Artificial Partners to Understand Joint Action: Representing Others to Develop Effective Coordination

Cecilia De Vicariis, Giulia Pusceddu, Vinil T. Chackochan, and Vittorio Sanguineti

***Abstract*— In the last years, artificial partners have been proposed as tools to study joint action, as they would allow to address joint behaviors in more controlled experimental conditions. Here we present an artificial partner architecture which is capable of integrating all the available information about its human counterpart and to develop efficient and natural forms of coordination. The model uses an extended state observer which combines prior information, motor commands and sensory observations to infer the partner's ongoing actions (partner model). Over trials, these estimates are gradually incorporated into action selection. Using a joint planar task in which the partners are required to perform reaching movements while mechanically coupled, we demonstrate that the artificial partner develops an internal representation of its human counterpart, whose accuracy depends on the degree of mechanical coupling and on the reliability of the sensory information. We also show that human-artificial dyads develop coordination strategies which closely resemble those observed in human-human dyads and can be interpreted as Nash equilibria. The proposed approach may provide insights for the understanding of the mechanisms underlying human-human interaction. Further, it may inform the development of novel neuro-rehabilitative solutions and more efficient human-machine interfaces.**

***Index Terms*— Joint action, human–robot interaction, partner model, game theory.**

I. Introduction

JOINT action is pervasive in our daily life – two persons carrying a heavy object together or folding a sheet, a therapist interacting physically with a patient are just a few examples [1]. Joint action implies a dynamic interplay between self and other [2]; within an ensemble, individuals adapt their behaviors to adeptly exchange information and coordinate with others [3], [4]. As their behaviors are interdependent, it is difficult to experimentally disentangle the individual contributions and the way each participant reacts to their partner's actions.

A way to overcome this problem is to design experiments in which human participants interact with simulated, human-like artificial partners (AP), whose personal traits and interaction strategies can be manipulated experimentally [5]. This approach is often referred to as Virtual Partner Interaction (VPI) [6] or Human Dynamic Clamp (HDC) [5], as it resembles the dynamic clamp technique in cellular neuroscience [7] which uses virtual, simulated ion channels with specific biophysical properties to understand the complexity of neuron dynamics. The original HDC approach relied on a theoretical model of coordination dynamics which involves coupled non-linear oscillators – the Haken-Kelso-Bunz (HKB) model [8]. Originally introduced to model the dynamics of the relative phase between two fingers or limbs performing rhythmic movements, the HKB model was later extended to capture basic aspects of social coordination between two individuals. Early applications of the HDC concept [5], [6] focused on real-time, rhythmic bidirectional interaction and involved human subjects coordinating their hand movements with an avatar. In conjunction with high-density EEG recordings, this paradigm has been used to investigate the neural foundations of social interaction [9]. The VPI/HDC paradigm has provided important contributions to the study of discrete joint action. However, crucial determinants of sensorimotor interaction, like perception, decision-making and control mechanisms need to be explicitly modeled to develop more versatile artificial agents. APs have been also used in applications involving decision-making, with no actual movements. For instance, when a human participant interacts with a simulated partner whose risk-sensitivity is modulated depending on model uncertainty [10].

Here we propose a biomimetic artificial partner (AP) architecture which can develop collaborative strategies with a human partner. The model builds upon previous work [11] on computational modeling of joint action. In particular, we assume that the AP behaves optimally at both perceptual and control level. These assumptions rely on a large body of evidence from perceptual and motor control literature [12], [13], but have been rarely applied to joint action [14]. We specifically argue that individuals involved in a joint action builds an internal representation of their partner's ongoing actions and/or intentions (partner model) and use it to establish an interaction [11]. The proposed AP architecture can be used to compare actual (human) and idealized (artificial) partners

Manuscript received November 1, 2021; revised March 30, 2022; accepted May 13, 2022. Date of publication May 18, 2022; date of current version June 3, 2022. *(Corresponding author: Cecilia De Vicariis.)*

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Competent Ethical Committee (Comitato Etico Regione Liguria), and performed in line with the Declaration of Helsinki.

Cecilia De Vicariis, Giulia Pusceddu, and Vittorio Sanguineti are with the Department of Informatics, Bioengineering, Robotics and Systems Engineering, University of Genoa, 16145 Genoa, Italy (e-mail: ceciliadevicariis@gmail.com; giulia.pusceddu96@gmail.com; vittorio.sanguineti@unige.it).

Vinil T. Chackochan is with the Biomedical Systems Engineering, Bournemouth University, Poole BH12 5BB, U.K. (e-mail: vthekkedathchackocha@bournemouth.ac.uk).

Digital Object Identifier 10.1109/TNSRE.2022.3176378

in order to test different assumptions on how we develop a coordination. We assessed the AP performance in an interaction paradigm where two agents are mechanically connected through a virtual spring, and are instructed to perform reaching movements from the same start and target points, but through different via-points.

II. Modeling Framework

In this section we provide the background ideas, and the empirical basis for the proposed AP architecture. We specifically focus on (i) optimality in perception and control; (ii) the need to represent partners in joint action; and (iii) the way a coordination strategy can be developed through repeated practice.

A. Optimality in Perception and Action

Sensorimotor control is the end result of two inter-related processes, i.e. estimation of the state of the body and the external environment, and movement control. Motor commands - muscle activations - are generated on the basis of our movement goals and our belief in the current state (position, velocity) of our own body and the external environment. Both state estimation and movement control can be understood in terms of optimality principles. There is indeed ample evidence that our nervous system exhibits a close-to-optimal performance in both sensorimotor integration [12] and movement control [13]. Hence the combination of optimal estimation and optimal control is the ideal framework for the development of biomimetic artificial partners.

Sensorimotor integration is believed to optimally combine prior beliefs and sensory feedback - e.g. visual, acoustic and/or haptic - to minimize the prediction uncertainty [12]. Prior belief about body and environment dynamics is believed to rely on an internal representation of the causal relationship between the motor command and the body and environment state – the notion of 'forward model' of the body – and between the body state and its sensory consequences – a model of the sensory system, or 'sensory model'.

Optimal control posits that humans aim at maximizing their subjective utility, i.e. a trade-off between task-dependent movement cost – in reaching movements, it may simply be the endpoint error – and the perceived effort associated to movement. Hence actions result from a subjective evaluation of their associated costs and/or benefits. This subjective aspect can be summarized by a cost function that incorporates the task goals and requirements and accounts for effort [15], [16]. It is natural to extend this framework to joint action [11], [17], [18].

B. Partner Representation

Sharing information is a key determinant of successful joint action [3]. Information exchange implies verbal or non-verbal communication. If the participants are physically connected, haptic channel constitues a powerful source of information about the partner [11], [19]–[21]. More reliable information generally leads to more efficient interaction [11]. However, stronger coupling provide more information but makes control more difficult, whereas weaker coupling facilitates coordination but provides less reliable information about partner actions [21].

During interactions, individuals not only observe and react to partner actions but they also need to predict the actions of their partner and their expected consequences [3], [22]. To this purpose, joint action crucially requires that participants develop an internal representation of their partner or, using terminology from the sensorimotor control field [23], a 'partner model'.

