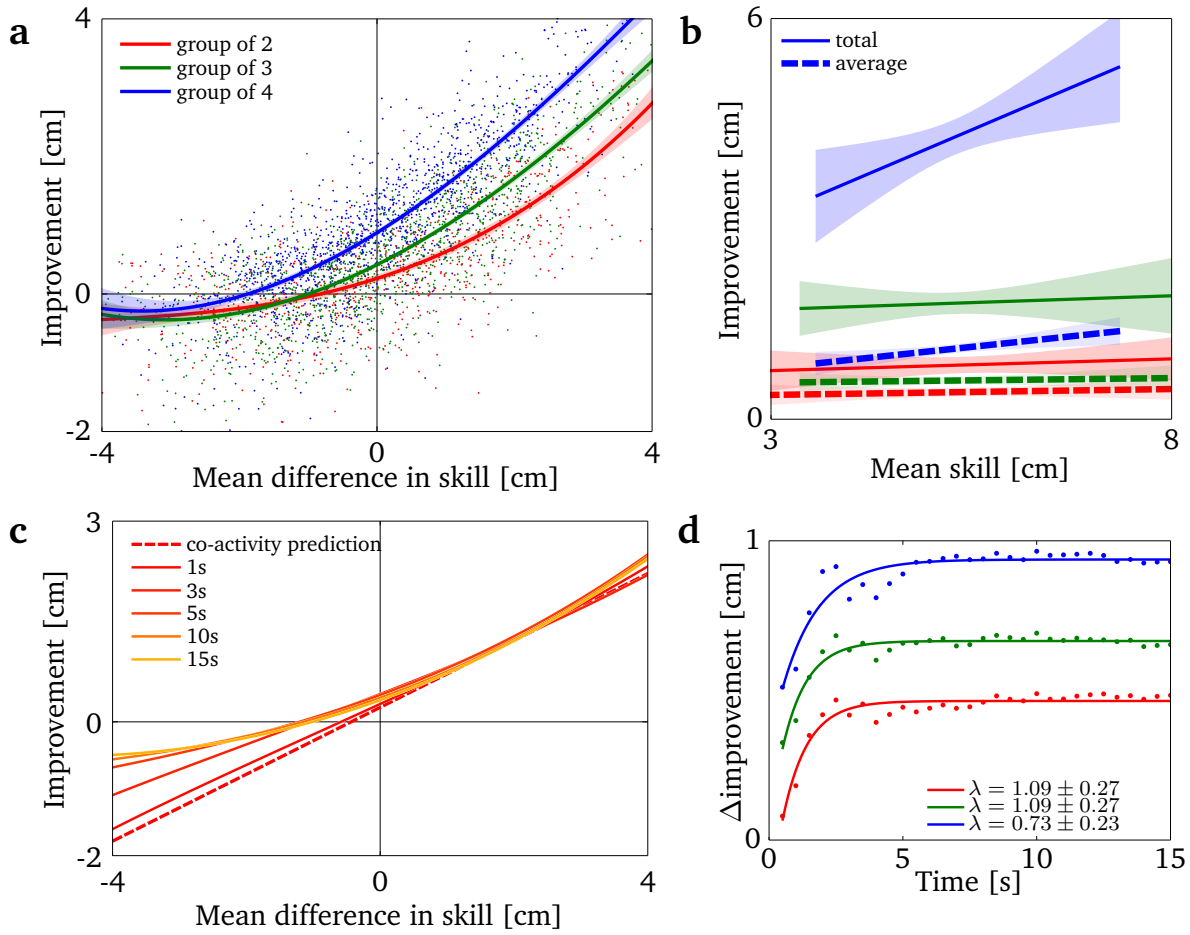


Figure 6.2 | Results from multi-human interaction. **a**, Performance improvement as a function of the mean difference in skill. Incredibly, the better partner is not hindered by additional worse partners. **b**, Total improvement and average improvement per person (solid and dashed traces respectively) as a function of the mean skill of the group. As groups got larger, the total improvement increased. The average improvement was also greater in larger groups. **c**, Change in improvement as the trial progresses. At first, the improvement and the prediction from *co-activity* are the same but the improvement rapidly changes as the partners identify their relative difference in skill. **d**, Change in improvement, with respect to the *co-activity* prediction as a function of time, where the solid line is an exponential fit to the data. The values of λ on the plot show the mean and standard deviation with a 95% confidence interval. Z-tests show that there is no significant difference between the λ for groups of 2, 3 and 4 people ($z_{\lambda_2\lambda_4} = 0.93$ and $z_{\lambda_3\lambda_4} = 0.95$), implying that the time taken to identify the other partners' error is not a function of group size. This demonstrates the scalability of interpersonal sensorimotor interaction as no extra time is taken to come to a consensus, even with more partners in the group.

6.2 Why do more partners help?

A possible reason for larger improvement with more partners could be explained by the normality of the trajectories. One person's trajectory would deviate from the goal, but the average of multiple partners' trajectories would be close to the target trajectory. So are the trajectories of each partner unbiased, or is there a consistent manner in the movement errors produced by humans tracking a target?

