This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101092889, Topic HORIZON-CL4-2022-HUMAN-01-14

SHARESPACE

Embodied Social Experiences in Hybrid Shared Spaces

Project Reference No	101092889
Deliverable	D5.4 Architecture design and validation (L3 avatars)
Workpackage	WP5: Cognitive Architecture
Nature	D (Deliverable)
Dissemination Level	PU – Public
Date	31/12/2023
Status	Draft v1.0
Editor(s)	Antonio Grotta (CRdC) Marco Coraggio (CRdC) Francesco De Lellis (CRdC) Mario di Bernardo (CRdC)
Involved Institutions	CRdC, UKE, DFKI, UM, Golaem
Document Description	This deliverable reports the first version of the Cognitive Architectures to be used to drive L3 virtual humans in the SHARESPACE project. We introduce algorithms and strategies to drive the motion of fully autonomous (L3) virtual humans in shared hybrid spaces. The architectures proposed are first described and then validated both numerically and experimentally, considering the specificities of the proof-of-principle scenarios.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

Term / Abbreviation	Definition
AI	artificial intelligence
CA	cognitive architecture
L0, L1, L2, L3	Level of autonomy
NN	neural network
PoP	Proof of Principle
RL	reinforcement learning
VH	virtual human

Table 1: List of Abbreviations

1 INTRODUCTION

1.1 PURPOSE OF THE DOCUMENT

This deliverable presents the design and first implementation of the *Cognitive Architecture* (CA) that can drive the behavior of the L3 virtual humans (see the living glossary, [deliverable D1.1]) in the SHARESPACE platform. In particular, we provide a description of the algorithms, datasets, and feedback control strategies collected and developed within the first 12 months of the project, used to design the first versions of the CA for L3 virtual humans. Moreover, we present a thorough numerical and experimental validation of the proposed architectures, also verifying their compatibility with the software and hardware constraints of the other components of the SHARESPACE platform as described in deliverable D1.7.

1.2 STRUCTURE OF THE DOCUMENT

The remainder of this document is organized as follows:

- **Section 2: State of the art on cognitive architectures for motor tasks:** Here, we provide an overview of the current state of the art on cognitive architectures to drive virtual humans in sensorimotor tasks.
- **Section 3: The SHARESPACE vision: Autonomy levels:** In this section, we recall the concept of autonomy levels of virtual humans in SHARESAPCE, classifying them into L1, L2 and L3. We focus on describing L3 as a fully autonomous virtual human capable of adjusting and adapting its actions in response to people in the group.
- **Section 4: L3 CA in the Proofs of Principle of Social Connectedness**: This section describes the AI-based methodology employed in the design of the Cognitive Architecture (CA) for the PoP of Social Connectedness. Additionally, it presents the validation of the CA both numerically (with real experimental data) and experimentally, proving its effectiveness in practical scenarios.
- **Section 5: L3 CA in the Proofs of Principle of Amplification**: Here, we present the implementation details of the CA for the PoP of Amplification. We outline the specific steps and strategies used in deploying the CA in this PoP, focusing on how it enhances the PoP's functionalities. Additionally, we also provide numerical examples to validate the CA's performance.
- **Section 7: Conclusions and future work:** In the final section, we draw conclusions based on the preceding discussions and outline the directions for future research and development.

1.3 METHODOLOGY AND APPROACH

The first version of the L3 CA presented here has been designed to satisfy the requirements of the two *Proof of Principles* (PoPs), presented in deliverable D1.2, and to allow compatibility and integration with the *System Architecture*, described in deliverable D1.7, taking into account possible advantages

and disadvantages of the various alternatives in the design process. The design and development of the CA was discussed with partners in the consortium operating in other relevant work packages (i.e., WP1: "System Architecture", WP2: "Sensorimotor primitives of social interaction", WP3: "Capturing", and WP4: "Rendering"). Furthermore, we report both numerical and experimental results concerning the PoP of Social Connectedness to study the performance of the strategies on which the CA is based, when combined with all other elements of the SHARESPACE environment.

2 STATE OF THE ART ON COGNITIVE ARCHITECTURES FOR MOTOR TASKS

In the literature, there are several examples of cognitive architectures designed to drive avatars in motor tasks while coordinating with humans.