C. Game Theory and Fictitious Play

In a joint action, the subjective utility of a participant typically depends on the state of all partners involved. Game theory provides the analytic and computational substrate for the underlying control processes . Therefore, it can be seen as an extension of optimal control to multi-agent scenarios. A game involves two or more individuals whose interests are neither completely opposed, nor completely coincident. In some joint action scenarios, agents agree on a shared strategy – for instance, through verbal communication – before the action takes place, thus behaving as a collective [24]. In other scenarios there is no explicit prior agreement on a shared strategy. In these situations, coordination emerges gradually as each agent collects information about their opponent's actions, their outcomes and possibly their ultimate goals, using various mechanisms [3] during continuous or repeated interaction. In motor control scenarios, differential non-cooperative games [25] can model situations in which humans deal with their counterpart without speaking and by communicating just through sensory cues (visual, acoustic or haptic), but they independently determine their actions. It has been observed [11], [17], [18], [26] that when both agents have perfect information about the partner and the environment, they converge to a Nash equilibrium, i.e. a situation in which none of the players can unilaterally improve their benefit [27]. When the players have competing goals, coordination develops gradually. They gradually gain knowledge about dyad dynamics, the task requirements, and partner's actions. Hence coordination (if any), is a result of learning and adaptation.

One possible solution of the problem of iteratively calculating a Nash equilibrium is represented by the classical learning process known as fictitious play [28], [29]. In fictitious play, two agents play the game repeatedly. After arbitrary initial moves in the first round, in every round each agent determines its best response against the empirical strategy distribution of their partner. Fictitious play posits that the players use a stationary strategy and only requires a model of the strategy distribution. However, the players do not have to know anything at all about their opponent's payoffs. All they do is to form beliefs about how their opponents will play [30]. Because of its minimal requirements, fictitious play has been proposed as a candidate mechanism for the development of optimal coordination in joint action, either involving movements [11] or not [10].

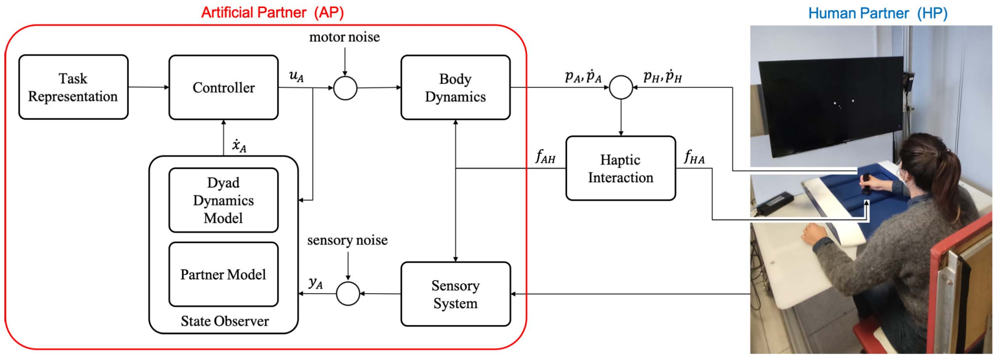

Fig. 1. Artificial partner architecture. The artificial partner has a body and a sensory system. The control architecture includes a task representation, a feedback controller and a state observer, which also includes a partner model. This is the minimal architecture to support the development of an interaction. The control architecture is integrated into a robotic platform which supports the physical interaction with the human partner.

III. Materials and Methods

A. Artificial Partner Architecture

The proposed artificial partner architecture directly derives from the computational model of joint action proposed in [11]. As in 'solo' behavior, each motor command is generated by a feedback controller which relies on a state observer which predicts the overall dyad state. The only addition to the 'solo' model is that, in order to correctly account for dyad dynamics, the state observer also needs to estimate the partner's ongoing actions and/or movements. Based on the above computational model, a minimal but general AP architecture – see Figure 1 – should involve a task representation, a feedback controller, a state observer and a partner model. Interaction with its human counterpart also requires that the AP have its own body and sensory system. These can take different forms. The AP body may be physical – for instance, a humanoid robot – or virtual – an avatar moving in a VR or AR environment. Interaction may be visual – this requires that the AP is equipped with an artificial visual system – and/or physical – in this case, the AP may also involve a haptic interface, which exchanges interaction forces with the human player. Haptic channel provides additional information about partner movements, hence it is also part of AP's sensory system. Although this AP concept is quite general, in this study we demonstrate it in a simplified situation. Both human and artificial partners move on a horizontal plane. The human partner (HP) grasps the handle of a planar robotic manipulandum, through which the AP and HP are physically coupled. Although very simple from the geometric and mechanical point of view, this scenario is central in many studies on computational motor control and in many rehabilitation robotics applications.

1) Dyad Dynamics: We describe the dynamics of both the human and the artificial partner as two planar arms with two degrees of freedom (q_s for shoulder, q_e for elbow). We linearized arm dynamics around configuration $q_0 = [45°, 90°]$ and expressed it in Cartesian space:

$$\begin{aligned} M_H \ddot{p}_H + b_H \dot{p}_H &= f_H + f_{AH} \\ M_A \ddot{p}_A + b_A \dot{p}_A &= f_A + f_{HA} \end{aligned} \tag{1}$$

where p_A and p_H are the hand position vectors of, respectively, the artificial (A) and the human (H) partner and M_A and M_H are their respective inertia matrices. We set $M_i = J_i(q_0)^{-T} M_i(q_0) J_i(q_0)^{-1}$ with $i \in \{A, H\}$, where $J_i(q)$ is the jacobian of the forward kinematic transformation and $M_i(q)$ is the arm inertia matrix in joint space. Eq. 1 reflects the assumption that both partners are subjected to a small viscous force accounting for muscles and soft tissue damping. The interaction force f_{HA}:

$$f_{HA} = k \cdot (p_H - p_A) \tag{2}$$

is applied to the AP. An opposite force $f_{AH} = -f_{HA}$, is applied to the human partner through the robot handle. The position of the human partner, p_H, coincides with the position of the robot handle. As in [13], we modelled the dynamics of muscle force generation (f_A and f_H) as a second order system:

$$\begin{aligned} \tau_H^2 \ddot{f}_H + 2\tau_H \dot{f}_H + f_H &= u_H \\ \tau_A^2 \ddot{f}_A + 2\tau_A \dot{f}_A + f_A &= u_A \end{aligned} \tag{3}$$

where u_A and u_H are the activation vectors of human and artificial partners, which are taken as the dyad inputs, and τ_A, τ_H are the activation time constants (we set $\tau_A = \tau_H = 40$ ms). Eqs. 1, 2 and 3 can be expressed in state-space form by defining $x = \left[x_A^T \; x_H^T\right]^T$, as the state vector, where $x_i = [p_i^T \; \dot{p}_i^T \; f_i^T \; \dot{f}_i^T]^T$, with $i \in \{A, H\}$. The system has two control inputs, namely AP and HP muscle activations (u_A and u_H). In state-space form, Eqns. 1, 2 and 3 can be expressed as:

$$\dot{x} = \begin{bmatrix} A_{AA} & A_{AH} \\ A_{HA} & A_{HH} \end{bmatrix} \cdot x + \begin{bmatrix} B_A \\ 0 \end{bmatrix} \cdot (u_A + w_A) + \begin{bmatrix} 0 \\ B_H \end{bmatrix} \cdot (u_H + w_H) \tag{4}$$

where $w_A \sim N(0, \Sigma_A^w)$ and $w_H \sim N(0, \Sigma_H^w)$ are the 'motor'noise processes associated to, respectively, the artificial and the human partner, and A_i, A_{ij}, B_i, with $i, j \in \{A, H\}$, $i \neq j$, are defined as follows:

$$A_{ii} = \begin{bmatrix} 0_2 & I_2 & 0_2 & 0_2 \\ -k \cdot M_i^{-1} & -b_i \cdot M_i^{-1} & M_i^{-1} & 0_2 \\ 0_2 & 0_2 & 0_2 & I_2 \\ 0_2 & 0_2 & -I_2/\tau_i^2 & -2I_2/\tau_i \end{bmatrix} \tag{5}$$

$$A_{ij} = \begin{bmatrix} 0_1 & 0_2 & 0_2 & 0_2 \\ k \cdot M_i^{-1} & 0_2 & 0_2 & 0_2 \\ 0_1 & 0_2 & 0_2 & 0_2 \\ 0_1 & 0_2 & 0_2 & 0_2 \end{bmatrix} \quad B_i = \begin{bmatrix} 0_2 \\ 0_2 \\ 0_2 \\ I_2/\tau_i^2 \end{bmatrix} \tag{6}$$

where I_2 and 0_2 denote, respectively, the 2×2 identity and zero matrices.

