Employing Models of Human Social Motor Behavior for Artificial Agent Trainers

Lillian M. Rigoli
Department of Psychology,
Macquarie University
Sydney, NSW, Australia

Patrick Nalepka
Department of Psychology,
Macquarie University
Sydney, NSW, Australia

Hannah Douglas
Department of Psychology,
Macquarie University
Sydney, NSW, Australia

Rachel W. Kallen
Department of Psychology,
Macquarie University
Sydney, NSW, Australia

Simon Hosking
Aerospace Division, Defence Science and Technology Group
Melbourne, Victoria, Australia

Christopher Best
Aerospace Division, Defence Science and Technology Group
Melbourne, Victoria, Australia

Elliot Saltzman
Department of Physical Therapy & Athletic Training, Boston University, MA & Haskins Laboratories, CT

Michael J. Richardson
Department of Psychology,
Macquarie University
Sydney, NSW, Australia

ABSTRACT

Many everyday tasks require individuals to work together as a team to achieve a task goal. For many complex or high-stakes multi-agent activities, team members are required to participate in simulated training exercises to develop the task- and team-work (coordination) skills needed to maximize task performance. Such training, however, can be both time- and labor-intensive, requiring the participation of full teams and expert instructors. One way to minimize these costs is to augment team training scenarios with interactive artificial agents (AAs) capable of robust, 'human-like' behavioral interaction. With regard to perceptual-motor tasks specifically, recent research suggests that this can be achieved using task dynamical models derived from the dynamical primitives of human motor behavior. To investigate the degree to which such models can be employed for human team training, we examined whether a task dynamical model of human herding behavior could be embedded in the control architecture of an AA to train novice human-actors to learn various simulated multi-agent herding tasks. Three experiments were conducted that (i) first modeled human team performance during a set of novel herding tasks adapted from [19, 21], (ii) tested an AA utilizing this model to complete the tasks with human novices, and (iii) demonstrated how this AA could effectively train novices in a manner comparable to a human-expert trainer.

KEYWORDS

Agents competing and cooperating with humans; Agents for improving human cooperative activities; Human-robot/agent interaction; social agent models

ACM Reference Format:
Lillian M. Rigoli, Patrick Nalepka, Hannah Douglas, Rachel W. Kallen, Simon Hosking, Christopher Best, Elliot Saltzman, and Michael J. Richardson. 2020. Employing Models of Human Social Motor Behavior for Artificial Agent Trainers. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 9 pages.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

1 INTRODUCTION

A wide range of activities require that individuals work together to achieve a shared task goal. Such multi-agent activity occurs when a team of doctors perform an emergency surgery or a military unit is engaged in a tactical action. In these and many other cases, a breakdown in coordination can have severe consequences. Thus, it is essential for team members to undergo training to ensure that they not only have the appropriate skills for effective task performance, but can perform these skills in a manner that facilities and ensures coordinated co-action. Indeed, key to effective team training is creating training scenarios that provide the opportunity for those individuals to learn both task- and team-work skills simultaneously [6, 31]. This requires that team training scenarios include intact teams and promote self-guided learning through the simulation of task contexts [7, 25, 26]. Combined with the fact that expert instructors are also required to complete team training exercises, the cost of conducting team training, both in terms of time and money, can be incredibly high.

One way to reduce the logistic and financial hurdles associated with team training is to incorporate artificial agents (AA) within team training contexts. Indeed, human-AA team training systems not only have the potential to reduce the costs of team training, but can also provide individuals with the opportunity to engage in more frequent and individualized skills training. Accordingly, there is now a growing body of research directed towards investigating the degree to which human team members and even human trainers could be replaced or enhanced by interactive AAs [16, 17, 34].

A central focus of this latter research is what kinds of AAs lead to better training performance outcomes (e.g., [36]). For example, certain tasks might be best enhanced by including a highly rigid AA [3], while other more stochastic tasks might be better enhanced using AAs that are more readily adaptive to changing task contexts

[1, 2]. In either case, the effectiveness of human-AA training often depends on the ability of AAs to respond to human co-actors in a seamless and effective manner. It has therefore been hypothesized that achieving optimal human-user experience and learning depends on AA technologies incorporating natural, human-like patterns of behavior, as more human-like AAs tend to result in a better transition from AA training to human co-action [35, 39].

1.1 Dynamical Motor Primitives

To develop AAs capable of human-like behavioral coordination requires first modelling the fundamental behavioral processes that define effective human behavior within a given task context. Of particular relevance here, is that despite the assumed complexity of human multi-agent activity, there is now a growing body of research demonstrating how the spatio-temporal patterning of many human perceptual-motor behaviors is typically low-dimensional and synergistic, and can be explained and modelled using a small, fundamental set of simple dynamical rules (e.g., [11, 23, 24, 28, 30, 37]), or *dynamical motor primitives* (DMPs; e.g., [8, 9, 29]). Specifically, this research has revealed that human actions are generally composed of two movement types: (i) discrete movements, such as reaching or throwing an object; and (ii) rhythmic movements, such as walking or hammering a nail. Moreover, these primitives can be modeled using fundamental behaviors of nonlinear dynamical systems, namely: point-attractive and limit cycle behaviors.

With regard to the specific formulation of these dynamical motor primitives (DMPs), several related approaches have been proposed (e.g., [9–11, 22, 27–29, 37]). In principle, they all capture discrete movements using a point-goal directed damped-mass spring system, and rhythmic movements using a forced (driven) damped-mass spring system or a nonlinear self-sustained oscillator (e.g., a Rayleigh or van der Pol oscillator). For instance, [28] have demonstrated how a wide range of reaching, rhythmic wiping, and cranking tasks can be modelled using simple task-specific systems of fixed-point and limit cycle attractors acting on corresponding end-effectors (e.g., hands for reaching). Similarly, [5, 37] have demonstrated how human goal-directed navigation within an obstacle-ridden environment can be modelled using a simple, elementary set of DMPs. Modelling by [29] and [33] has demonstrated similar possibilities with regard to drumming and racket ball bouncing.

With regard to multi-agent behavior, recent research has demonstrated how dynamical models composed of DMPs can effectively capture the end-effector trajectories of co-actors' hands when completing cooperative object sorting and passing tasks [13, 14], as well as the behavioral dynamics of effective interpersonal collision avoidance [22] and social rhythmic coordination tasks [4, 12, 38]. However, perhaps the best example of how complex multi-agent activity can be modeled using systems of fixed-point and limit cycle dynamical primitives stems from the recent work of [19, 21], in which the authors demonstrated how a relatively simple model composed of DMPs is able to capture the behavior of human-actors performing complex multi-agent herding tasks.