Let me examine the likelihood that all partners are on the same side of the target at the same time instance (see Figure 6.3a). The trajectories in the data reveal that these chances are $69.5 \pm 0.5\%$, $55.8 \pm 0.5\%$ and $46.8 \pm 0.5\%$ for 2, 3 and 4 partners respectively. I can simulate trajectories based on the assumption that each trajectory is normally distributed and independent from others; the simulated trajectories reveal that the chances of all partners being on the same side are 50, 25 and 12.5%, which are considerably lower than the data. The higher likelihood of all partners being on the same side of the target is due to the correlated nature of the trajectories (Figure 6.3b are sample trajectories from a solo trial); humans overshoot and undershoot the target in predictable ways with respect to the target. The correlation between the trajectories is found to be $\rho = 0.6$, and simulating normally distributed but correlated trajectories predicts the same likelihood as seen in the data.

What does this imply for the mechanism that a simple averaging of trajectories will result in improvement? Figure 6.3c shows the predicted performance improvement if the trajectories, with a correlation $\rho = 0.6$, from all partners are averaged. The better partner is more hindered by more worse partners, thus trajectory-averaging cannot explain the improvement in performance. This trajectory-averaging is the same as *co-activity* with correlated partners. Other mechanisms, such as follow the better and multi-sensory integration, rely on the worse partner being on the other side of the target to help the better partner. However, the high correlation between the trajectories means that this situation rarely occurs; therefore, follow the better and multi-sensory integration are, as was shown in Chapter 4, incapable of explaining the results in multi-human interaction.

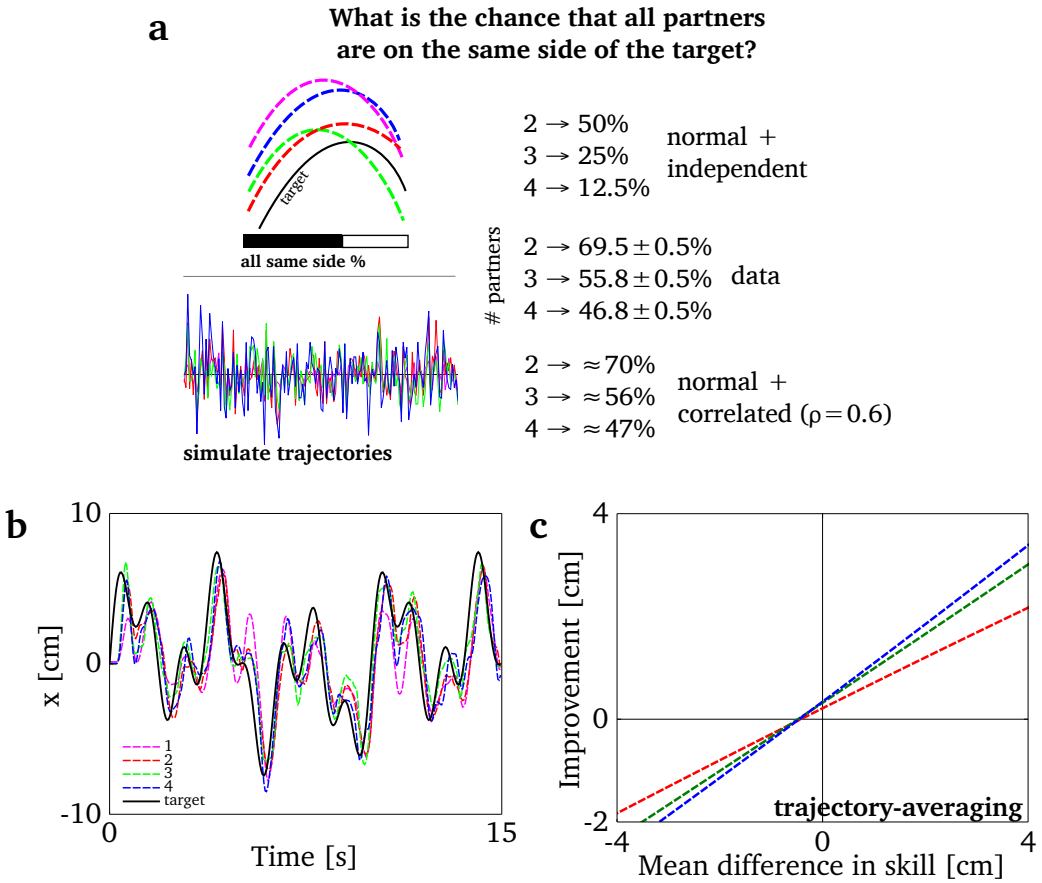


Figure 6.3 | Normality of trajectories does not explain improvement. **a**, The chance of all partners in a group to be on the same side of the target over a whole trial. If the simulated trajectories are normally distributed with no correlation, the chance does not match the data. Trajectories with a correlation of $\rho = 0.6$ between each partner matches the chances in the data. This high correlation, which does not change with the size of the group, explains why *co-activity*, *follow the better* and *multi-sensory integration* cannot explain the improvements. **b**, Sample trajectories from one trial showing the highly correlated trajectories. From this sample, it is clear why taking the mean of these correlated trajectories, or following another partner is a bad strategy as the predictions for the performance improvement shown in **c**.

6.2.1 Interpersonal goal integration for multiple partners

The mechanisms of *co-activity, follow the better* and *multi-sensory integration* cannot explain the improvement observed during sensorimotor interaction with multiple partners. But what does the *interpersonal goal integration* model predict?