2) Artificial Partner Dynamics and Sensory System: Eq. 4 captures the dynamics of both the artificial and the human partner and their physical coupling. This general formulation will be used to simulate the artificial partner dynamics, which can be expressed as:

$$\dot{x}_A = A_{AA} \cdot x_A + A_{AH_1} \cdot p_H + B_A \cdot [u_A + w_A] \tag{7}$$

where A_{AH_1} denotes the first column of A_{AH}. For simulation and control purposes, the above equations were translated in discrete-time form:

$$\begin{aligned} x_A(t+1) = A^d_{AA} \cdot x_A(t) + A^d_{AH_1} \cdot p_H + B^d_A \\ \cdot [u_A(t) + w_A(t)] \end{aligned} \tag{8}$$

where A^d_{AA}, $A^d_{AH_1}$ and B^d_A are calculated from their continuous counterpart. To simplify the notation, in the following we will drop the 'd' but we will always refer to the discretized equations. The artificial partner's sensory system provides information about itself, its human counterpart (i.e., the whole dyad) and the external environment. Information about the human partner can be gathered both directly, e.g. through vision, thus observing p_H and $\dot{p}_H$, and through the interaction force f_{HA}. Hence the vector y_A of the sensory signals takes the form $y_A = [p_A, \dot{p}_A, p_H, \dot{p}_H, f_{HA}]^T$ and the sensory system is described as:

$$y_A = H_A \cdot x + v_A \tag{9}$$

The H_A matrix reflects the portion of sensory information which depends on system state:

$$H_A = \begin{bmatrix} I_2 & 0_2 & 0_2 & 0_2 & 0_2 & 0_2 \\ 0_2 & I_2 & 0_2 & 0_2 & 0_2 & 0_2 \\ 0_2 & 0_2 & 0_2 & I_2 & 0_2 & 0_2 \\ 0_2 & 0_2 & 0_2 & 0_2 & I_2 & 0_2 \\ -kI_2 & 0_2 & 0_2 & kI_2 & 0_2 & 0_2 \end{bmatrix} \tag{10}$$

The reliability of the sensory information is determined by the magnitude of the sensory noise, assumed to be Gaussian, $v_A(t) \sim N(0, \Sigma^v_A)$, with $\Sigma^v_A = \text{diag}(\sigma^2_x \cdot I_2, \sigma^2_{xd} \cdot I_2, \sigma^2_x \cdot I_2, \sigma^2_{xd} \cdot I_2, \sigma^2_f \cdot I_2)$. We set $\sigma_x = 1.7$ mm, $\sigma_{xd} = 35$ mm/s and $\sigma_f = 2$ N.

3) State Observer and Partner Model: Consistent with a large body of literature [12], we assume that the artificial partner optimally integrates sensory information its and own motor command (efferent copy) in order to estimate its own state, the state of its human partner, and possibly the state of the environment (for instance, target positions). Estimating human partner state is crucial in order to establish a successful interaction. While many studies agree that humans can form models of their opponents and/or they 'understand'their intentions, the exact nature of these models remains elusive. Interacting agents may simply estimate the ongoing partner movements. Or, they may additionally account for the partner's body dynamics, e.g. inertial properties, thus estimating the past motor commands [11]. Further, interacting agents may be able to infer their partner's control policy, i.e. the mapping between the state of both players and the partner's motor command [18]. This can be seen as a representation of the partner's ultimate goal or intentions [31]. Whatever its form, estimating partner state requires a forward model of the human partner's dynamics and the availability of suitable sensory information – e.g. vision, proprioception, audition, etc.

In the current AP implementation, the partner model fully accounts for the human partner's dynamic behavior. In particular, we assume that the state observer maintains a model of the whole dyad dynamics – which includes both artificial and human partner and their mechanical coupling. The state observer estimates the AP state x_A and the HP state, x_H. However, the latter is determined by the HP's motor command, u_H. Hence the state observer must be extended to predict the partner's motor command. Motor commands have sensory consequences, therefore the sensory consequences of movement carry information about the 'past' motor command. The complete set of 'prior'assumptions is summarized as:

$$\begin{aligned} x_A(t+1) &= A_{AA}x_A(t) + A_{AH}x_H(t) + B_A\,[u_A(t) + w_A(t)] \\ x_H(t+1) &= A_{HA}x_A(t) + A_{HH}x_H(t) + B_H\,[u_H(t) + w_H(t)] \\ u_H(t+1) &= u_H(t) + w_u(t) \end{aligned} \tag{11}$$

where $w_u(t) \sim N(0, \Sigma^u_x)$. The last equation reflects the prior belief that the human partner's input is an integrated Gaussian noise. In all experiments we set $\Sigma^u_x = 0.5\text{N}^2$. The state observer is defined in terms of the augmented state vector $X_A = \left[x^T_A\ x^T_H\ u^T_H\right]^T$ as:

$$\begin{aligned} \hat{X}_A(t+1) = A_A \cdot \hat{X}_A(t)\ + B_A \cdot u_A(t) \\ + K_A(t) \cdot \left\{ y_A(t) - [H_A\ 0]\,\hat{X}_A(t) \right\} \end{aligned} \tag{12}$$

where A_A is defined as:

$$A_A = \begin{bmatrix} A_{AA} & A_{AH} & 0 \\ 0 & A_H & B_H \\ 0 & 0 & A_u \end{bmatrix} \tag{13}$$

The Kalman gain $K_A(t)$, $t = 1, \cdots, T$ of the innovation term reflects the trade-off of the reliability of the prediction and correction terms.

4) Task Representation, Optimal Controller and Fictive Play: The AP task is specified by a quadratic cost functional:

$$\begin{aligned} J_A[u_A, u_H] = x(T)^T \cdot Q_T \cdot x(T) \\ + \frac{1}{T}\sum_{t=0}^{T-1} x(t)^T \cdot Q(t) \cdot x(t) \\ + \frac{1}{T}\sum_{t=0}^{T-1} u_A(t)^T \cdot R(t) \cdot u_A(t) \end{aligned} \tag{14}$$

The above cost functional completely specifies AP's control policy if the human motor command is known.

Consistent with the fictitious play notion, we assumed that the HP's motor commands (dynamic partner model) or the movements (kinematic partner model) do not change from the previous trial. In other words, at each trial the AP 'sees'a plant that incorporates the partner's state estimated at the previous trial. This results in an affine linear dynamical system, and the resulting LQG controller has a feedback and a feedforward component:

$$u_A(t) = -L_A^{tr}(t) \cdot x(t) - l_A^{tr}(t) \tag{15}$$

In conclusion, the artificial partner autonomously determines its own control policy with no explicit agreement with the human counterpart, which in game theory is called a non-cooperative scenario. This implementation of fictitious play only uses the most recent estimate of partner's input. This is less robust than estimating the distribution of partner inputs over multiple repetitions, but may be adequate for practical purposes.