1.2 Multi-agent Herding Task

The multi-agent herding task [19, 21] is a game that requires two players to control the movements of *herding agents* (HAs) in order to corral and contain a herd of *target agents* (TAs) within a defined containment area. The task is typically played on a large table-top display, with players standing on opposite sides of the display and controlling the movements of the HAs using a motion tracking device or pen (see Figure 1); although the task has also been investigated in a fully immersive, 3D virtual environment [20]. When left undisturbed, the TAs roam randomly around the game field. When an HA comes within a specified distance, TAs will move away from the HA. The goal of the task is for players to corral and contain the TAs within a specified time (e.g, 120 s), with containment deemed successful if the TAs are contained within the containment area for a certain proportion of time (e.g, 70% of a trial).

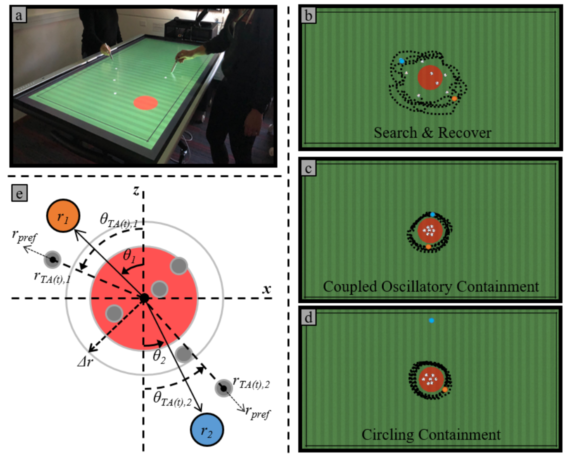

Figure 1: (a) The herding task and setup employed for human-human and human-artificial agent testing. (b-d) The behaviors expected to be observed in the current study: (b) Search and Recover (S&R); (c) Coupled Oscillatory Containment (COC); and (d) Circling Containment (CIR). The black trajectories represent approximately 10 s of behavior by the blue and orange herding agents (HAs). The white dots are the target agents (TAs) being herded. (e) Illustration of the task space employed for the [19, 21] herding model (see text for more details).

Initially, all human players adopt a behavioral mode termed *search and recover* (S&R), in which players discretely herd the TA furthest from the containment area on their side of the game field. While this strategy is useful for corralling TAs towards the containment area, it is difficult to keep a herd of TAs contained using S&R behavior alone (see Figure 1b). Accordingly, a subset of players spontaneously discover and adopt a less obvious, yet far superior mode of TA containment termed *coupled oscillatory containment* (COC). As illustrated in Figure 1c, COC corresponds to players encircling the TAs by moving in a rhythmic, semicircular manner around the herd. Discovering COC reflects a moment of sudden insight, with those players that use COC achieving nearly 100% containment success after its discovery. Moreover, the key realization for players that discover COC, is not the performance benefits of COC itself, but the understanding that oscillatory movements provide the most effective way of containing TAs together as a herd. Indeed, [19] revealed how players who realize the effectiveness of employing oscillatory behavior to contain a TA-herd, also appear to simultaneously realize that a single HA can contain a TA-herd by moving in a circle around the herd (see Figure 1e).

1.3 Task Dynamic Herding Model

Of particular relevance here, is that [19, 21] demonstrated how human S&R and COC reflect task-specific realizations of environmentally coupled point-attractor and limit cycle dynamics. Using a task-dynamic approach [27, 28, 37], the behaviour of each HA-i's (where $i = 1$ or 2) can be defined in polar task-space coordinates, where $(r_{pole}, \theta_{pole}) = (0, 0)$ is the (x, z) center of the containment area (see Figure 1e). Given this task space and the data-driven assumption that at each time step (t), the TA that HA-i selects to pursue (i.e, $TA(t), i$) is (i) furthest from the specified containment location and (ii) closer to HA-i's current field position. The dynamics of each HA-i's radial distance was defined as,

$$\ddot{r}_i + b_r \dot{r}_i + \varepsilon_r \left(r_i - \left(r_{TA(t),i} + r_{pref} \right) \right) = 0 \tag{1}$$

where $\dot{r}_i$, and $\ddot{r}_i$ represent the velocity and acceleration, respectively, of HA-i's radial distance with respect to the center of the containment area; $r_{TA(t),i}$ is the radial coordinate of the TA that HA-i is currently pursuing; r_{pref} is a fixed parameter that specifies HA-i's minimum preferred radial distance from TAs; b_r is a damping term; and ε_r scales the strength of the centrally-directed radial force attracting HA-i's to $TA(t), i$. Essentially, Eq. (1) exhibits fixed-point attractor dynamics and operates to minimize the difference between HA-i's current radial distance, r_i, and the radial distance, $r_{TA(t),i}$, of the TA currently being pursued.

With regard to the dynamics of each HA-i's radial angle (θ_i), this was defined using a hybrid non-linear oscillator [10] capable of exhibiting both point-attractor and limit cycle dynamics. More specifically, the angular dynamics (θ_i) of each HA-i was defined using the equation

$$\ddot{\theta}_i + b_{\theta_i} \dot{\theta}_i + \beta \dot{\theta}_i^3 + \gamma \theta^2{}_i \dot{\theta}_i + \varepsilon_\theta (\theta_i - \theta_{TA(t),i}) = \left(\dot{\theta}_i - \dot{\theta}_j \right) \left(A + B \left(\theta_i - \theta_j \right)^2 \right) \tag{2}$$

where $\dot{\theta}_i$ and $\ddot{\theta}_i$ represent the velocity and acceleration, respectively, of HA-i's radial angle with respect to the center of the containment area; b_{θ_i} and ε_θ are linear damping and stiffness terms; and $(\beta \dot{\theta}_i^3)$ and $(\gamma \theta^2{}_i \dot{\theta}_i)$ are Rayleigh van der Pol nonlinear escapement terms, respectively.

Ignoring the right side of (2) for the moment, (2) results in fixed-point like dynamics when $b_{\theta_i} \geq 0$ and limit cycle dynamics when $b_{\theta_i} < 0$. That is, when $b_{\theta_i} \geq 0$, (2) minimizes the difference between HA-i's current radial angle, θ_i and the radial angle, $\theta_{TA(t)}, i$ of $TA(t), i$. However, when $b_{\theta_i} < 0$, (2) results in HA-i oscillating back and forth about $\theta_{TA(t),i}$.

Key to the realization of S&R versus oscillatory (OC) and COC behavior at any time step (t) is the value of b_θ. Based on the assumption that individuals begin to adopt oscillatory behaviour when the furthest TA, $TA(t), i$, is sufficiently close to the containment area, [19, 21] defined the parameters dynamics of b_θ as

$$\dot{b}_{\theta_i} + \delta \left(b_{\theta_i} - \alpha \left(r_{TA(t),i} - r_\Delta \right) \right) = 0 \tag{3}$$

where $\dot{b}_{\theta_i}$ is the rate of change of b_{θ_i}, and γ and α are fixed parameters that determine the rate at which b_{θ_i} changes as a function of the radial distance of $TA(t), i$ relative to the radial distance, r_Δ, that an HA considers close enough to the containment location for $TA(t), i$ to be considered 'sufficiently' corralled. When considered in conjunction with (1) and (2), (3) results in S&R like behavior when $TA(t), i$ is far from the containment area (i.e., when $r_{TA(t),i} \geq r_\Delta$), and OC behavior when $TA(t), i$, and thus all TAs close to HA-i, are also very close to or within the containment area (i.e., when $r_{TA(t),i} < r_\Delta$).