The force experienced at the hand, as described in eq. (6.3), is the sum of the compliant connection to each partner. Thus, the goal that subjects identify through the interaction force should be equivalent to a partner whose skill is the average of all other partners. Hence, the i^{th} partner should improve by

$$\Delta e_i = e_i - \sqrt{\frac{e_i^2\, e_j^2}{e_i^2 + e_j^2}} \quad \text{with} \quad e_j = \frac{1}{N-1}\sum_{k\neq j}^{N} e_k \tag{6.8}$$

where e_i is the solo error of the i^{th} partner. Figure 6.4a shows the predictions from the model, which predicts more improvement for groups of 2 and 3. This is likely due to a simultaneous change in the compliance of the connection. Recall in Chapter 5 that the compliance in the connection modulated the improvement during interaction. Likewise, additional partners cause the connection to be stiffer. Unfortunately, the haptic tracking experiment could not be conducted due to limitations in time. Interaction in group sizes of 2, 3 and 4 partners would be equivalent to haptic tracking with a connection to the target of stiffness $100,\ 200,$ and $300\,\text{Nm}^{-1}$ respectively. Figure 6.4b shows predictions from the model when corrected for the compliance by assuming that haptic tracking error is $2\,$cm for $100\,\text{Nm}^{-1}$, $1\,$cm for $200\,\text{Nm}^{-1}$ and $0.5\,$cm for $300\,\text{Nm}^{-1}$. However, the haptic tracking control experiment should be conducted to find the actual errors for each compliance.

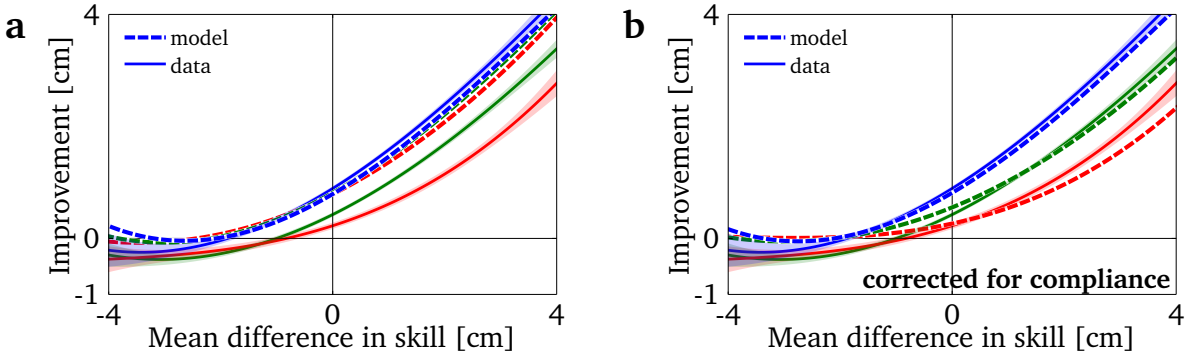


Figure 6.4 | Statistical model of interpersonal goal integration for multi-human interaction. **a**, The statistical predictions for the performance improvement from the interpersonal goal integration model is shown. As the interaction group size increases, the model predicts a slight increase in performance in the worse partner, and the better partner is able to maintain their performance. In comparison to the data, the model predicts less difference in the improvement of the worse partner; this is due differences in the compliance of the connection in groups of 2, 3 and 4. **b** shows the predictions for the model corrected for the compliance with *guesses* for the haptic tracking errors for $100,\ 200,$ and $300\,\text{Nm}^{-1}$ as $2,\ 1$ and $0.5\,$cm respectively.

6.2.2 Other partners' performance modulates improvement

If multi-human interaction is the same as interacting with a single partner whose skill is the average of all partners, the difference in skill between the other partners should not affect the improvement. In other words, if I interacted with 2 other partners and their solo errors were $e_2 = 3\,\text{cm}$ and $e_3 = 5\,\text{cm}$ or $e_2 = 4\,\text{cm}$ and $e_3 = 4\,\text{cm}$, my improvement would still be the same as the average solo error of my partners is equal. Let me define the mean absolute difference between the other partners in the group as

$$d_i = \frac{1}{n(D_i)} \sum_{x \in D_i} x \tag{6.9}$$

where D_i is a set of size $n(D_i)$ that contains all the differences in errors between the other partners,

$$D_i \equiv |e_j - e_k|, \text{ with } j, k \neq i\,. \tag{6.10}$$

Let me give an example: I analyse the improvement of partner 1 in a group of 4. In this case,

$$D_1 = \{\ |e_2 - e_3|,\ |e_2 - e_4|,\ |e_3 - e_4|\ \} \tag{6.11}$$

and

$$d_1 = \frac{1}{3}\left(|e_2 - e_3| + |e_2 - e_4| + |e_3 - e_4|\right)\,. \tag{6.12}$$

I only look at the trials where the other partners 2, 3 and 4 have a mean absolute difference in error greater than $2\,\text{cm}$, i.e. $d_1 > 2\,\text{cm}$. The improvement from only these trials is taken and plotted accordingly. If the mean skill of the other partners is all that matters, the improvement from these trials should be equivalent to the improvement from all trials.