A number of examples are concerned with applications of the *mirror game*, originally introduced in (Noy et al., 2011), where two people perform a jointly improvised dance-like motor task in front of each other, even as simple as just moving an arm, with the aim of synchronizing their motion in the absence of a designated leader. In particular, people are told that "the mirror game is a collaborative game whose purpose is to enjoy creating motion together that is synchronized and interesting". To the best of our knowledge, the first cognitive architecture to drive an autonomous avatar interacting with a human in the mirror game was presented in (Zhai et al, 2014) by researchers at the University of Naples Federico II; therein, a CA was created by combining a Haken-Kelso-Bunz (HKB) model (Haken et al., 1985) with a nonlinear feedback controller to either lead or follow the motion of a human being. In (Zhai et al, 2016), a model of interaction is identified to describe the motion of two people playing the mirror game. Then, the model is used to drive a computer avatar able to successfully improvise joint motion with a human; in particular, interaction with the avatar exhibits the three characteristics of joint improvisation, which are high movement synchronization, behavioral plasticity, and absence of a clear leader. Later, in (Zhai et al., 2018a) the same problem was solved via two strategies, one based on adaptive control and the other on optimal control. In these papers, a deterministic and relatively accurate model describing the motion of a human player is assumed to be known. A later architecture able to generate original motion (without exploiting pre-recorded motion) was presented in (Zhai et al, 2018b). A model-free cognitive architecture, based on reinforcement learning was introduced in (Lombardi et al., 2021b). Specifically, therein a virtual player establishes its motion to mimic a certain person, through a (tabular) Q-learning algorithm, trained by observing the movement of that person playing the mirror game with another human.

This dyadic game has been extended to a multi-player experimental setup, where multiple people are asked to move a finger in an oscillatory manner, with the aim to synchronize their motion with that of the others. In (Alderisio et al., 2017), it was shown that this task can be modeled by a Kuramoto network model (Kuramoto, 1975). Moreover, in (Alderisio et al., 2017b), a software platform named Chronos was presented to perform coordination and synchronization experiments via "leap motion" controllers and a local Wifi network (for more details see https://dibernardogroup.github.io/Chronos/index.html).

Later, further investigations were carried out to better characterize the dynamics of groups of people performing this synchronization task. In particular, in (Bardy et al., 2020), the capability was studied of people to maintain coordination after perceptual contact was lost, finding that coordination can be held for about 7 s; in (Calabrese et al., 2021), the emergence of leadership was investigated through an analysis based on causation entropy, with the observed result that different kinds of leaders are possible: those who place themselves at the front of the group, those who seat at the back, and combinations of the two. In (Calabrese et al. 2022), a reduction of the frequency when participants are synchronized was observed with respect to the average of the frequencies observed in isolated motion of the participants. Finally, in (Bienkiewicz et al., 2023) it was observed that positive and negative emotional stimulation can alter the motor dynamics of the participants.

Cognitive architectures that drive virtual humans in this multi-agent group interaction task, with the goal of showing a specific *motion signature* (which is a distribution of frequencies occurring during motion), were presented in (Lombardi et al., 2019) and in (Lombardi et al., 2021a), using deep reinforcement learning; in the former paper, the AI-based agent is trained by taking as input the mean position of its neighbors, whereas in the latter it takes directly all the positions from its neighbors. This architecture is trained using a combination of data from real experiments and a sim-to-real paradigm where synthetic data is generated via appropriate models of multiple agents playing the mirror game.

Studies conducted on human participants point to human intention as an important factor in determining the emergence of collective behavior in human interactions. For instance, (Babajanyan et al., 2022) conducted a comprehensive analysis of human participants engaged in a "pick and place" task. Their research illuminates how hidden intentions can be discerned from the emergent behavioral patterns when individuals collaborate to achieve a shared objective. Furthermore, (Scaliti et al., 2023) revealed the feasibility of discerning human intentions from recorded kinematic data when individuals are tasked with completing specific assignments.

Recently, human coordination in human groups was studied in-depth in (Shalal et al., 2020), where the authors investigated the emergence of synchronization in a group of expert violin players playing a repeating phrase. The effects of delays in the perception of music and different topologies were studied, to characterize the ability of the players to maintain synchronization and the strategies they used for this.