Returning to the right side of (2), this function couples the angular movements of HA-i to partner HA-j, with the coupling strength parameters A and B determining whether HAs exhibit stable in-phase or anti-phase rhythmic coordination during COC behaviour. Note that this latter coupling function was included by [21] because human actors tend to produce stable in-phase and anti-phase COC behaviour. However, such behavioural synchrony is not required for effective herd containment[1].

Finally, although [19, 21] never attempted to model the kind of *Circling containment* (CIR) behavior illustrated in Figure 1d, this behaviour is akin to end-effector "cranking", which [28] have previously demonstrated can be modeled using a reduced version of (2). Specifically,

$$\ddot{\theta}_i + b_\theta \dot{\theta}_i + \beta \dot{\theta}_i^3 = 0 \tag{4}$$

where the frequency of, $\dot{\theta}_i$, is equal to $\pm \sqrt{\frac{b_\theta}{\beta}}$, and direction is dependent on the initial sign of $\dot{\theta}_i$. Hence, a Heaviside function can be employed to transition the model between potential OC and CIR behaviour (see subsection 3.1.1 for details).

1.4 Current Study

The current study aimed to demonstrate how task dynamical models of human behavior (composed of DMPs) can be employed to develop interactive AAs capable of training novice human-actors in a manner comparable to expert human trainers. To achieve this aim, three experiments were conducted to investigate whether an HA controlled by the above multi-agent herding model could be employed to train novice human-actors for different herding task scenarios. **Experiment 1 *(Novice Human-Human Performance)*** explored the behavioral dynamics exhibited by human-human pairs when learning to perform several new versions of the multi-agent herding task. **Experiment 2 *(Novice Human-Artificial Agent Performance)*** then explored the degree to which a task specific variant of the herding model proposed in [21] was able to effectively capture the behavior of human players across these new tasks by comparing human-(model)AA performance to human-human performance. Finally, **Experiment 3 *(Human-Artificial Agent to Human-Human Transfer)*** tested whether novice human-actors could not only learn the fundamental modes of behavior required for successful herding performance (i.e., OC, COC, circling [CIR] behavior) from a model controlled AA trainer, but were also able

[1] Parameter settings employed in the current study: $b_r = 45$; $\varepsilon_r = 45$; $r_{pref} = 0.03m$. $\beta = 0.161641$; $\gamma = 7.22282$; $\varepsilon_\theta = \omega^2$; $\omega = 6.364866716 \frac{rad}{s}$ during COC behavior. $\omega = 12.729733432 \frac{rad}{s}$ during S&R behavior; $A = -0.2$; $B = 0.2$; $\delta = 23.08993$; $\alpha = 138.262$

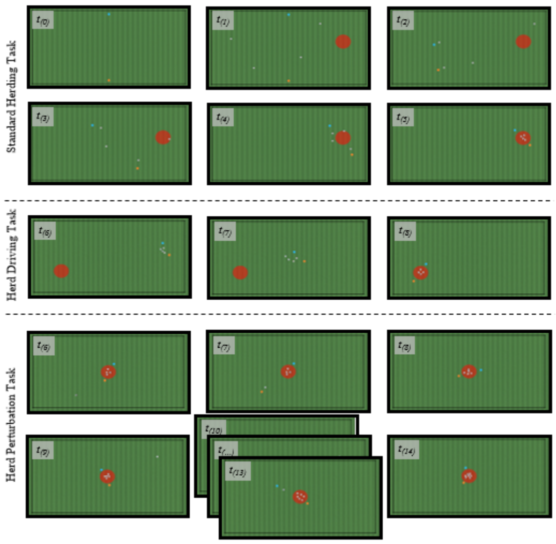

Figure 2: Illustration of the three herding task levels. The orange and blue discs are the herding agents (HAs), the white dots are the targets (TAs), and the containment area is specified in red. See text for more details.

to transfer and adapt these behavioral skills to human-human task performance.

1.4.1 General Method and Apparatus. As illustrated in Figure 1a, the herding task employed here was an extension of the herding task developed by [19, 21] and required two HAs, represented on the game field by blue or orange discs, to corral and contain a herd of autonomous TAs, in this case, a herd of 'cow' textured spheres. HAs were tasked to contain the TAs within a 19.2 cm diameter red containment area positioned at various locations around a 45 cm x 155 cm game field. The herding task was displayed to participants on a 75 inch tabletop touch screen (Model SBID-MX075; SMART Technologies Inc, Calgary, Canada), such that human players could control the position/motion of HAs using a stylus pen. When not within the vicinity of an HA, the TAs' movements were defined by Brownian motion dynamics, such that the TAs moved randomly around the play field with a maximum force of 0.12 N (TAs had a mass of 1 kg) per time step (0.02 s). When a TA came within 12 cm of an HA, however, a TA's Brownian motion dynamics was replaced by a repulsive force, f, that moved the TA directly away from the HA's position at a force proportional to the distance of the HA. More specifically, at any time point, t, if a TA_j (where $j = 1, 2, ..7$) was within $R = 12$ cm of a HA, HA_i (where $i = 1, 2$), TA_j was repelled directly away from HA_i at a force equal to

$$|f_{TA_j}| = \frac{R}{|D_{HA_i}|} \quad (5)$$

where D is the current distance of HA_i and $max(f) = 0.36$, which equates to a maximum acceleration of 36 $\frac{cm}{s^2}$ (see [19, 21] for more details).

Three herding task levels were developed for the current study (see Figure 2). Level 1 was almost identical to the herding task developed by [19] and was considered the ***Standard Task***. This level required two HAs to corral a herd of four TAs inside a red containment area (located either centrally $[x, z] = [0, 0]$, or slightly to the left $[-0.3, 0]$ or right $[0.3, 0]$) of the game field and contain the TA-herd within this area for 20 consecutive seconds. If at any time any TA moved outside the containment area, the containment timer would reset and restart once all TAs were re-contained. HAs had a maximum of 2 minutes to achieve this goal.

Level 2 ***Perturbation Task*** started with the containment area centrally located (at $[x, z] = [0, 0]$) and had a 15-second containment timer. Once the first herd of 4 TAs were successfully corralled and contained, an additional TA would randomly appear outside the containment area. Once this new TA had been corralled, and the new 5-TA herd contained for another 15 seconds, a second and then a third TA appeared, requiring HAs to complete a total of four 15-second containments (for the 4-, 5-, 6- and 7-TA herds). A trial was considered successful only when HAs were able to corral and contain all three additional TAs, including containing the final 7-TA herd for 15 consecutive seconds. HAs had 3 minutes to reach task success.