The mean absolute difference between the other partners had no effect on the *interpersonal goal integration* model (see Figures 6.5c and 6.5d), but the data shows that the improvement was greater when the other partners had larger differences in error (Figures 6.5a and 6.5b). Is the interaction model incorrect, or is something missing from it?

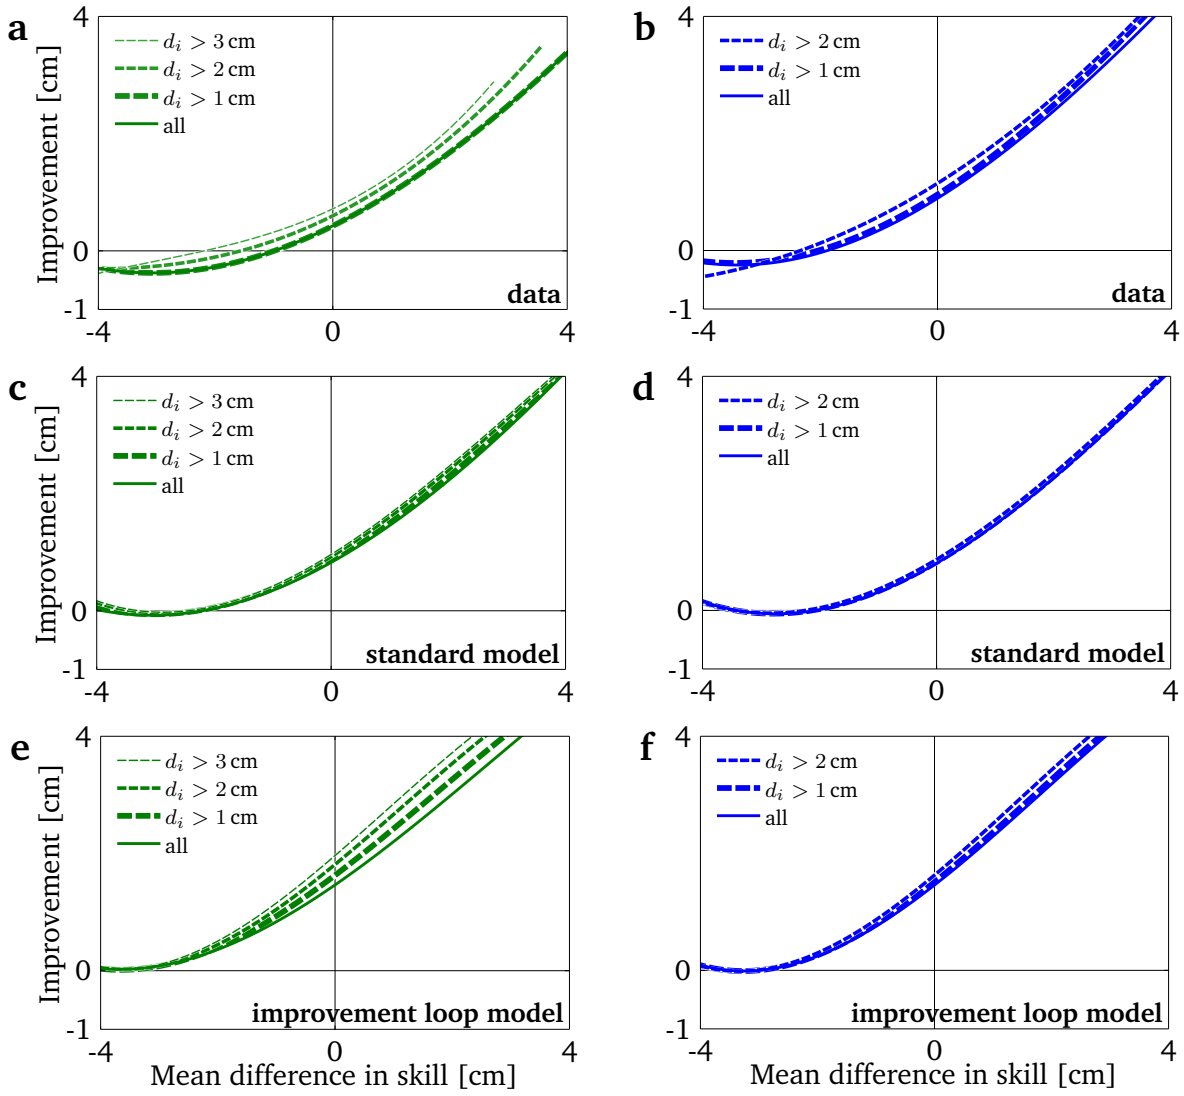


Figure 6.5 | How the other partners' performance affect improvement in the interpersonal goal integration model and data. When interacting in groups of 3 or 4, there is a redundancy in the improvement plot as the mean difference in skill cannot distinguish between 3 equally skilled partners or two good and one bad partner. **(a, b)** are from data, **(c, d)** are predictions from the *standard interpersonal goal integration* model, assuming no improvement loop occurs and **(e, f)** are from the same model but considering the cyclic improvement of the partners. The solid trace is when all points are taken into account and the dashed lines are only from points where the mean absolute difference between the other partners (d_i) is greater than 1, 2 or 3 cm. The data from multi-human interaction is sensitive to the other partners' errors as larger d_i led to greater improvements; this is predicted by the *interpersonal goal integration* model only when the improvement loop is considered.