In (McKee at al., 2023), a network was considered of people playing a monetary game, where selfishness typically led to high short-term gains, but altruism led to higher long-term ones; a reinforcement learning-based agent was trained to suggest the creation or suppression of interaction links between the participants: the AI agent outperformed other state-of-the-art strategies in spreading altruism in the human network. Additionally, in (Patil et. al, 2023) a shepherding task was considered, and autonomous RL-based agents were trained, with hierarchical strategies. These agents outperformed heuristic-based ones used in previous research in human-agent coordination, also when the latter were based on models synthesized from expert human behavior.

In conclusion, the existing literature provides a foundation for the development of cognitive architectures in collaborative motor tasks, but open gaps remain concerning the creation of cognitive architectures that have wide applicability, are robust to perturbation, uncertainty, and noise, and are flexible to adapt to different goals, in dynamic virtual environments.

3 THE SHARESPACE VISION: AUTONOMY LEVELS

Let us recall the definitions given in deliverable D1.1 concerning the levels of autonomy of virtual humans in SHARESPACE. In particular, the notation "L#", where # is an integer number between 1 and 3 (included), quantifies how autonomous a virtual human (VH) is. Specifically, as described in D1.1, L3 denotes "a VH whose movements are the sole result of a computation by the cognitive architecture and not the unaltered or altered version of the motion of any single human and are performed with the objective to achieve a collective goal".[1]

The inputs, outputs, and main functions of the L3 CA, in a high-level description, can be found in Figure 1 and Table 2 and in the block diagram in Figure 2. These specifications have been discussed and agreed with other partners of the Consortium, namely those operating in workpackages WP1 ("System Architecture"), WP3 ("Capturing"), and WP4 ("Rendering"). In particular, the goal of the L3 CA is to generate a realistic motion, so that the L3 virtual human driven by the CA can be successfully integrated in the social group; moreover, the motion is generated to attain a specific goal, such as amplifying or blocking the propagation of social information or help the group in performing a collective task. To generate this motion, the L3 CA requires: (i) data about the current motion of all virtual humans it is connected to in the hybrid space; (ii) a criterion to guide the CA in the generation of the motion (this typically being some metrics to assess the compliance of the generated motion with some goal); (iii) either an exhaustive library of possible task-specific motions (currently under development in task T2.5 [Social interaction dataset]) or a model of the motion being performed. Note that the availability of a movement primitives library is necessary to inform the CA on what kind of motion is expected in the current task. Also, the motion generated by the CA is original, as the CA is not merely selecting a motion from a database of possible ones. Then, the specific way in which the motion is generated depends on the implementation at hand, which might depend also on the scenario being considered.

Goal	Inputs	Outputs
Synthesize a motion that achieves a specific collective goal, while appearing realistic and human-like	• Motion of VHs to which L3 is connected to in hybrid space • Metrics to assess the compliance of the input and output motion with a goal • movement primitives library or model to guide the synthesis of realistic motion	Motion of a VH in hybrid space

Table 2: High-level description of the cognitive architectures for different levels of autonomy of virtual humans.

[1]Additionally, L0 refers to a person in a physical space, L1 is the level of autonomy of a fully human-driven VH (i.e., whose motion faithfully replicates that of the driving person), whereas L2 is a semi-autonomous VH (i.e., whose motion is the altered version of that of a person). CAs able to drive L1 and L2 VHs are discussed in detail in deliverable D5.1.

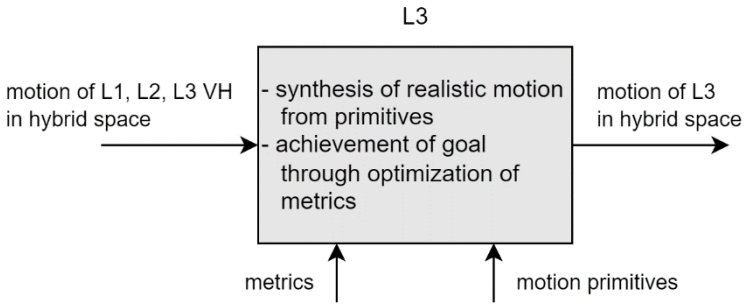

Figure 1: Abstract schematics of the input/output and tasks for the L3 cognitive architecture.