For Level 3 ***Driving Task***, the first containment area was biased towards either the left or right ($x = \pm 0.6$) of the game field with respect to the initial HA positions ($z = \pm 0.3$), and had a 15-second containment timer. After the first successful containment period, the containment area moved to a second random ($z = \pm 0.3$) location on the opposite (x) side of the game field. HAs were then required to transport the TAs to this new containment area and contain the herd there for an additional 15 seconds. Players were given 3 minutes to reach task success.

The tasks were developed using Unity 2018.2.21f (Exp. 1 and 2) and 2018.4.11f (Exp. 3) (Unity Technologies Ltd, San Francisco, CA). All relevant game data was sampled at 50 Hz (e.g., $[x, z]$ position and velocity of the HAs and TAs). For Exp. 3 (see subsection 4.1), two touchscreens were connected via LAN using the Mirror high-level networking API (source: https://mirror-networking.com/). As illustrated in Figure 1a, human co-actors were required to stay on opposite sides of the game field (touchscreen) during game play, but were free to move along their side of the game field (touchscreen) area. Similarly, during human-AA game play, human-actors were assigned to one of the two sides, with the AA "assigned" to the other side of the game area.

At the start of each trial, only the blue and orange HAs were displayed on the screen and were positioned on opposite sides of the field, 10 cm from the edge of the field (see Figure 2a). A trial started when the participant(s) touched their assigned HA with their pen, which triggered the appearance of the the TAs and the red containment area. A trial ended automatically after the trial time had expired or the success criteria was achieved. At that time, the TAs and target containment area would disappear and either 'FAILURE' or 'SUCCESS' was displayed on the screen. A level was successfully 'completed' after four successful trials.

Importantly, participants were only informed about the basic success criteria for each level. All participants were students and staff from Macquarie University. At the end of all experimental sessions, participants were debriefed about the study's purpose and received course credit or monetary compensation. All participants were unfamiliar with the herding tasks for all experiments. Finally, all of the methods and procedures employed for this study were approved by the Macquarie University's Human Sciences Ethics board.

1.4.2 Behavioral Mode Classification. Based on the previous findings of [19, 21] we expected (and observed) three possible modes of behavior across the different task. First, we expected that all participants would exhibit S&R behavior when corralling or driving the TAs to the containment location. We also expected that some pairs would adopt this mode of behavior when containing the TAs within the containment area. Second, we expected that a subset of pairs would discover and employ OC and COC behaviour when containing the TAs within the containment area. Finally, we expected that some participants in a pair might adopt circling (CIR) behavior in order to keep the TAs contained while their co-participant retrieved/corralled the new TA(s) that appear during the Perturbation Task (Level 2).

To classify these different modes of behavior, a two-step classification procedure was employed. The first step involved determining if and when either HA exhibited CIR behavior. This was achieved by extracting the (x, z) planar movements of the two HAs from the trial recordings, filtering these movement time-series using a 10 Hz low-pass Butterworth filter and then transforming them into polar coordinates centered about the TA-herd's mean position. For each HA, the unwrapped angular position time-series was obtained and rotated so that the first sample $\theta_{t=0} = 0$. Rotations were counted as the number of 2π rotations (in either direction) from $\theta_{t=0}$.

Once the periods of CIR behavior had been identified, a similar procedure employed by [19, 21] was used to classify OC from S&R behavior. Ignoring the periods of time already classified as CIR, a windowed frequency analysis was conducted to determine when each HA in a pair was exhibiting OC behavior, and in turn, when HA pairs were exhibiting COC. A window length of 5 s was employed, with a HA deemed to be exhibiting OC behavior during a 5 s window when they exhibited a dominant peak frequency above 0.5 Hz (consistent with [19, 21]). From this windowed analysis, two measures of oscillatory performance were then extracted for each trial. First, the average proportion of time that each HA in a pair exhibited OC during a trial (*PropOC*) and second, the proportion of time that COC was exhibited by HAs during a trial (*PropCOC*) (i.e., periods where both HAs were adopting OC simultaneously). Finally, note that the proportion of S&R behavior exhibited by an HA during a trial is equal to $1 - (Circ + PropOC)$.

1.4.3 Herding Performance Measures. Across all three levels we employed two measures of general task performance: *completion time* or the time (s) to complete a trial goal as a proportion of the maximum trial time; and *corral time*, which was equal to the amount of time (s) from the beginning of the trial to the beginning of the first successful containment. We also employed two measures of herd containment performance, all of which were calculated during the periods of successful containment: *herd-spread*, which equaled the average herd spread (in cm^2) measured by computing the convex hull formed with respect to all the TAs in the herd (i.e., the smallest convex polygon that encompasses the entirety of an entire herd); and *herd-travel*, which was the cumulative distance (cm) that the herd's mean position (i.e., center-of-mass (COM)) moved during containment.

With regard to the performance of HAs when new TAs appeared in the Perturbation Task, herd-spread and herd-travel were also calculated for the period between when a new TA appeared and when the next successful containment period began. To assess TA retrieval performance, we also calculated new TA *retrieval time*, which was the time (s) from when a new TA appeared to the beginning of the next successful containment period.

Finally, in the Driving Task the performance of HAs when moving the TAs from the first target containment area to the second target containment area was also assessed by means of herd-spread and herd-travel during this driving phase. We also calculated *driving time*, which was the time (s) from when the new containment area appeared to the beginning of the next successful containment period.

2 EXPERIMENT 1

Before we could test whether the AA model proposed by [21] could be employed to train novice human (NH) players to perform these novel herding tasks, we first needed to determine whether the same behavioral dynamics exhibited by human-actors in [19, 21] were observed. Of particular interest was the behaviors exhibited by human-actors during TA retrieval in the Perturbation Task and while driving the herds between containment areas in the Driving Task.

2.1 Method

2.1.1 Participants and Procedure. Sixteen pairs participated in Exp.1. All pairs completed the Standard Task first, then the Perturbation Task, followed by the Driving Task. Participant pairs could only move on to the next level once they had completed four successful trials in the previous level. There was no restriction on the number of trials that participant pairs could attempt in each level. However, participants were given a maximum of 1.5 hours to complete all three levels.

2.1.2 Data Reduction and Analysis. Only 10 of the 16 pairs were able to complete all three task levels. Given that the aim of Exp.1 was to identify the performance of successful human-human behavior, only the following data from successful pairs was retained for analyses: (i) successful trials in the Standard Herding Task, (ii) trials where pairs successfully completed the first herd containment during the Perturbation and Driving Task (i.e., the containment before TA perturbation or containment area change, respectively), (iii) any instances of successful re-containment following TA perturbations in the Perturbation Task, and (iv) only successful trials were used as candidates to investigate driving behavior during the Driving Task.

2.2 Results

As can be discerned from an inspection of Table 1, successful NH-NH pairs discovered and adopted both OC and COC behaviour by the end of the Standard Task. Consistent with [19, 21], the more OC and COC behaviour exhibited by pairs, the better they were at containing the TA-herd within the specified containment area. This is illustrated in Figure 3a, where the proportion of oscillatory behaviour (PropOC) exhibited by a pair for each of the successful containments achieved in the Standard Task is plotted as a function

Table 1 Mean performances measures for Experiments 1 and 2.