In this formulation of the controller movement duration is specified in advance. This constraint may be relaxed by including temporal discount terms in an infinite-horizon cost functional.

5) Temporal Synchronization: Temporal synchronization is one basic feature of any joint action [32] – first and foremost, partners synchronize the start of their movements. Entrainment - i.e. the tendency to fall on the same rhythm - has been observed in unintentional action or when people are not required to synchronize. However, in intentional joint action, entrainment cannot fully explain temporal synchronization as the latter requires mutual adaptation [33]. To provide the artificial partner with the capability of synchronizing their movements with its human counterpart, we defined a simple mechanism which adapts the AP's start time, on a trial by trial basis, to the human start time:

$$TS^A(\mathrm{tr}+1) = TS^A(\mathrm{tr}) + \alpha \cdot \left[TS^H(\mathrm{tr}) - TS^A(\mathrm{tr})\right] \tag{16}$$

where tr is the trial number, and TS^A and TS^H are the start times (measured from target onset) of, respectively, artificial and human partner. The artificial partner uses a simple threshold mechanism (speed grater than 0.1 m/s) to estimate $TS^A(\mathrm{tr})$ at each trial. Parameter α specifies the trade-off between stability and plasticity. In all experiments we set $\alpha = 1$, i.e. the AP sets its start time to that observed in the human partner on the previous trial.

6) Implementation: The AP architecture has been implemented with two distinct haptic interfaces: a planar serial manipulandum with two degrees of freedom (braccio di ferro, [34]) and a planar cartesian manipulandum (H-MAN, Articares Ltd) [35]. In both cases, the implementation consists of two interconnected Simulink Desktop Real-Time and MATLAB applications. The Simulink application carries out real-time control of the artificial partner; the MATLAB application updates the optimal controller at the end of each trial on the basis of the estimated partner model. Within this implementation, the update rate of the artificial partner was 1 kHz. In the case of the H-MAN robot, the AP position was updated at 200 Hz. The refresh rate of the graphic display was 40 Hz and the data (position of both human and artificial partner and interaction forces) were saved at 100 Hz for further analysis.

B. Experiments

Although very simple, the proposed AP architecture is suitable for use in all finite-horizon scenarios that involve

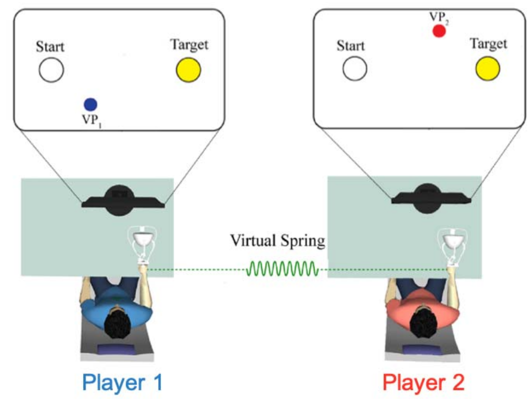

Fig. 2. Two-via point joint action task. Each player has the same starting and end point but different via-points (VP_1, VP_2). Each player seat in front of a computer screen and has to grasp the handle of a 3D haptic interface (Novint Falcon). They are mechanically coupled by a virtual spring and instructed to perform planar point-to-point movements through different via-points. Each player can only see the final target and his/her own via-point. The players can not see each other or communicate verbally. Modified from [11].

planar movements in the horizontal plane. These include situations investigated in several experimental studies of joint action and many robot-assisted rehabilitation protocols. In this preliminary study, we focused on two aspects of the AP: (i) its ability to unveil the ongoing movements of the human partner, and (ii) its ability to adapt its behaviors in order to optimize the interaction. To test these functionalities, we focused on a joint action scenario in which two partners – mechanically connected through a virtual spring – are required to perform reaching movements from the same starting point to a common target, by crossing an intermediate via-point (VP) which is different for each subject. The players can only see their own VP but cannot see or speak to each other and they are instructed to keep the interaction force to a minimum, so that they must somehow coordinate in order to accomplish their respective goals. The task can be interpreted as a non-cooperative coordination game. A recent study [11] reported that human players gradually develop a form of coordination which tends to the theoretical Nash equilibrium for that game – corresponding to both partners following the same trajectory through both VPs. The experimental setup and task are summarised in Figure 2.

In all experiments, the human participant sits in front of a computer screen. The subject grasps the robot handle and so that, when the hand is located approximately at the center of the workspace (origin of the reference system), the shoulder joint is flexed at about 45° with respect to the left-right shoulder direction and the forearm is flexed at about 90° with respect to the elbow. As in the human-human interaction experiments [11], the subject can see the start and target point, its own via-point and their own hand position (a cursor on the screen), and is instructed to grasp the robot's handle with the right hand and to control the cursor motion, moving it from the start point to the target by crossing the via-point and keeping the interaction force low. Participants were not informed that he/she would have performed a joint task with a partner – either human or artificial. The AP has identical requirements, which are completely specified by the following cost function:

$$\begin{aligned} J_A[u_A, u_H] = {} & w_p \cdot \|p_T - p_A(T)\|^2 \\ & + w_v \cdot \|\dot{p}(T)\|^2 \\ & + w_{vp} \cdot \|p_{VP_A} - p(TC_A)\|^2 \\ & + w_f \cdot \frac{1}{T} \sum_{t=1}^{T} \|p_H(t) - p_A(t)\|^2 \\ & + r \cdot w_u \cdot \frac{1}{T} \sum_{t=1}^{T} u_A(t)^2 \end{aligned} \tag{17}$$

The cost function has five terms. The first two terms enforce stopping on target at the end of the movement. The third term reflects the requirement to pass through the via-point. The fourth term accounts for minimising the distance between agents throughout the movement. The last term penalises the effort incurred during the movement, i.e. it encourages the AP to keep the interaction force low. The weight coefficients determine the relative importance of the corresponding constraint. We set these weights by assuming (Bryson's rule) a maximum acceptable displacement (in the via-point and in the final target) equal to, respectively, the radius of the via-point ($r_T = 2.5$ mm) and that of the target ($r_{VP} = 5$ mm). Similarly, we calculated the value of the 'velocity' weight by assuming a maximum acceptable speed at the target of $d_v = 5$ mm/s. We made a similar normalisation in the maximum inter-agent distance ($r_{AH} = 15$ mm) and maximum activation ($u_{max} = 10$ N). In all experiments we used the following weights: $w_p = 1/r_T^2$, $w_{vp} = 1/r_{VP}^2$, $w_v = 1/d_v^2$, $w_f = 1/r_{AH}$ and $w_u = 1/u_{max}^2$. The scalar coefficient r - which specifies the trade-off between task-related accuracy and effort - was set to 1. Greater r implies a greater sensitivity to effort and therefore lower controller gain. The time of crossing of the via-point is also part of the optimization. Based on some prior evidence that VP crossing times are approximately proportional to the fraction of path length covered at the time of crossing (PL), at each trial we set the crossing time as $TC^{tr+1} = PL^{tr} \cdot MT_A^{tr}$, where MT_A^{tr} is the total movement duration.