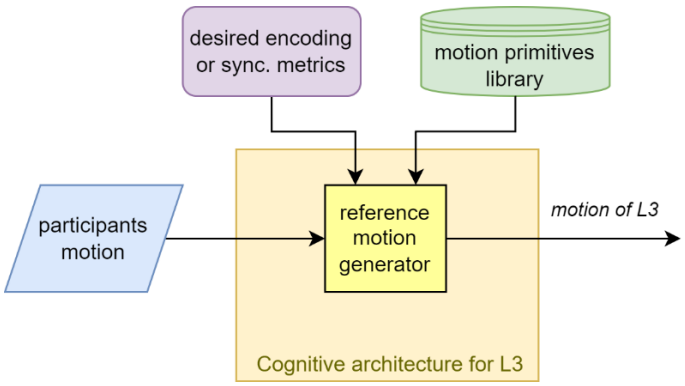

Figure 2: High level block diagram of the L3 cognitive architecture.

3.1 Assumptions of the Cognitive Architecture

Next, we list the assumptions that were used in the design of the L3 CA.

A1. Motions that are input and output to the CA are multidimensional discrete-time signals, of dimension equal to $3d$ over a fixed timespan $T \in \mathbb{R}_{>0}$, where $d \in \mathbb{N}_{\geq 1}$ is a number of points in space (e.g., an arm position, a finger position).

A2. For each task and associated goal, at least one metric (i.e., a function that produces a real number) is available to assess how good a certain motion is with respect to the collective goal.

A3. For each task, either a movement primitive library or a model of motion is provided.

- The library is a dataset of paradigmatic motions, called *movement primitives*; the library is exhaustive with respect to the kinds of motion that can be performed in a given task, in the sense that for each motion that can reasonably be performed in the task, there is a motion in the library which is sufficiently close to it (according to some distance function).
- The motion model is a set of mathematical relationships that can map the motion in the 3d space to a low-dimensional representation of it, and the low-dimensional representation back to a motion in the 3d space.

A4. L3 cognitive architectures generate only the motion of one or a few degrees of freedom of the skeleton representation of a virtual human (e.g., the end-effector position, such as the position of a fingertip or of a wrist); the inverse kinematics being left as part of the animation.

4 L3 CA IN THE PROOFS OF PRINCIPLE OF SOCIAL CONNECTEDNESS

4.1 ROLE AND REQUIREMENTS OF THE L3 CA IN THE PROOF OF PRINCIPLE OF SOCIAL CONNECTEDNESS

In this PoP, several human participants are asked to perform a synchronization task. Specifically, standing in a circle, facing its center, they are instructed to oscillate their preferred hand as smoothly as possible back and forth on a plane parallel to the ground, along a straight trajectory (see Figure 3). The participants are told that their goal is synchronizing their motion while oscillating. The PoP has several phases, in some of which at least one of the participants can receive a mild aversive stimulus.

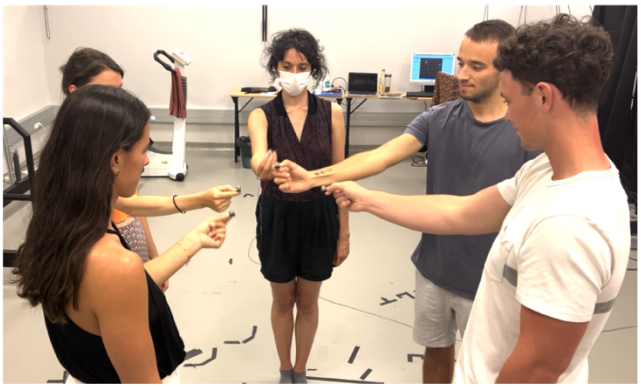

Figure 3: A picture showing an example of the scenario. People are moving their arms to synchronize their motion.