Improvement loop

What is missing from the *interpersonal goal integration* model? One possibility is that subjects *could* distinguish each partner's actions separately and create an internal model for each partner, essentially demultiplexing the haptic information channel. This multi-goal integration does predict that larger differences in error between the other partners will result in more improvement, but how could the CNS distinguish each partner through one force channel?

The other, more plausible explanation is that partners were continuously improving from the interaction. Taking interaction in pairs as an example, partner A will improve from partner B's goal, but partner B will also improve from A; this results in an improvement loop where each partner continuously improves from the other. The improvement for the i^{th} partner at loop l is given by

$$\Delta e_i = e_i - \left(\frac{1}{e_i^{-1} + \frac{1}{N-1} \sum_{j \neq i}^{N} e_{j,l}} \right)^{1/2} \tag{6.13}$$

where

$$e_{j,l} = \begin{cases} e_j & \text{if } l = 1 \\ \left(\frac{1}{e_j^{-1} + \frac{1}{N-1} \sum_{k \neq i}^{N} e_{k,l-1}} \right)^{1/2} & \text{otherwise.} \end{cases} \tag{6.14}$$

Figures 6.5e and 6.5f show the improvement predicted by the second loop $l = 2$ by the *interpersonal goal integration* model. When only the first loop was considered, corresponding to the maximum likelihood estimate (in Figures 6.5c and 6.5d), there was no difference in the improvement even if the maximum difference between the partners, d_i, was large. However, at the second loop the improvement does distinguish between small and large d_i. Just as in the data, the improvement loop model predicts more improvement with larger d_i or separation between the other partners.

How exactly does such an improvement loop occur? The *interpersonal goal integration* mechanism relies on identifying the partner's goal through the interaction forces. The noise in the partner's goal is a function of the variance in their trajectory, as was shown in the earlier Chapter 4 in Figure 4.6b, such that less variable partners yielded less noisy measurements of the partner's goal. So if partner A improves from the interaction, their variability decreases and so the partner B can extract a less noisy measurement of partner A's goal, enabling partner B to improve even more. I also examined the correlation between the goal and the partners' trajectories as a function of the tracking error of the partner; as expected, the correlation decreases as the subject's tracking error increases, implying that a better partner does yield a more accurate measurement of the goal.

Figure 6.6b shows the improvement for two partners as a function of the number of improvements loops. As the best partner approaches zero error, the improvement converges

after several loops. In theory, partners could have improved significantly more from the interaction if this process could occur indefinitely, but the noise in proprioception and the haptic measurement of the interaction force will impose an upper limit on the improvement from interaction. However, humans interacting with a robot improving in such a manner are likely to improve more than with a human partner, achieving super-human performance. However, they are still limited by the noise in the sensorimotor system.

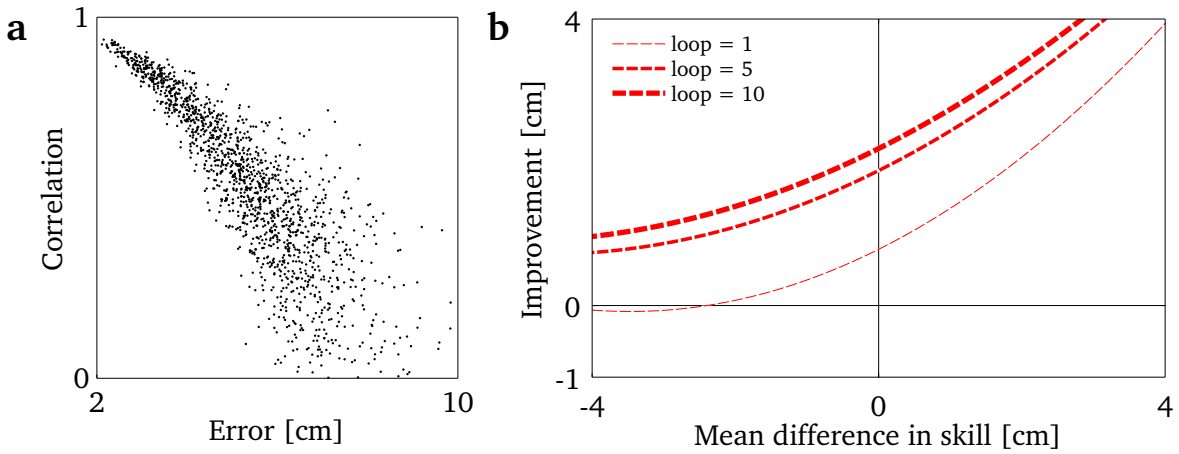


Figure 6.6 | Evidence of improvement loop during interaction. **a**, Correlation between goal and subject trajectories as a function of the subject's tracking error. Any decrease in tracking error will result in an increase in correlation with the goal, such that the goal extracted from an improved subject will be less noisier than before, enabling a continuous improvement loop during interaction. **b**, After every improvement loop, all partners improve slightly more than before. However, the improvement converges as the partners converge to zero tracking error. The improvement loop allows subjects to improve more than the maximum likelihood estimate suggests.