Standard Task	Comp. Time	Corral Time	Prop OC	Prop COC	Herd Spread	Herd Travel
NH-NH	0.54	44.53	0.38	0.22	14.32	48.06
AA-NH	0.39	27.12	0.80	0.63	11.25	42.29

Perturbation Task			Initial Containment Phase				TA Retrieval					
	Comp. Time	Corral Time	Prop OC	Prop COC	Herd Spread	Herd Travel	Retrieval Time	Prop OC	Prop COC	Prop CIR	Herd Spread	Herd Travel
NH-NH	0.72	25.10	0.51	0.33	15.16	33.28	10.73	0.11	0.01	0.22	42.42	25.06
AA-NH	0.64	24.89	0.57	0.55	12.69	31.98	8.35	0.18	0.01	0.55	26.11	12.80

Driving Task			Initial Containment Phase				Herd Driving				
	Comp. Time	Corral Time	Prop OC	Prop COC	Herd Spread	Herd Travel	Driving Time	Prop OC	Prop COC	Herd Spread	Herd Travel
NH-NH	0.46	28.59	0.46	0.31	13.49	31.19	32.53	0.08	0.01	44.11	128.17
AA-NH	0.53	33.98	0.70	0.67	11.34	27.41	31.72	0.08	0.00	87.76	128.86

NH = Novice Human; AA = Artificial Agent. Completion (Comp.) time is reported as a proportion of max trial time.
For AA-NH, propOC and PropCIR only correspond to NH behavior. Corral and Retrieval time in seconds. Herd Travel in cm. Herd Spread in cm^2.

of herd-spread and herd-travel. Indeed, a mixed effects regression analysis (with a random intercept for pair) revealed that the relationship between PropOC and Herd Spread (β = -5.93, z = -3.65, p < 0.001) and PropOC and Herd Time (β = -33.24, z = -4.79, p < 0.001) illustrated in Figure 3b were significant, with more oscillatory behavior resulting in smaller herd-spread and less herd-travel during containment.

NH-NH pairs also employed OC and COC behaviour to contain TAs in the Perturbation and Driving Tasks. With regard to TA retrieval, 60% of the successful NH-NH pairs employed circling (CIR) behaviour (at least once). Interestingly, although most pairs continued to transition back to OC and COC behaviour during each containment period, those participants that didn't discover or employ CIR tended to employ S&R containment during TA retrieval, despite the fact the CIR resulted in better herd containment (see Figure 3b). For the Driving Task, pairs adopted S&R behavior to transport the TA-herd towards the new containment area, moving in a way that prioritized driving the TA-Herd's COM towards the new containment area.

3 EXPERIMENT 2

This experiment was designed to provide a preliminary test of whether an HA whose movements and actions were controlled by the adapted version of the DMP-based herding model, detailed in subsection 1.3, could (i) help train a NH to complete the three herding tasks and (ii) result in human-AA performance at or above human-human performance. This was achieved by having NH participants play the herding tasks together with the model-controlled AA and then assessing the performance of these NH-AA teams with respect to NH-NH performance.

3.1 Method

3.1.1 Artificial Herding Agent Control. Given the behaviors in Exp. 1, the following modifications were made to the model described in subsection 1.3 to enable the AA to complete the herding tasks introduced in this study.

Target Agent Selection. At each fixed time step, the TA the AA pursued was updated. The AA selected the TA that was (i) closer to it than the other HA and (ii) was furthest from the containment area. If all the TAs were closer to the AA's partner, then the AA selected the TA that was closest to it. Compared to previous work which fixed candidate TAs to pursue to either side of the game field [21], this selection criteria allowed for the AA to behave adaptively regardless of the location of the containment area, or the workload taken on by the AA's partner.

Circling Behavior. A key insight from Experiment 1 was that the emergence of CIR behavior was an effective control strategy for keeping the TA-herd contained during new TA retrieval. To account for this, the model included a conditional term to specify whether the AA implemented (2) or (4) when computing $\ddot{\theta}_i$. Here, if the AA's partner HA-j's $r_j > r_{max}$ (where r_{max} = 0.288 m), the AA solved for $\ddot{\theta}_i$ using (4) - otherwise (2) was implemented. See Algorithm 1 for pseudocode summarizing AA behavior.

Miscellaneous. The AA control architecture was adapted to (i) account for features of natural human behavior and (ii) ensure the AA was robust to differences between task conditions. These adjustments were tied to whether all TAs were contained ($r_{TA(t),i} < r_\Delta$). Specifically, when $r_{TA(t),i} < r_\Delta$ (i.e., during oscillatory behavior), oscillatory movements were centered at $\theta_{TA(t),i} = 0$, which was centered on the sagittal plane in the direction of the AA's starting position. During oscillatory behavior, the AA kept a minimum distance from the TA herd ($r_{min} = 0.096m$). These changes reflected behaviours observed in previous work [19, 21] and Exp 1.

Finally, the polar task axis was linearly interpolated at each timestep to prevent abrupt changes to the AA's behavior (due to a change

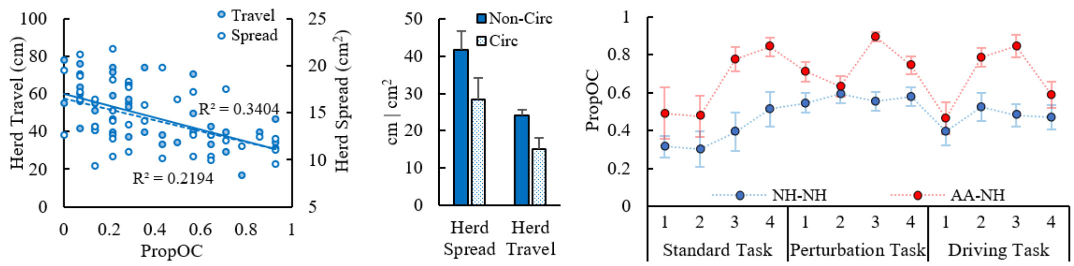

Figure 3: (a) Left: Relationship between PropOC and Herd Containment Performance. (b) Middle: Effects of circling (CIR) behavior on herd-spread and herd-travel during successful TA retrieval in the Perturbation Task. (c) Right: The average magnitude of PropOC behaviour exhibited by i) novice human participant pairs (NH-NH) in Exp.1, and ii) NH-AA: participants playing with the artificial agent in Exp 2. Error bars represent standard error.

in the number of TAs as in the *Perturbation Task* or a change in the containment location (i.e, during the *Driving Task*). Additionally, random acceleration noise ($\ddot{r}_{noise} = \pm 1 \frac{m}{s^2}$ and $\ddot{\theta}_{noise} = \pm 1 \frac{rad}{s^2}$) was added to the AA's control to produce more human-like behavior.