The role of L3, which is maneuvering an autonomous virtual human (when present), is to increase the level of synchronization as much as possible. Additional details concerning the setup of the PoP and the use of L3 are given in deliverable D1.2. To carry out this task, L3 needs to have access to the current situation in the group, that is the positions of the hands of other virtual humans (including L1 participants, semi-autonomous L2 VHs, and any other autonomous L3 VHs, if present). Moreover, the assumptions in § 3.1 need to hold. In particular, Assumption A1 holds with $d = 1$, since the CA receives the position of the index finger of each VH; Assumption A2 is satisfied by the fact that at least one metric of synchronization is available (see § 4.2.4 below); Assumption A3 also holds, as a model of motion is available that can be used to generate motion compatible with the task at hand; finally, Assumption A4 is verified as the CA will generate the motion of just the index finger of its VH.

4.2 IMPLEMENTATION OF L3 CA IN THE PROOF OF PRINCIPLE OF SOCIAL CONNECTEDNESS

4.2.1 DESCRIPTION AND MODELING OF THE ENVIRONMENT

To make the CA capable of responding to, or generating a large variety of periodic motion, we have designed an implementation that considers just the *phase* of the oscillatory motion of interest rather

than its full description in the three-dimensional space. The formal meaning of phase, for a motion in the 3d space is explained in § 4.2.2 below.

4.2.2 PHASE REDUCTION

Given a dynamical system $\frac{dx(t)}{dt} = f(x(t))$ with $x \in \mathbb{R}^n$ exhibiting a *stable limit cycle* $\Gamma \subset \mathbb{R}^n$ (isolated periodic orbit) of period T (Strogatz, 2015), according to (Kralemann et al., 2008), we can define a *bijective* function $f_\theta: \Gamma \to \mathbb{S}$ (where $\mathbb{S}$ is the circle, i.e., $[0, 2\pi)$ or $[-\pi, \pi)$) that maps each point of the limit cycle Γ to a phase $\theta(x)$, with $\frac{d\theta(x(t))}{dt} = \omega = \frac{2\pi}{T}$, where ω is the *natural frequency* of the periodic orbit. Let $B(\Gamma)$ be the *basin of attraction* of the limit cycle Γ. Given a point $x \in B(\Gamma)$, there is a set of points, say $\Phi(x)$, called an *isochron*, such that the trajectories starting from $\Phi(x)$ converge asymptotically to the same trajectory (that is, to the same phase). Thus, it is possible to extend f_θ to a function $\hat{f}_\theta: B(\Gamma) \to \mathbb{S}$, so that all points in an isochron are mapped to the same so called *latent phase*; clearly, $\hat{f}_\theta$ remains a surjective function, but is not injective, and thus it is not invertible. Practical methods that can be used to estimate the (latent) phase of the motion of the index fingers of VHs are reported below in § 4.2.5.1.

4.2.3 PHASE MODELLING OF A PERIODIC MOTOR TASK

In Alderisio et al. (2017), it was shown that when people are performing the oscillating motion used in this PoP in isolation, their motion tends to be periodic and can be associated to a phase, changing with an approximately constant natural frequency, which can be estimated from recorded data of the participants. Moreover, when people are interacting with each other, as happens in the PoP, the evolution of their phases $\theta_i(t)$ can be described by a network of heterogeneous Kuramoto oscillators, as

$$\dot{\theta}_i(t) = \omega_i + \frac{c}{N} \sum_{j \in \mathcal{N}_i} a_{ij} \sin\left(\theta_j(t) - \theta_i(t)\right), \qquad i = 1,2,\dots,N, \tag{1}$$

where $\omega_i \in \mathbb{R}$ is the natural frequency of the i-th oscillator, $\theta_i(t) \in \mathbb{S}$ approximates the phase of the i-th participant's motion, $c \in \mathbb{R}_{>0}$ is the coupling strength between node i and node j, $\mathcal{N}_i$ is the set of nodes that are neighbors of node i and a_{ij} is 1 if node i is coupled with node j and 0 otherwise, and N is the number of all agents (i.e., oscillators) in the group (including L3, if present). The value of $\omega_i(t)$ can be estimated by having participant i carry out the oscillatory task in isolation. To allow compatibility with the animation and rendering equipment, we discretize (*1*) and consider a discrete-time dependent phase $\theta_i(k)$ for each participant, evolving according to the following equation:

$$\theta_i(k+1) = \theta_i(k) + \Delta t \left(\omega_i(k) + \frac{c}{N} \sum_{j \in \mathcal{N}_i} a_{ij} \sin\left(\theta_j(k) - \theta_i(k)\right) \right), \tag{2}$$