6.3 Discussion

The *interpersonal goal integration* model in Chapter 4 made the counter intuitive prediction that the performance of a better partner would not be hindered even when interacting with two or three worse partners. I tested this prediction by conducting a tracking experiment where groups of 2, 3 or 4 partners interacted through a compliant spring. As predicted by the interaction model, the best partner was not hindered by the addition of extra worse partners. The total improvement of the group was greatest with 4 partners, and a partner's expected improvement was also highest with larger group sizes. The performance improvement could not be explained by simply taking the average of all partners' trajectories, as they were found to be correlated; partners were often on the same side of the target, disproving the explanation that more partners help due to an averaging of trajectories, which converges to the target trajectory. The correlation between the partners' trajectories also explains why follow the better and multi-sensory integration do not predict the mutually beneficial effects of interaction; as the partners are usually on the same side of the target, taking any weighted average between the trajectories does not always result in performance improvement.

The improvement was found to change over time; at the beginning of the trial, the improvement was close to the prediction of *co-activity* as the partners had not identified each others' errors. However, it rapidly changed with time, plateauing about 3 seconds after the start of the trial. No statistical difference was found in how fast the improvement converged, so subjects took the same time to identify the errors of one, two or three other partners. This suggests that subjects do not identify each partner separately, but feel the forces equivalent to a single partner whose skill is the same as the average of all other partners. This demonstrates the scalability of sensorimotor interaction in comparison to verbal communication which would require more time to come to a consensus in larger groups.

I also found that subjects improved more when interacting with one good and one bad partner, rather than two partners with the same skill as the average of the good and bad partners. This could be explained by a model where subjects could distinguish each partner's actions separately, and create an internal model for each partner. Due to the implausibility of demultiplexing the haptic channel, this multi-goal integration hypothesis was rejected. The other explanation was that partners continuously improved from one another. As the noise in the partner's goal is related to the variance in the partner's trajectory, a better partner yields a less noisier measurement of the goal. So if partner A improves from partner B, partner B will improve even more thanks to partner A becoming better. This improvement loop can cause interacting partners to improve more than the maximum likelihood estimate predicts, enabling super-human performance.

7 Conclusion

Our lives begin inside the safety of our mother's womb. During fetal development, we respond more to our mother's touch than to her voice (Marx & Nagy 2015). After birth, our parents physically help us to take our first steps and as we mature, we perfect our coordination to an incredible degree, accomplishing feats of synchrony, like in dance, that are beyond the reach of robots today. Interpersonal sensorimotor interaction plays a critical role in our lives, yet we understand it so little.

I began my investigation into interpersonal sensorimotor interaction by highlighting how cognitive influences such as eye-contact and past relationships could affect the behaviour of partners during interpersonal interaction. I conducted a study where two partners exchanged the same force with one another, and found differences between partners who interacted face-to-face and those who were separated by a curtain. By comparing these conditions to individuals who did the task alone with a robot, it was clear that blocking the partner's gaze was critical to eliminating cognitive influences during interpersonal sensorimotor interaction.

A pioneering study in sensorimotor interaction investigated pairs who rotated a crank as fast as possible to multiple targets in rapid succession (Reed & Peshkin 2008). This study found signs of partners specialising into roles, which the authors proposed were negotiated through the forces. To shed light on roles and their emergence, I repeated their reaching task with dyads who made one movement at a time to analyse the trial-by-trial adaptation of behaviour. My study revealed that dyadic behaviour was consistent and did not change after 30 movements. However, there were significant differences between dyads. These two observations were explained by a mechanism of *co-activity*, where partners independently plan and execute their movement whilst ignoring the forces (Jarrassé *et al.* 2012). The variability between dyads was due to differences in the partners' motion plan. In other words, the "roles" are simply by-products of partners with different motion plans, and do not emerge out of interaction. Therefore, dyadic reaching is not a task of sensorimotor interaction as the partners

do not adapt their behaviour in any way.

In contrast to a dyadic reaching task, a continuous dyadic tracking task with compliantly coupled partners enabled them to feel the partner's movement, but ignore them and do the task independently if they wished; Burdet's group found evidence of adaptations in partners' behaviour during this compliant tracking task (Ganesh *et al.* 2014). This study revealed that both the worse and better partners at the task improved thanks to interaction. Roles or *co-activity* could not reproduce this result, and neither could other mechanisms from psychology and neuroscience. The only mechanism that reproduced this result was one where subjects built an internal model of their partner to extract an extra sensory measurement of the goal from them. As the predictions from this mechanism matched the experimental data, it was formulated as computational algorithm and implemented as a robotic partner. Humans interacting with this robot algorithm improved in a similar manner to those who interacted with a human, thereby validating this mechanism of sensorimotor interaction.

Although this model led to an understanding of the changes occurring on a sensory level, the kinematics cannot reveal adaptations in motor force and impedance. Thus, I repeated the dyadic tracking task on an interface dedicated to measuring electromyography (Melendez-Calderon *et al.* 2011). Large, rapid changes in motor force and impedance were observed; worse partners in the dyad generally relaxed, and the better partners increased both motor force and impedance. Simulations of the mechanism of sensorimotor interaction predicted these changes in the motor force. Hence, the changes in muscular activity did not occur due to some high-level decision to relax and follow a better partner, but are the by-product of the CNS' strategy to optimally combine visual and haptic information of the goal.

Finally, if interpersonal sensorimotor interaction caused both partners in a tracking task to improve, does this generalise to larger groups? I conducted an experiment where groups of 2, 3 and 4 partners interacted together in a tracking task. Incredibly, the benefits of sensorimotor interaction hold true for all group sizes, enabling all partners in a group to either improve or maintain their task performance.