Algorithm 1: Artificial Herding Agent Control

Initialize variables;
while *While task active* **do**
  Select targeted agent (TA), **return** $(r_{f(t),i}, \theta_{f(t),i})$;
  if $r_{TA(t),i} < r_{\Delta}$ **then**
    $(r_{pole}, \theta_{pole}) = (0, 0)$ at containment location;
    $(r_{TA(t),i}, \theta_{TA(t),i}) = (r_{min}, 0)$;
    if $r_j > r_{max}$ **then**
      Implement (4);
    else
      Implement (2);
    end
  else
    $(r_{pole}, \theta_{pole}) = (0, 0)$ at TA-herd position;
  end
  Linearly interpolate pole towards $(r_{pole}, \theta_{pole})$;
  Solve for $\dot{b}_{\theta}$, $\ddot{r}_i$, $\ddot{\theta}_i$. Update b, r, $\dot{r}$, θ, $\dot{\theta}$;
end

3.1.2 Participants and Procedure. Ten participants participated in Exp. 2. The experimental procedure was the same as in Exp.1 with the following exceptions: NH participants completed the task with the AA as their partner. The AA controlled one of the colored disc HAs. The maximum number of attempts to complete the Standard task was limited to ten attempts, with a total of six attempts allowed for the other two tasks. If participants could not successfully complete a level within the maximum number of attempts, they proceeded to the next level. Following completion of the Standard Task, the order of the Perturbation and Driving Tasks was counterbalanced.

3.2 Results

The modified AA described in subsection 3.1.1 was successful in helping NH participants complete the herding tasks. Indeed, 90% of NH participants completed all 3 levels with the AA (compared to only 60% in Exp. 1) (see Table 1 and Figure 3c). Again, the more participants utilized OC behavior, the better they contained the TA-Herd. More importantly, participants in the current experiment learned to utilize OC and COC solutions to contain the TA-herd to a greater extent than NH participants in Exp. 1. Indeed, between-subjects t-tests revealed that during the Standard Task, NH participants working with the AA exhibited significantly greater magnitudes of PropOC ($t(17) = 5.42$, $p < 0.001$) and PropCOC ($t(17) = 3.75$, $p < 0.002$) compared to NH participants in Exp. 1, implying that the AA did in-fact help NHs learn to employ OC and COC behaviour. The positive effects of the AA on NH learning was further highlighted by finding that during the Perturbation Task, NH participants in the current experiment exhibited significantly more instances of CIR behaviour compared to NH participants in Exp. 1 ($t(17) = 2.15$, $p < 0.05$).

Additionally, AA-NH pairs were significantly faster at retrieving TAs than NH-NH pairs ($t(17) = 3.57$, $p < 0.01$). AA-NH pairs were also better at containing the herd during TA retrieval compared to NH-NH pairs, with the difference in herd-travel between AA-NH and NH-NH pairs being significant ($t(17) = 3.14$, $p < 0.01$). There were, however, no significant differences between AA-NH and NH-NH pairs with respect to driving performance (both $t(17) < 1.80$, $p > 0.09$). This was likely due to the fact that in Exp. 1. all NH participants completed the Driving Task after completing both the Standard and Perturbation Tasks, whereas half of NH pairs in the current study completed the Driving task before the Perturbation task (due to counterbalancing) and thus had less task practice.

4 EXPERIMENT 3

This experiment sought to explicitly test whether the AA controlled by the herding model detailed in subsections 1.3 and 3.1.1 could provide training equivalent to an expert human (EH), both in terms of task performance and the ability of NH participants to transfer learning to novel human-human task scenarios. Accordingly, Exp. 3 first had novice human (NH) participants learn to perform the Standard Task with either an AA or one of two EH trainers. Following training, pairs of similarly-trained NH participants then completed the Perturbation and Driving Tasks together as a new NH-NH team.

Table 2 Mean performances measures for Experiment 3.

Standard Task							
NH Trainer	# Attempts	Comp. Time	Corral Time	Prop OC	Prop COC	Herd Spread	Herd Travel
AA	5.2	0.49	38.64	0.62	0.60	10.16	42.76
EH	4.4	0.44	32.61	0.59	0.54	12.80	43.85

Perturbation Task				Initial Containment Phase				TA Retrieval				
NH Trainer	# Attempts	Comp. Time	Corral Time	Prop OC	Prop COC	Herd Spread	Herd Travel	# Retrieved	Retrieval Time	Prop CIR	Herd Spread	Herd Travel
AA	5.6	0.89	40.47	0.56	0.42	12.85	33.65	1.83	12.19	0.34	45.93	30.93
EH	5.0	0.81	28.47	0.61	0.42	12.34	28.93	2.26	13.29	0.21	53.98	32.30

Driving Task				Initial Containment Phase				Herd Driving		
NH Trainer	# Attempts	Comp. Time	Corral Time	Prop OC	Prop COC	Herd Spread	Herd Travel	Driving Time	Herd Spread	Herd Travel
AA	4.7	0.61	41.32	0.56	0.39	13.34	36.93	34.83	52.36	134.43
EH	5.2	0.63	47.42	0.58	0.42	12.17	37.69	35.13	44.61	132.92

For the Stadrard Task, PropOC and PropCIR values only correspond to the behavior of NH. Herd Travel in cm. Herd Spread in cm^2.
Completion (Comp.) time is reported as a proportion of max trial time. Corral and Retrieval time in seconds.

4.1 Method

4.1.1 Participants and Procedure. Twenty pairs participated in Exp. 3. Prior to arrival, participant pairs were randomly assigned to either the AA or EH training condition. In both conditions, participants in a pair were taken to separate laboratory rooms to complete AA or EH training in the Standard task. Following training, participants completed Levels 2 and 3 together, but remotely (in separate rooms). The trial limits were the same as those employed in Exp. 2, and level order following training in Level 1 was counterbalanced.

4.2 Results

A summary of the results of Experiment 3 can be found in Table 2. Participants in both the AA and EH training conditions completed the Standard Task within the maximum allowed 10 attempts, confirming that both the AA and EH were equally effective in training NH participants. Moreover, analysis of the different classification and performance measures revealed no differences in performance between AA and EH trained NH participants, other than a significant difference in herd-spread during containment, $t(18) = -4.46$, $p < .001$, such that herd-spread was lower for AA-NH pairs than for EH-NH. With regard to how the pairs of NH participants performed the Perturbation and Driving Task following training, there were no significant differences in performance between AA and EH trained participants. That is, NH participants trained by the AA demonstrated equal performance to NH pairs trained by an EH, indicating that NH participants were able to transfer the skills learned during both AA and EH training to novel NH-NH task contexts. Note also, that despite never seeing CIR behaviour performed during training, NH participants in both training groups realized that CIR behaviour could be employed to contain a herd during TA retrieval, with 60% of pairs in both the AA and EH training conditions employing CIR during TA retrieval.