where $\Delta t \in \mathbb{R}$ is the time step of the discretization. To have the motion of the L3 VH(s) be intrinsically similar to those of people, we propose to generate the VH's phase, say $\theta_{L3}(k)$, using a model similar to (*2*), given by

$$\theta_{\mathrm{L3}}(k+1) = \theta_{\mathrm{L3}}(k) + \Delta t\left(\omega_{\mathrm{L3}}(k) + \frac{c}{N}\sum_{j\in\mathcal{N}_{\mathrm{L3}}} a_{\mathrm{L3}j}\sin\left(\theta_j(k) - \theta_{\mathrm{L3}}(k)\right)\right), \tag{3}$$

where $\omega_{\mathrm{L3}}(k)$ is decided at each time step k by the CA.

4.2.4 METRICS

We consider three main kinds of metrics to assess the quality of the synchronization between the participants from (Alderisio et al., 2017): the *order parameter*(s), the *group synchronization index* (indices) and the *individual synchronization index* (indices). The first quantifies phase synchronization, the second quantifies frequency synchronization, while the third quantifies how much agent i is phase-synchronized with the rest of the group. All the metrics are described in more detail below.

4.2.4.1 ORDER PARAMETER

We quantify the level of phase synchronization among a group of N agents, whose phases at time k are denoted as $\theta_1(k), \dots, \theta_N(k) \in [0, 2\pi]$ via the *order parameter*. More specifically, we define the *total* order parameter and the *net* order parameter. The *total order parameter* at time k, denoted by $r_{\mathrm{tot}}(k) \in [0, 1]$, is given by

$$r_{\mathrm{tot}}(k) = |q_{\mathrm{tot}}(k)|, \tag{4}$$

where $q_{\mathrm{tot}}(k) \in \mathbb{C}$ is the *average phasor* associated to all agents (including L3), and is given by

$$q_{\mathrm{tot}}(k) = \frac{1}{N}\sum_{i=1}^{N} e^{\mathrm{j}\theta_i(k)}. \tag{5}$$

On the other hand, the *net order parameter* at time k, denoted by $r_{\mathrm{net}}(k) \in [0, 1]$, is given by

$$r_{\mathrm{net}} = |q_{\mathrm{net}}(k)|, \tag{6}$$

where $q_{\mathrm{net}}(k) \in \mathbb{C}$ is the average phasor associated to all agents except for L3, and is given by

$$q_{\mathrm{net}}(k) = \frac{1}{N-1}\left(\sum_{\substack{i=1,\\ i\neq \mathrm{L3}}}^{N} e^{\mathrm{j}\theta_i(k)} - e^{\mathrm{j}\theta_{\mathrm{L3}}(k)}\right). \tag{7}$$

In general, phase synchronization is high when r_{tot} (or r_{net}) is close to 1, and low when it is close to 0.

4.2.4.2 GROUP SYNCHRONIZATION INDEX

To quantify the level of frequency synchronization of the ensemble, we use the *group synchronization index* (Alderisio et al., 2017). As before, we consider both a *total* group synchronization index and a net synchronization index. In particular, the total group synchronization index is computed as follows. Let $\psi_{\mathrm{tot}}(k) \in [0, 2\pi]$ denote the phase of the average phasor of the group:

$$\psi_{\mathrm{tot}}(k) = \tan^{-1}\frac{\mathrm{Im}(q_{\mathrm{tot}}(k))}{\mathrm{Re}(q_{\mathrm{tot}}(k))}. \tag{8}$$

Then, define the difference of the phase of agent i with respect to the group as

$$\phi_{\mathrm{tot},i}(k) = \theta_i(k) - \psi_{\mathrm{tot}}(k). \tag{9}$$

Averaging the phasor of this difference over a time window of length T, we obtain

$$\bar{\phi}'_{\mathrm{tot},i} = \frac{1}{T}\sum_{l=1}^{T} e^{\mathrm{j}\phi_{\mathrm{tot},i}(l)}, \tag{10}$$

And taking again the phase, we get

$$\bar{\phi}_{\mathrm{tot},i} = \tan^{-1}\frac{\mathrm{Im}\left(\bar{\phi}'_{\mathrm{tot},i}\right)}{\mathrm{Re}\left(\bar{\phi}'_{\mathrm{tot},i}\right)} \in [0,2\pi], \tag{11}$$