1) Experiment 1 – Partner Model: We expect that AP's ability to correctly estimate HP movements is affected by the strength of the physical connection and by the reliability of the haptic channel. We systematically varied the stiffness of the virtual spring k and the variance σ_f^2 of the haptic sensory channel. During this experiment, the human subject could see both via-points and was instructed to perform reaching movements from the start to the end point, by crossing both VPs. The robot force was switched off, i.e. $f_{AH} = 0$. In this way, the human subject was free to move with no AP intervention, but physical connection with the partner was always present on the AP side (so that $f_{HA} \neq 0$). The AP was programmed to make reaching movements through their own VP according to the cost functional of Eq.17. The simulated AP sensory system provided information about own movements, own VP, the target and the interaction force with the human subject, f_{HA}, but no direct information about position and velocity of the human subject. In summary, the human subject performs unperturbed reaching movements through two VPs but does not participate in the interaction. In contrast, the AP aims at adapting its movements in order to keep its goals (crossing its own VP) and to minimize the interaction force. As HP performance is almost stationary, accomplishing the task by the AP solely depends on its ability to predict own and partner actions. The experiment was organized in 30 epochs of 10 trials each. Each epoch corresponded to a different combination of stiffness and noise variance, in random balanced order. Specifically, we used six stiffness values: 100, 150, 200, 300, 400, 500 N/m and five noise standard deviation values: 0.1, 1, 2, 3, or 4 N ($6 \times 5 = 30$). Only one human subject (28 y, female) participated in this experiment as the focus here was on AP performance.

2) Experiment 2 – Learning Coordination: The goal of this experiment was to test various aspects of AP performance, in particular its capability to develop coordinated movements with a HP and whether they resemble those observed in two human partners. The experimental protocol was identical to [11]. Mechanical connection between AP and HP was now bi-directional, so that their movements were inter-dependent. Both AP and HP could only see their own via-point, and the only information about their partner was provided by the interaction force. The experiment was organized in 13 epochs of 12 trials each (a total of 156 trials). In the first epoch (baseline phase), mechanical coupling was turned off, and AP and HP acted alone. During epochs 2-11 the partners were mechanically connected. During epochs 12-13 (after-effect phase) the force was removed again. At the end of each trial, the human subject was provided a 0-100 score reflecting their performance (a combination of distance to own via-point and average interaction force - the lower the better). The experimental session lasted approximately one hour. The only difference with respect to the original experimental protocol was an additional control on movement time. This was necessary because of the finite-horizon implementation of the current AP ($MT = 2.5\ s$ in the current experiment). If the human participant's movement time was $2s < MT < 3s$, at the end of the trial the target color turned green (appropriate duration); otherwise it turned red (wrong duration). A total of four subjects (25-27 y, 3M+1F) participated in this experiment. Their performance was compared with the data from five human-human dyads [11], haptic (H) condition (25 ± 5 y, 9 M + 1 F).

The research conforms to the ethical standards laid down in the 1964 Declaration of Helsinki and was approved by the competent ethical committee (Comitato Etico Regione Liguria). Each participant signed a consent form conforming to these guidelines.

C. Data Analysis

The recorded AP and HP movements and forces were smoothed and differentiated (4th order Savitzky-Golay, window size 370 ms).

1) Experiment 1: AP and human player are connected through a virtual spring. As a measure of interaction, for each

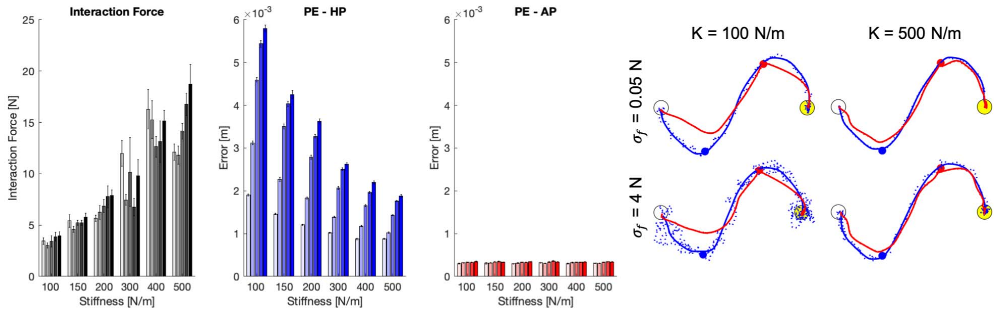

Fig. 3. Experiment 1 - Partner model. Left: The interaction force increases with stiffness. Each bar corresponds to an individual epoch and therefore to a specific stiffness-noise pair. The bars were reordered for increasing stiffness and noise values. Noise magnitude is denoted by grey level (light grey: $\sigma_f = 0.5$ N; dark grey: $\sigma_f = 4$ N). Middle: Trajectory estimation accuracy (PE), for HP (red) and AP (blue) . Data are reordered as for interaction force. All error bars denote standard errors. Right: AP (red) and HP (blue) trajectories in selected subject and trials, with different stiffnesses (K) and different noise standard deviation value σ_f. The continuous and dashed lines denote, respectively, the actual trajectories and the corresponding AP predictions.

trial (tr) we calculated the mean interaction force:

$$IF(tr) = \frac{1}{T(tr)} \sum_{t=0}^{T(tr)-1} |f_{HA}(t)| \tag{18}$$

where $T(tr)$ is the number of samples at trial tr. For each combination of stiffness and haptic noise magnitude, we took the IF average, standard deviation and coefficient of variation.

For each trial tr, we also computed the trajectory error PE between the actual and predicted trajectory of both the human and the artificial partner:

$$PE_i(tr) = \sqrt{\frac{1}{T(tr)} \sum_{t=0}^{T(tr)-1} \left|p_i(t) - \hat{p}_i(t)\right|^2} \tag{19}$$

where $p_i(t)$ is the position of player i, with $i \in \{A, H\}$, and $\hat{p}_i(t)$ is the position predicted by the AP's state observer. Likewise, we calculated the speed error SE between the actual and predicted speed profile of both the human and the artificial partner:

$$SE_i(tr) = \sqrt{\frac{1}{T(tr)} \sum_{t=0}^{T(tr)-1} \left|v_i(t) - \hat{v}_i(t)\right|^2} \tag{20}$$

where $v_i(t) = |\dot{p}_i(t)|$ is the speed profile. To examine the dependence of the above indicators on stiffness and noise level of the haptic channel, we used a repeated-measures ANOVA with two factors (stiffness and haptic noise magnitude). Due to the stochastic nature of AP behavior, we took each individual repetition of each condition ($N = 10$ trials per epoch) as an independent sample of AP behavior. We also corrected (Bonferroni-Holm) for multiple comparisons.

2) Experiment 2: We first assessed whether AP's temporal synchronization mechanism induced a mutual adaptation of the starting times of the two players, similar to that observed in human dyads. To do this, we compared the start times over trials of both artificial (TS^A) and human (TS^H) partners with those observed in previously reported human dyad experiments [11]. To analyze the overall dyad performance and convergence to joint coordination, we also evaluated the time course of spatial variability and of the minimum via-point distances.

Spatial variability was calculated as in [36]. We re-sampled all movements for a given subject at a fixed number of points equally spaced along the path, and found the average trajectory. Then, for each point along the average trajectory, we found the nearest sample point from each individual trajectory. These nearest points were averaged to recompute the corresponding point along the average trajectory, and the procedure was repeated until convergence. Hence this is a measure of path spatial variability, independent of time fluctuations.