5 CONCLUSIONS

The current study employed a multi-agent herding task to demonstrate that task dynamical models of human behaviour (comprised of environmentally coupled DMPs) can be employed to develop AAs that are not only capable of robust, 'human-like' behavioural interaction, but can also provide adaptive training comparable to an expert human trainer. Such AAs could play a critical role in enhancing the effectiveness of team training across a range of industrial, medical, sport and military contexts by reducing the costs of such training, and providing individuals with the opportunity to engage in more frequent and targeted training. It is also likely that AAs embodying human models of perceptual-motor behavior will lead to higher levels of trust between humans and interactive AAs [18].

Future research could employ machine learning approaches (such as imitation and reinforcement learning) to further enhance the capabilities of DMP based AAs by parameterizing DMP-based control architectures [9] to adopt different play-styles or "personalities" [7, 15, 32]. Finally, although the current study investigated team transfer training using the smallest team size possible (*N = 2*), the current modeling methodology can be employed to capture the behaviour of human actors engaged in larger team contexts, including competitive multi-agent tasks. Thus, future research should investigate the efficacy of utilizing DMP-based models in activities involving larger teams, where there is a greater potential for complex patterns of behavior to emerge.

ACKNOWLEDGMENTS

We would like to thank Kristof Bertalan, Dean Bevitt, Yu-Ying Mak, Jonathan Park and Mathew Prants for help with data collection. This work was supported by Australian Research Council Future Fellowship (FT180100447) awarded to Michael Richardson and the Australian Department of Defence, Science and Technology (DST) group (partnership grant MyIP8655).

REFERENCES

[1] Elvio G. Amparore, Marco Beccuti, Marco Botta, Susanna Donatelli, and Fabio Tango. 2018. Adaptive artificial co-pilot as enabler for autonomous vehicles and intelligent transportation systems. In *Proceedings of the Tenth International Workshop on Agents in Traffic and Transportation (ATT 2018)*, Ana Lúcia C. Bazzan, Luca Crociani, Ivana Dusparic, and Sascha Ossowski (Eds.). CEUR Workshop Proceedings, 70–77.

[2] Hendrik Buschmeier and Stefan Kopp. 2018. Communicative listener feedback in human-agent interaction: Artificial speakers need to be attentive and adaptive: Socially interactive agents track. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*, M. Dastani, G. Sukthankar, E. André, and S. Koenig (Eds.). IFAAMAS, Stockholm, Sweden, 1213–1221.

[3] Jessie Y. C. Chen, Shan G. Lakhmani, Kimberly Stowers, Anthony R. Selkowitz, Julia L. Wright, and Michael Barnes. 2018. Situation awareness-based agent transparency and human-autonomy teaming effectiveness. *Theoretical Issues in Ergonomics Science* 19, 3 (5 2018), 259–282. https://doi.org/10.1080/1463922X.2017.1315750

[4] Guillaume Dumas, Gonzalo C. de Guzman, Emmanuelle Tognoli, and J. A. Scott Kelso. 2014. The human dynamic clamp as a paradigm for social interaction. *Proceedings of the National Academy of Sciences* 111, 35 (9 2014), E3726–E3734. https://doi.org/10.1073/pnas.1407486111

[5] Brett R. Fajen and William H. Warren. 2003. Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance* 29, 2 (2003), 343–362. https://doi.org/10.1037/0096-1523.29.2.343

[6] Gerald F. Goodwin, Nikki Blacksmith, and Meredith R. Coats. 2018. The science of teams in the military: Contributions from over 60 years of research. *American Psychologist* 73, 4 (5 2018), 322–333. https://doi.org/10.1037/amp0000259

[7] Jamie C. Gorman, Nancy J. Cooke, and Polemnia G. Amazeen. 2010. Training Adaptive Teams. *Human Factors* 52, 2 (4 2010), 295–307. https://doi.org/10.1177/0018720810371689

[8] Neville Hogan and Dagmar Sternad. 2007. On rhythmic and discrete movements: reflections, definitions and implications for motor control. *Experimental Brain Research* 181, 1 (6 2007), 13–30. https://doi.org/10.1007/s00221-007-0899-y

[9] Auke Jan Ijspeert, Jun Nakanishi, Heiko Hoffmann, Peter Pastor, and Stefan Schaal. 2013. Dynamical movement primitives: Learning attractor models formotor behaviors. (2013). https://doi.org/10.1162/NECO{_}a{_}00393

[10] B. A. Kay, J. A.S. Kelso, E. L. Saltzman, and G. Schöner. 1987. Space-time behavior of single and bimanual rhythmical movements: Data and limit cycle model. *Journal of Experimental Psychology: Human Perception and Performance* 13, 2 (5 1987), 178–192. https://doi.org/10.1037/0096-1523.13.2.178 arXiv:cond-mat/0212203v1

[11] J. A. Scott Kelso. 1995. *Dynamic patterns: The self-organization of brain and behavior.* MIT Press. 358 pages.

[12] J. A. Scott Kelso, Gonzalo C. de Guzman, Colin Reveley, and Emmanuelle Tognoli. 2009. Virtual Partner Interaction (VPI): Exploring Novel Behaviors via Coordination Dynamics. *PLoS ONE* 4, 6 (6 2009), e5749. https://doi.org/10.1371/journal.pone.0005749

[13] Maurice Lamb, Tamara Lorenz, Stephen J. Harrison, Rachel Kallen, Ali Minai, and Michael J. Richardson. 2017. PAPAc: A Pick and Place Agent Based on Human Behavioral Dynamics. In *Proceedings of the 5th International Conference on Human Agent Interaction - HAI '17*. ACM Press, New York, New York, USA, 131–141. https://doi.org/10.1145/3125739.3125771

[14] Maurice Lamb, Patrick Nalepka, Rachel W. Kallen, Tamara Lorenz, Steven J. Harrison, Ali A. Minai, and Michael J. Richardson. 2019. A Hierarchical Behavioral Dynamic Approach for Naturally Adaptive Human-Agent Pick-and-Place Interactions. *Complexity* 2019 (6 2019), 1–16. https://doi.org/10.1155/2019/5964632

[15] Maria Lombardi, Davide Liuzza, and Mario di Bernardo. 2019. Deep learning control of artificial avatars in group coordination tasks. (6 2019). http://arxiv.org/abs/1906.04656

[16] Souvik Maity. 2019. Identifying opportunities for artificial intelligence in the evolution of training and development practices. *Journal of Management Development* 38, 8 (2019), 651–663. https://doi.org/10.1108/JMD-03-2019-0069

[17] Shiwali Mohan, Kalai Ramea, Bob Price, Matthew Shreve, Hoda Eldardiry, and Les Nelson. 2019. Building Jarvis - A learner-aware conversational trainer. In *Joint Proceedings of the ACM IUI 2019 Workshops*, Christoph Trattner, Denis Parra, and Nathalie Riche (Eds.). ACM, New York, NY.