Contributions and future work

In a simple task where two partners exchanged the same force, I found no difference between individuals who did the task alone with a robot, and pairs who were separated by a curtain. Pairs were explicitly told about the interaction with the partner, but behaved similar to individuals; thus, the knowledge of interacting with another subject does not change your behaviour. However, pairs sitting face-to-face without a curtain behaved differently to individuals. In other words, the change in behaviour occurred due to the partner's gaze, whose effect was modulated by the partner's sex. These results challenge the interpretation of Shergill *et*

al. who claimed that the sensory cancellation of self-produced actions caused pairs to change their behaviour (Shergill *et al.* 2003). Luckily for researchers in sensorimotor interaction, my results imply that blocking the partner's gaze is sufficient to eliminate the social, cognitive influences of interpersonal interaction.

Although the study of Reed *et al.* (2008) that examined dyadic reaching movements was pioneering, my experiments revealed that the roles they saw were a mere by-product of two independently moving partners who ignored the interaction force. Even signs of effort optimisation were missing; partners stuck to one successful motor behaviour or motor memory which was globally suboptimal (Ganesh & Burdet 2013). It's likely that partners could not optimise globally as they had insufficient knowledge of the system; if dyadic reaching was done bimanually, where the motor force of each hand (or "partner") is known to the CNS, subjects could have eliminated all wasted forces that occur between the hands. The lack of adaptations in behaviour in the dyadic reaching task demonstrates that sensorimotor interaction arises only when there is something to gain from the partner, and will not emerge by simply coupling people together. Sensorimotor interaction occurs in the dyadic tracking task because useful sensory information can be gained from the partner's actions.

What information was gained from sensorimotor interaction, and how was it obtained? The only mechanism to explain the mutually beneficial effects of dyadic tracking was one where subjects formed an internal model of the partner's task to extract their goal; the partner's goal was optimally combined with own visual estimate of the goal, resulting in improved task performance even when interacting with a partner worse than oneself. This *interpersonal goal integration model* also predicted the effect of changing the coupling strength between partners on the improvement and motor force during interaction. Furthermore, it could predict the outcome of sensorimotor interaction with multiple partners.

It is unclear how the internal model of the partner could be implemented by the brain. Future work would build on the experimental observations and computational models provided here to analyse the neural circuits involved in processing the haptic information coming from the partner's movement, perhaps using transcranial magnetic stimulation or functional magnetic resonance imaging.

Bibliography

1. Argyle, M. & Dean, J. Eye-Contact, Distance and Affiliation. *Sociometry* **28,** 289–304 (Sept. 1965).
2. Bahrami, B. *et al.* Optimally Interacting Minds. *Science* **329,** 1081–1085 (Aug. 2010).
3. Bailenson, J, Blascovich, J, Beall, A & Loomis, J. Equilibrium Theory Revisited: Mutual Gaze and Personal Space in Virtual Environments. *Presence* **10,** 583–598 (Dec. 2001).
4. Baron-Cohen, S. The extreme male brain theory of autism. *Trends in Cognitive Sciences* **6,** 248–254 (June 2002).
5. Basdogan, C., Ho, C.-H., Srinivasan, M. A. & Slater, M. An Experimental Study on the Role of Touch in Shared Virtual Environments. *ACM Transactions on Computer-Human Interaction* **7,** 443–460 (2000).
6. Blakemore, S.-J., Wolpert, D. M. & Frith, C. D. Central cancellation of self-produced tickle sensation. en. *Nature Neuroscience* **1,** 635–640 (Nov. 1998).
7. Bock, O. Load compensation in human goal-directed arm movements. en. *Behavioural Brain Research* **41,** 167–177 (Dec. 1990).
8. Buchanan, T. S. & Lloyd, D. G. Muscle activity is different for humans performing static tasks which require force control and position control. *Neuroscience Letters* **194,** 61–64 (July 1995).
9. Conditt, M. A., Gandolfo, F. & Mussa-Ivaldi, F. A. The Motor System Does Not Learn the Dynamics of the Arm by Rote Memorization of Past Experience. en. *Journal of Neurophysiology* **78,** 554–560 (July 1997).
10. Dingwell, J. B., Mah, C. D. & Mussa-Ivaldi, F. A. Experimentally Confirmed Mathematical Model for Human Control of a Non-Rigid Object. en. *Journal of Neurophysiology* **91,** 1158–1170 (Mar. 2004).

11. Dingwell, J. B., Mah, C. D. & Mussa-Ivaldi, F. A. Manipulating Objects With Internal Degrees of Freedom: Evidence for Model-Based Control. en. *Journal of Neurophysiology* **88,** 222–235 (July 2002).

12. Ernst, M. O. & Banks, M. S. Humans integrate visual and haptic information in a statistically optimal fashion. en. *Nature* **415,** 429–433 (Jan. 2002).

13. Evans, G. W. & Howard, R. B. Personal space. *Psychological Bulletin* **80,** 334–344 (1973).

14. Franklin, D. W. *et al.* CNS Learns Stable, Accurate, and Efficient Movements Using a Simple Algorithm. en. *The Journal of Neuroscience* **28,** 11165–11173 (Oct. 2008).