The minimum via-point distances, MD_{ij}, were defined as the minimum distances at which player i gets closest to the j-th via-point.

To examine the dependence of the above indicators on training, we used paired-sample t-test comparing the indicators averaged over the initial and final epoch of the connected phase - epoch 2 and 11 respectively. As in Experiment 1, we corrected (Bonferroni-Holm) for multiple comparisons.

IV. Results

A. Experiment 1 - Partner Model

We first assessed the AP ability to predict its own movements and those of the human partner under various conditions (stiffness and noise level of the haptic channel). Figure 3 summarizes the partner model performance over conditions. As expected, the interaction force increases with stiffness; see Figure 3 (left). The observation was confirmed by statistical analysis. We observed a significant effect of stiffness ($p < 0.0001$), but no effect of noise or stiffness $\times$ noise interaction. The AP's state observer reliably predicts its own movements. As expected, prediction of the trajectory of the HP is significantly less reliable (greater PE: 0.25 ± 0.13 cm for HP vs 0.03 ± 0.0018 cm for AP); see Figure 3 (center). Stiffness and noise level affect prediction accuracy. As regards HP trajectories, the prediction error decreases significantly as stiffness increases ($p < 0.0001$) and noise decreases ($p < 0.0001$). We also found a significant stiffness $\times$ noise interaction ($p < 0.0001$). The AP prediction error increases as the sensory noise increases ($p < 0.0001$), but is not affected by stiffness. Similar results were observed for speed. The prediction accuracy is less reliable for HP (speed prediction error, SE, is about 5.76 ± 2.25 cm/s for HP vs 0.72 ± 0.22 cm/s for

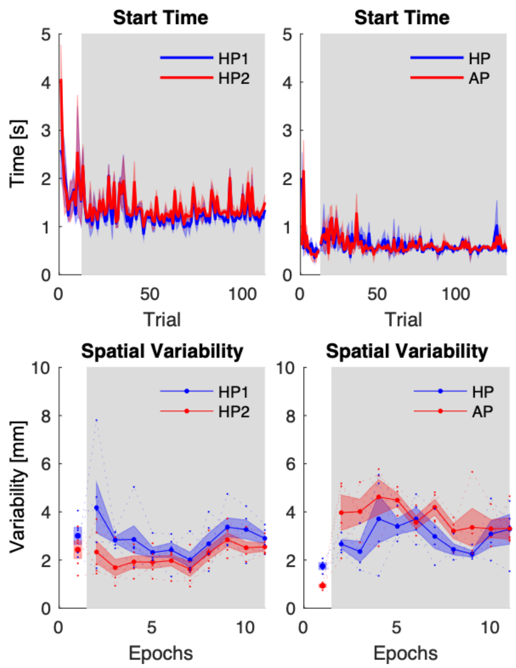

Fig. 4. Top: Artificial and human partner gradually learn to synchronize, i.e. to start their movements at the same time (left) in a way that is qualitatively similar to H-H dyads (right). Bottom: temporal evolution of spatial variability over epochs. The artificial partner (right) is in red. H-H data taken from [11]. All panels report mean $\pm$ SE.

AP). The speed prediction error (SE) for HP decreases as the stiffness increases ($p < 0.0001$) and the noise level decreases ($p < 0.0001$). The error also increases with both stiffness ($p < 0.0001$) and noise level ($p < 0.0001$) for AP's speed estimation. We also observed a significant stiffness $\times$ noise interaction for both HP ($p < 0.0001$) and AP ($p = 0.0014$).

B. Experiment 2 - Learning Coordination

The ability to synchronize the timing of the movements is crucial for the development of collaborative strategies [3]. Figure 4 (top) compares the effect of AP-HP (A-H) synchronization with corresponding observations from human-human dyads. Synchronization is qualitatively similar in A-H and H-H dyads. A-H and H-H also exhibit similar amounts of inter-trial variability; see Figure 4 (bottom).

We then looked at the trajectories and speed profiles that the A-H dyad exhibits over repeated trials. During the baseline phase, AP and HP are not mechanically coupled and they simply perform reaching movements through their respective via-points; see Figure 5 a,c (top). The A-H and H-H dyads exhibit movements that are qualitatively similar in path, speed profile and inter-trial variability. During the adaptation phase, when mechanical coupling was turned on, both partners in a dyad gradually tend to move along similar paths by crossing their VP and getting closer to their partner's; see Figure 5a,c (bottom). Again, the behavior of A-H dyads is qualitatively similar to their H-H counterparts.

In H-H dyads, both subjects decrease their distance from their opponent's via-point and maintain the distance from own via-point low along all epochs; see Figure 5b. This indicates that both players prioritize their own goals and gradually improve coordination with their partners; see [11]. This behavior is a signature of the development of a coordination.

The same behavior is observed in A-H dyads. Over trials, both the AP and the HP gradually decrease their distance from their opponent's via-point, while at the same time keeping the distance from their own via-point low; see Figure 5d. During interaction the HPs behave as in H-H dyads, by significantly decreasing their minimum distance to their partners ($p < 0.0001$); in APs the improvement is only partial (no significant decrease). At population level, we also found no significant change in start time and spatial variability.

V. Discussion

We presented a general computational framework for the development of artificial partners which are capable of establishing a coordination with a human partner. The architecture involves three main components: (i) a simulated body and sensory system, (ii) a state and partner observer, which includes an internal representation of the partner (partner model); and (iii) a feedback controller, based on a representation of the task in terms of a quadratic cost function. This formulation is rooted on probabilistic (bayesian) sensorimotor integration [12] and optimal control [13]. We extend these principles – summarized in section II.A – in two directions. First, we argue that during joint action each participant infers their partner's intentions and/or ongoing actions; see II.B [3], [4]. Building upon an experimentally confirmed model [11], we formulate a sensorimotor integration model which predicts not only the dyad state, but also the ongoing partner actions. Second, consistent with the optimal control framework, we propose that coordination strategies emerge gradually through a simple adaptive process (fictitious play) which, in the case of perfect information, leads to a Nash equilibrium [30]. We present a specific implementation, which focuses on planar hand movements with finite duration , where the human and artificial partner are mechanically connected through a haptic interface. Under these assumptions, dyad dynamics and their sensory systems are described by linear time invariant dynamical models, with Gaussian noise. The artificial partner's task is described by a quadratic cost functional. Although simple, this formulation is applicable to a variety of experimental joint action scenarios, and its predictions can be directly compared to the observed human-human interaction outcomes.