[18] Alexander Mörtl, Tamara Lorenz, and Sandra Hirche. 2014. Rhythm patterns interaction - Synchronization behavior for human-robot joint action. *PLoS ONE* 9, 4 (1 2014), e95195. https://doi.org/10.1371/journal.pone.0095195

[19] Patrick Nalepka, Rachel W. Kallen, Anthony Chemero, Elliot Saltzman, and Michael J. Richardson. 2017. Herd those sheep: Emergent multiagent coordination and behavioral-mode switching. *Psychological Science* 28, 5 (2017), 630–650. https://doi.org/10.1177/0956797617692107

[20] Patrick Nalepka, Rachel W. Kallen, Anthony Chemero, Elliot Saltzman, and Michael J. Richardson. 2019. Practical Applications of Multiagent Shepherding for Human-Machine Interaction. In *Advances in Practical Applications of Survivable Agents and Multi-Agent Systems*, Yves Demazeau, Eric Matson, Juan Manuel Corchado, and Fernando De la Prieta (Eds.). Springer Nature, 168–179. https://doi.org/10.1007/978-3-030-24209-1{_}14

[21] Patrick Nalepka, Maurice Lamb, Rachel W. Kallen, Kevin Shockley, Anthony Chemero, Elliot Saltzman, and Michael J. Richardson. 2019. Human social motor solutions for human–machine interaction in dynamical task contexts. *Proceedings of the National Academy of Sciences* 116, 4 (1 2019), 1437–1446. https://doi.org/10.1073/PNAS.1813164116

[22] Michael J. Richardson, Steven J. Harrison, Rachel W. Kallen, Ashley Walton, Brian A. Eiler, Elliot Saltzman, and R. C. Schmidt. 2015. Self-organized complementary joint action: Behavioral dynamics of an interpersonal collision-avoidance task. *Journal of Experimental Psychology: Human Perception and Performance* 41, 3 (6 2015), 665–679. https://doi.org/10.1037/xhp0000041

[23] Michael J. Richardson and Rachel W. Kallen. 2015. Symmetry-Breaking and the Contextual Emergence of Human Multiagent Coordination and Social Activity. *World Scientific Review* 6 (2015), 1–57. https://doi.org/10.1142/9789814730617

[24] Michael J. Richardson, Rachel W. Kallen, Patrick Nalepka, Steven J. Harrison, Maurice Lamb, Anthony Chemero, Elliot Saltzman, and R. C. Schmidt. 2016. Modeling embedded interpersonal and multiagent coordination. In *Proceedings of the 1st International Conference on Complex Information Systems*, Víctor Méndez Muñoz, Oleg Gusikhin, and Victor Chang (Eds.). SciTePress, Setúbal, Portugal, 155–164. https://doi.org/10.5220/0005878101550164

[25] Eduardo Salas, Nancy J. Cooke, and Jamie C. Gorman. 2010. The Science of Team Performance: Progress and the Need for More.... *Human Factors: The Journal of the Human Factors and Ergonomics Society* 52, 2 (4 2010), 344–346. https://doi.org/10.1177/0018720810374614

[26] Eduardo Salas, Deborah DiazGranados, Cameron Klein, C. Shawn Burke, Kevin C. Stagl, Gerald F. Goodwin, and Stanley M. Halpin. 2008. Does Team Training Improve Team Performance? A Meta-Analysis. *Human Factors:* 50, 6 (12 2008), 903–933. https://doi.org/10.1518/001872008X375009

[27] Elliot Saltzman and David Caplan. 2015. A Graph-Dynamic Perspective on Coordinative Structures, the Role of Affordance-Effectivity Relations in Action Selection, and the Self-Organization of Complex Activities. *Ecological Psychology* 27, 4 (10 2015), 300–309. https://doi.org/10.1080/10407413.2015.1086228

[28] Elliot Saltzman and J. A. Kelso. 1987. Skilled actions: A task-dynamic approach. *Psychological Review* 94, 1 (1987), 84–106. https://doi.org/10.1037/0033-295X.94.1.84

[29] Stefan Schaal, Shinya Kotosaka, and Dagmar Sternad. 2000. Nonlinear dynamical systems as movement primitives. In *Proceedings of the 1st IEEE-RAS International Conference on Humanoid Robotics.* IEEE, Piscataway, NJ.

[30] Richard C. Schmidt and Michael J. Richardson. 2008. Dynamics of Interpersonal Coordination. In *Coordination: Neural, Behavioral and Social Dynamics*, Armin Fuchs and Viktor K. Jirsa (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 281–308. https://doi.org/10.1007/978-3-540-74479-5{_}14

[31] Marissa L. Shuffler, Davin Pavlas, and Eduardo Salas. 2012. Teams in the military: A review and emerging challenges. In *Oxford library of psychology. The Oxford handbook of military psychology*, J. H. Laurence and M. D. Matthews (Eds.). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780195399325.013.0106

[32] Piotr Słowiński, Chao Zhai, Francesco Alderisio, Robin Salesse, Mathieu Gueugnon, Ludovic Marin, Benoit G. Bardy, Mario Di Bernardo, and Krasimira Tsaneva-Atanasova. 2016. Dynamic similarity promotes interpersonal coordination in joint action. *Journal of the Royal Society Interface* 13, 116 (2016), 20151093. https://doi.org/10.1098/rsif.2015.1093

[33] Dagmar Sternad, Hiromu Katsumata, Marcos Duarte, and Stefan Schaal. 2001. Bouncing a ball: Tuning into dynamic stability. (2001), 1163–1184 pages. https://doi.org/10.1037/0096-1523.27.5.1163

[34] Konstantinos Tsiakas, Maher Abujelala, Alexandros Lioulemes, and Fillia Makedon. 2016. An intelligent Interactive Learning and Adaptation framework for robot-based vocational training. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 1–6. https://doi.org/10.1109/SSCI.2016.7850066

[35] Annemarie Turnwald and Dirk Wollherr. 2019. Human-Like Motion Planning Based on Game Theoretic Decision Making. *International Journal of Social Robotics* 11, 1 (1 2019), 151–170. https://doi.org/10.1007/s12369-018-0487-2

[36] Jurriaan van Diggelen, Mark Neerincx, Marieke Peeters, and Jan Maarten Schraagen. 2019. Developing Effective and Resilient Human-Agent Teamwork Using Team Design Patterns. *IEEE Intelligent Systems* 34, 2 (3 2019), 15–24. https://doi.org/10.1109/MIS.2018.2886671

[37] William H. Warren. 2006. The dynamics of perception and action. *Psychological Review* 113, 2 (2006), 358–389. https://doi.org/10.1037/0033-295X.113.2.358

[38] Auriel Washburn, Rachel W. Kallen, Maurice Lamb, Nigel Stepp, Kevin Shockley, and Michael J. Richardson. 2019. Feedback delays can enhance anticipatory synchronization in human-machine interaction. *Plos One* 14, 8 (8 2019), e0221275. https://doi.org/10.1371/journal.pone.0221275

[39] Logan Yliniemi and Kagan Tumer. 2016. Using Awareness to Promote Richer, More Human-Like Behaviors in Artificial Agents. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Nardine Osman and Carles Sierra (Eds.). Springer, 122–133. https://doi.org/10.1007/978-3-319-46840-2{_}8