15. Ganesh, G. *et al.* Two is better than one: Physical interactions improve motor performance in humans. en. *Scientific Reports* **4** (Jan. 2014).

16. Ganesh, G. & Burdet, E. Motor planning explains human behaviour in tasks with multiple solutions. *Robotics and Autonomous Systems. Models and Technologies for Multi-modal Skill Training* **61,** 362–368 (Apr. 2013).

17. Groten, R. *et al. Experimental analysis of dominance in haptic collaboration* in *The 18th IEEE International Symposium on Robot and Human Interactive Communication, 2009. RO-MAN 2009* (Sept. 2009), 723–729.

18. Hastie, R. & Kameda, T. The Robust Beauty of Majority Rules in Group Decisions. *Psychological Review* **112,** 494–508 (2005).

19. Jarrassé, N., Charalambous, T. & Burdet, E. A Framework to Describe, Analyze and Generate Interactive Motor Behaviors. *PLoS ONE* **7,** e49945 (Nov. 2012).

20. Kenny, D. A. & Acitelli, L. K. Accuracy and bias in the perception of the partner in a close relationship. *Journal of Personality and Social Psychology* **80,** 439–448 (2001).

21. Körding, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. en. *Nature* **427,** 244–247 (Jan. 2004).

22. Laughlin, P. R. & Ellis, A. L. Demonstrability and social combination processes on mathematical intellective tasks. *Journal of Experimental Social Psychology* **22,** 177–189 (1986).

23. Little, K. B. Personal space. *Journal of Experimental Social Psychology* **1,** 237–247 (Aug. 1965).

24. Ljung, L. Asymptotic behavior of the extended Kalman filter as a parameter estimator for linear systems. *IEEE Transactions on Automatic Control* **24,** 36–50 (Feb. 1979).

25. Maccoby, E. E. *The Two Sexes: Growing Up Apart, Coming Together* en (Harvard University Press, 1998).

26. Madan, C., Kucukyilmaz, A., Sezgin, T. & Basdogan, C. Recognition of Haptic Interaction Patterns in Dyadic Joint Object Manipulation. *IEEE Transactions on Haptics* **8,** 54–66 (Jan. 2015).

27. Marx, V. & Nagy, E. Fetal Behavioural Responses to Maternal Voice and Touch. *PLoS ONE* **10,** e0129118 (June 2015).

28. Mehra, R. Approaches to adaptive filtering. *IEEE Transactions on Automatic Control* **17,** 693–698 (Oct. 1972).

29. Mehrabian, A. Relationship of attitude to seated posture, orientation, and distance. *Journal of Personality and Social Psychology* **10,** 26–30 (1968).

30. Melendez-Calderon, A., Bagutti, L., Pedrono, B. & Burdet, E. *Hi5: A versatile dual-wrist device to study human-human interaction and bimanual control* in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on* (Sept. 2011), 2578 –2583.

31. Melendez-Calderon, A., Komisar, V. & Burdet, E. Interpersonal strategies for disturbance attenuation during a rhythmic joint motor action. *Physiology & Behavior* **147,** 348–358 (Aug. 2015).

32. Reed, K. *et al.* Haptically Linked Dyads: Are Two Motor-Control Systems Better Than One? *Psychological Science* **17,** 365–366 (May 2006).

33. Reed, K. & Peshkin, M. Physical Collaboration of Human-Human and Human-Robot Teams. *IEEE Transactions on Haptics* **1,** 108–120 (July 2008).

34. Schilbach, L. *et al.* Toward a second-person neuroscience. *Behavioral and Brain Sciences* **36,** 393–414 (Aug. 2013).

35. Sebanz, N, Bekkering, H & Knoblich, G. Joint action: bodies and minds moving together. *Trends in Cognitive Sciences* **10,** 70–76 (Feb. 2006).

36. Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. en. *The Journal of Neuroscience* **14,** 3208–3224 (May 1994).

37. Shergill, S. S., Bays, P. M., Frith, C. D. & Wolpert, D. M. Two Eyes for an Eye: The Neuroscience of Force Escalation. en. *Science* **301,** 187–187 (July 2003).

38. Simon, D. *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches* en (John Wiley & Sons, June 2006).

39. Stefanov, N., Peer, A. & Buss, M. *Role determination in human-human interaction* in *EuroHaptics conference, 2009 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009. Third Joint* (Mar. 2009), 51 –56.

40. Stengel, R. F. *Optimal Control and Estimation* en (Courier Corporation, Oct. 2012).

41. Todorov, E. Stochastic Optimal Control and Estimation Methods Adapted to the Noise Characteristics of the Sensorimotor System. *Neural Computation* **17,** 1084–1108 (2005).

42. Van der Wel, R. P. R. D., Knoblich, G. & Sebanz, N. Let the force be with us: Dyads exploit haptic coupling for coordination. *Journal of Experimental Psychology: Human Perception and Performance* **37,** 1420–1431 (2011).

43. Vogeley, K. & Fink, G. R. Neural correlates of the first-person-perspective. *Trends in Cognitive Sciences* **7,** 38–42 (Jan. 2003).

44. Walsh, L. D., Taylor, J. L. & Gandevia, S. C. Overestimation of force during matching of externally generated forces. en. *The Journal of Physiology* **589,** 547–557 (2011).