A. The Artificial Partner Estimates the Actions of Its Opponent

We tested AP's ability to estimate the actions of a human opponent. As expected, we found that estimation improves as the strength of the coupling increases and as the noise variance decreases; see Figure 3. In the current implementation the AP keeps a dynamic model of the human, which fully accounts for its dynamic behavior. This is motivated by our previous report that players simply estimate the ongoing partner actions [11]. Other studies suggested that players develop more general partner representations, also accounting for intentions and ultimate goals [31]. Humans are indeed very good at extrapolating higher order information by observing the motion of

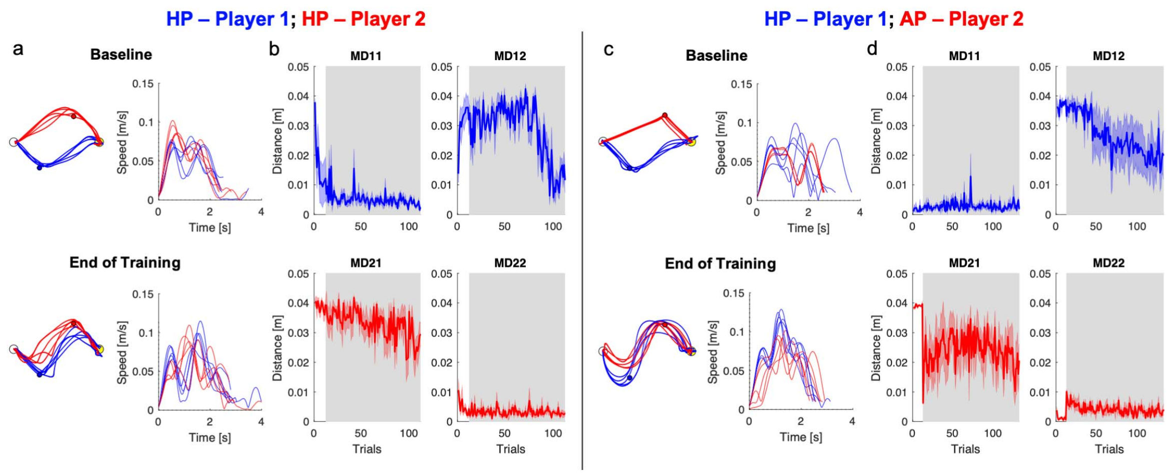

Fig. 5. Trajectories and speed profiles during baseline and the last training epoch, for typical H-H (a) and A-H (c) dyads. In the A-H dyad, AP is depicted in red, HP in blue. Both H-H and A-H dyads gradually develop a coordination. Temporal evolution of the minimum distance from via-point in human-human (b) and human-artificial (d) partner interaction (respectively, in blue and in red).

their peers [37]. Alternative partner model formulations – for instance, the AP could use a 'kinematic' representation which only accounts for some specific features of human movements, like smoothness, but does not explicitly account for the underlying mechanisms. APs could serve as general tools to investigate what we represent about a partner and how we develop that representation.

B. Artificial and Human Partner Learn to Coordinate

Consistent with previous work [11], [20], we found that mechanical coupling can be exploited by interacting partners to exchange information and achieve a certain degree of coordination. Players in both A-H and H-H dyads tend to synchronise their start times and exhibit similar amounts of inter-trial variability; see Figure 4. In H-H dyads, over trials the participants gather information about dyad dynamics and about the partner and adjust their actions accordingly. In the 2-via point task, this is reflected in the gradual decrease of the trajectory distances from both via-points; see Figure 5. Although very simple, fictitious play reproduces this behavior in A-H dyads. Human participants interacting with the AP develop a coordination which is very similar to that observed in human-human dyads. However, APs converge less often to the trajectories corresponding to Nash equilibria. In fact, the observed trajectories suggest that the AP does not fully compensate for the HP. A systematic exploration of this issue – for instance, by systematically varying stiffness and sensory noise levels – is beyond the scope of this study which only aims at a technical validation of the basic AP architecture.

C. A General Platform to Study Joint Action

The proposed experiments focus on a specific sensorimotor interactive task, but the architecture and underlying assumptions – see Figure 1 – can be extended to investigate different tasks and scenarios. AP behavior is completely specified by a set of parameters, related to body dynamics (A_{AA}, A_{AH} and B_A), sensory system (H_A), perceptual and motor uncertainty (Σ_A^v, Σ_A^w), assumptions on partner model (dynamic, kinematic, or other), and personal traits like vigor (r) and willingness to establish a coordination. Other personal traits like risk sensitivity [38] can be easily incorporated. Parameters selection is crucial for determining the AP attitude toward the interaction. For instance, manipulating the weights of the 'effort' component of the AP cost function would lead the AP to prefer more or less effortful movements (and therefore smoother trajectories). Further, manipulating the start time adaptation rate (parameter α) would control AP's ability to synchronize with the human partner. The use of realistic APs with different attitudes and personal traits toward the interaction simplifies the study of joint action. In the A-H experimental paradigm one mind is a completely specified, thus allowing to characterize how the human counterpart responds.

D. Artificial Partners as Diagnostic and Rehabilitation Tools

Biomimetic partners with a realistic, predictable behavior may be used to characterize the ability of a human partner to establish an interaction [5].The inherently human capability to build representations about the partner and integrate them with own internal representation is altered in some pathologies, e.g. autism spectrum disorders or schizophrenia [39], and a better understanding of these alterations would help improving the diagnosis and suggest ways to contrast their consequences. However, APs may be useful not only to investigate joint action. They can also be used to facilitate skill learning and neuromotor recovery. Rehabilitation robots have been often described as 'artificial therapists' and many of the developed technological solutions are somehow inspired by the observed mechanics of therapist-patient interaction. Patient and therapist constantly exchange information during a session. Looking at their actions and at the response to micro-perturbations, the therapist gradually develops an understanding of patient impairment and recovery potential. In the case of rehabilitation, interaction should aim at maximising recovery, making it faster and more durable. When interacting with assistive robots, humans tend to incorporate assistive forces in their motor plan, thus reducing their active contribution to movement, which may have adverse effects on recovery [40]. To counteract this, several heuristic mechanisms have been proposed to provide 'assistance-as-needed' [41]. A robot with an inherent ability to develop optimal forms of interaction would provide assistance 'as needed' by design. Also, it would automatically adapt to patient recovery.

VI. Conclusion

The proposed AP architecture is intended as a general modeling framework. The current formulation makes a number of specific assumptions – on representation of the partner's actions , on the way coordination is achieved – which clearly require further empirical examination. To this purpose, the model may constitute a valuable experimental tool to test different hypotheses. The presented model formulation is limited to planar arm movements and linear dynamics, but can be extended to more complex, higher dimensional non-linear scenarios, for which mathematical tools and numerically tractable implementations are now available [42].

References

[1] N. Sebanz, H. Bekkering, and G. Knoblich, "Joint action: Bodies and minds moving together," *Trends Cognit. Sci.*, vol. 10, no. 2, pp. 70–76, Feb. 2006.

[2] G. Dumas and M. T. Fairhurst, "Reciprocity and alignment: Quantifying coupling in dynamic interactions," *Roy. Soc. Open Sci.*, vol. 8, no. 5, May 2021, Art. no. 210138.

[3] C. Vesper *et al.*, "Joint action: Mental representations, shared information and general mechanisms for coordinating with others," *Frontiers Psychol.*, vol. 7, p. 2039, Jan. 2017.

[4] A. Curioni, G. Knoblich, N. Sebanz, A. Goswami, and P. Vadakkepat, "Joint action in humans: A model for human–robot interactions," in *Humanoid Robotics: A Reference*. Dordrecht, The Netherlands: Springer 2019, pp. 2149–2167.

[5] G. Dumas, G. C. de Guzman, E. Tognoli, and J. A. S. Kelso, "The human dynamic clamp as a paradigm for social interaction," *Proc. Nat. Acad. Sci. USA*, vol. 111, no. 35, pp. E3726–E3734, Sep. 2014.

[6] J. A. S. Kelso, G. C. de Guzman, C. Reveley, and E. Tognoli, "Virtual partner interaction (VPI): Exploring novel behaviors via coordination dynamics," *PLoS ONE*, vol. 4, no. 6, p. e5749, Jun. 2009.

[7] A. A. Sharp, M. B. O'Neil, L. F. Abbott, and E. Marder, "Dynamic clamp: Computer-generated conductances in real neurons," *J. Neurophysiol.*, vol. 69, pp. 992–995, Mar. 1993.

[8] H. Haken, J. A. S. Kelso, and H. Bunz, "A theoretical model of phase transitions in human hand movements," *Biol. Cybern.*, vol. 51, no. 5, pp. 347–356, Feb. 1985.

[9] G. Dumas, Q. Moreau, E. Tognoli, and J. A. S. Kelso, "The human dynamic clamp reveals the fronto-parietal network linking real-time social coordination and cognition," *Cerebral Cortex*, vol. 30, no. 5, pp. 3271–3285, May 2020.

[10] J. Grau-Moya, E. Hez, G. Pezzulo, and D. A. Braun, "The effect of model uncertainty on cooperation in sensorimotor interactions," *J. Roy. Soc. Interface*, vol. 10, no. 87, Oct. 2013, Art. no. 20130554.

[11] V. T. Chackochan and V. Sanguineti, "Incomplete information about the partner affects the development of collaborative strategies in joint action," *PLOS Comput. Biol.*, vol. 15, no. 12, Dec. 2019, Art. no. e1006385.

[12] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan, "An internal model for sensorimotor integration," *Science*, vol. 269, no. 5232, pp. 1880–1882, 1995.

[13] E. Todorov and M. I. Jordan, "Optimal feedback control as a theory of motor coordination," *Nature Neurosci.*, vol. 5, no. 11, pp. 1226–1235, 2002.

[14] D. M. Wolpert, K. Doya, and M. Kawato, "A unifying computational framework for motor control and social interaction," *Phil. Trans. Roy. Soc. London B, Biol. Sci.*, vol. 358, no. 1431, pp. 593–602, Mar. 2003.

[15] E. Todorov, "Optimality principles in sensorimotor control," *Nature Neurosci.*, vol. 7, no. 9, pp. 907–915, Sep. 2004.

[16] R. Shadmehr and J. W. Krakauer, "A computational neuroanatomy for motor control," *Exp. Brain Res.*, vol. 185, no. 3, pp. 359–381, 2008.

[17] D. A. Braun, P. A. Ortega, and D. M. Wolpert, "Nash equilibria in multi-agent motor interactions," *PLoS Comput. Biol.*, vol. 5, no. 8, Aug. 2009, Art. no. e1000468.

[18] Y. Li, G. Carboni, F. Gonzalez, D. Campolo, and E. Burdet, "Differential game theory for versatile physical human–robot interaction," *Nature Mach. Intell.*, vol. 1, no. 1, pp. 36–43, 2019.

[19] R. P. R. D. van der Wel, G. Knoblich, and N. Sebanz, "Let the force be with us: Dyads exploit haptic coupling for coordination," *J. Exp. Psychol., Hum. Perception Perform.*, vol. 37, no. 5, pp. 1420–1431, 2011.

[20] G. Ganesh, A. Takagi, R. Osu, T. Yoshioka, M. Kawato, and E. Burdet, "Two is better than one: Physical interactions improve motor performance in humans," *Sci. Rep.*, vol. 4, no. 1, pp. 1–7, May 2015.

[21] A. Takagi, F. Usai, G. Ganesh, V. Sanguineti, and E. Burdet, "Haptic communication between humans is tuned by the hard or soft mechanics of interaction," *PLOS Comput. Biol.*, vol. 14, no. 3, Mar. 2018, Art. no. e1005971.

[22] G. Knoblich and J. S. Jordan, "Action coordination in groups and individuals: Learning anticipatory control," *J. Exp. Psychol., Learn., Memory, Cognition*, vol. 29, no. 5, p. 1006, 2003.

[23] D. W. Franklin and D. M. Wolpert, "Computational mechanisms of sensorimotor control," *Neuron*, vol. 72, no. 3, pp. 425–442, 2011.

[24] M. Bacharach, "Interactive team reasoning: A contribution to the theory of co-operation," *Res. Econ.*, vol. 53, no. 2, pp. 117–147, Jun. 1999.

[25] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. Philadelphia, PA, USA: SIAM, 1998.

[26] N. Jarrassé, T. Charalambous, and E. Burdet, "A framework to describe, analyze and generate interactive motor behaviors," *PLoS ONE*, vol. 7, no. 11, Nov. 2012, Art. no. e49945.

[27] J. Nash, "Non-cooperative games," *Ann. Math.*, vol. 54, no. 2, pp. 286–295, 1951.

[28] U. Berger, "Brown's original fictitious play," *J. Econ. Theory*, vol. 135, no. 1, pp. 572–578, Jul. 2007.

[29] G. W. Brown, "Iterative solution of games by fictitious play," *Activity Anal. Prod. Allocation*, vol. 13, no. 1, pp. 374–376, 1951.

[30] D. Fudenberg and D. Levine, "Learning in games," *Eur. Econ. Rev.*, vol. 42, no. 3, pp. 631–639, May 1998.

[31] N. Sebanz, G. Knoblich, and W. Prinz, "How two share a task: Corepresenting stimulus-response mappings," *J. Exp. Psychol., Hum. Perception Perform.*, vol. 31, no. 6, p. 1234, 2005.

[32] C. Vesper, S. Butterfill, G. Knoblich, and N. Sebanz, "A minimal architecture for joint action," *Neural Netw.*, vol. 23, nos. 8–9, pp. 998–1003, Oct. 2010.

[33] N. Sebanz and G. Knoblich, "Prediction in joint action: What, when, and where," *Topics Cognit. Sci.*, vol. 1, no. 2, pp. 353–367, Apr. 2009.

[34] M. Casadio, V. Sanguineti, P. G. Morasso, and V. Arrichiello, "Braccio di Ferro: A new haptic workstation for neuromotor rehabilitation," *Technol. Health Care*, vol. 14, no. 3, pp. 123–142, Jan. 2006.

[35] D. Campolo, P. Tommasino, K. Gamage, J. Klein, C. M. L. Hughes, and L. Masia, "H-man: A planar, H-shape cabled differential robotic manipulandum for experiments on human motor control," *J. Neurosci. Methods*, vol. 235, pp. 285–297, Sep. 2014.

[36] D. Liu and E. Todorov, "Evidence for the flexible sensorimotor strategies predicted by optimal feedback control," *J. Neurosci.*, vol. 27, pp. 9354–9368, Aug. 2007.

[37] A. Zunino *et al.*, "Predicting intentions from motion: The subject-adversarial adaptation approach," *Int. J. Comput. Vis.*, vol. 128, no. 1, pp. 220–239, Jan. 2020.

[38] J. Grau-Moya, P. A. Ortega, and D. A. Braun, "Risk-sensitivity in Bayesian sensorimotor integration," *PLoS Comput. Biol.*, vol. 8, no. 9, Sep. 2012, Art. no. e1002698.

[39] J. Izawa, T. Asai, and H. Imamizu, "Computational motor control as a window to understanding schizophrenia," *Neurosci. Res.*, vol. 104, pp. 44–51, Mar. 2016.

[40] J. L. Emken, R. Benitez, A. Sideris, J. E. Bobrow, and D. J. Reinkensmeyer, "Motor adaptation as a greedy optimization of error and effort," *J. Neurophysiol.*, vol. 97, no. 6, pp. 3997–4006, 2007.

[41] L. Marchal-Crespo and D. J. Reinkensmeyer, "Review of control strategies for robotic movement training after neurologic injury," *J. Neuroeng. Rehabil.*, vol. 6, p. 20, Jun. 2009.

[42] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Oct. 2012, pp. 5026–5033.