Which can be understood as the average distance of the phase of agent i with respect to the group, in the time-window of length T. Then, $\Delta\phi_{\mathrm{tot},i}(k) = \phi_{\mathrm{tot},i}(k) - \bar{\phi}_{\mathrm{tot},i}$ can be intended as a measure of how much agent i is changing its phase with respect to the average phase of the group: for instance, $\Delta\phi_{\mathrm{tot},i}(k) = 0$ means that agent i is maintaining, at time k, the same displacement it has had with respect to the group over a time-window of length T; on the other hand, the larger $\left|\Delta\phi_{\mathrm{tot},i}(k)\right|$, the more agent i is changing its displacement, that is, its frequency is different from that of the group. Finally, the (*total*) *group synchronization index* at time k, denoted by $\rho_{\mathrm{tot}}(k) \in [0, 1]$, is given by

$$\rho_{\mathrm{tot}}(k) = \frac{1}{N}\left|\sum_{i=1}^{N} e^{\mathrm{j}(\Delta\phi_{\mathrm{tot,i}}(k))}\right|. \tag{12}$$

Values close to 1 (resp. 0) mean that the frequency of agent i is close (resp. far) to the average frequency of the group.

The net group synchronization index is computed in a similar way. Namely we let

$$\psi_{\mathrm{net}}(k) = \tan^{-1}\frac{\mathrm{Im}(q_{\mathrm{net}}(k))}{\mathrm{Re}(q_{\mathrm{net}}(k))} \in [0, 2\pi], \tag{13}$$

$$\phi_{\mathrm{net},i}(k) = \theta_i(k) - \psi_{\mathrm{net}}(k), \tag{14}$$

$$\bar{\phi}'_{\mathrm{net},i} = \frac{1}{T}\sum_{l=1}^{T} e^{\mathrm{j}\phi_{\mathrm{net},i}(l)}, \tag{15}$$

$$\bar{\phi}_{\mathrm{net},i} = \tan^{-1}\frac{\mathrm{Im}\left(\bar{\phi}'_{\mathrm{net},i}\right)}{\mathrm{Re}\left(\bar{\phi}'_{\mathrm{net},i}\right)} \in [0,2\pi], \tag{16}$$

$$\Delta\phi_{\mathrm{net},i}(k) = \phi_{\mathrm{net},i}(k) - \bar{\phi}_{\mathrm{net},i}, \tag{17}$$

and the *net group synchronization index* at time k, denoted by $\rho_{\mathrm{net}}(k) \in [0, 1]$, is given by

$$\rho_{\mathrm{net}}(k) = \frac{1}{N-1}\left|\sum_{\substack{i=1 \\ i\neq \mathrm{L3}}}^{N} e^{\mathrm{j}\left(\Delta\phi_{\mathrm{net,i}}(k)\right)} - e^{\mathrm{j}\left(\Delta\phi_{\mathrm{net,L3}}(k)\right)}\right|. \tag{18}$$

4.2.4.3 *INDIVIDUAL SYNCHRONIZATION INDEX*

To quantify the phase synchronization degree of the i-th participant with respect to the group, we define the *total* and *net individual synchronization index*, respectively, as

$$\rho_{\mathrm{tot},i} = \left|\bar{\phi}'_{\mathrm{tot},i}\right|, \qquad \rho_{\mathrm{net},i} = \left|\bar{\phi}'_{\mathrm{net},i}\right|. \tag{19}$$

Values close to 1 (resp. 0) mean that agent the phase of agent i is close (resp. far) to the average phase of the group.

4.2.5 FROM MOTION TO PHASE AND FROM PHASE TO MOTION

4.2.5.1 *FROM MOTION TO PHASE*

To estimate the phase of a motion, there are many different methods (e.g., [Varlet et al. 2011], [Pikovskij et al. 2003]). We will focus on two of them: the Hilbert transform, which is more precise but can be used only when the complete motion is already recorded, and an ad-hoc algorithm, which is in general less precise, but can be used online, when only limited portions of the motion are available.

According to *analytic signal theory* (Panter, 1965), for a one-dimensional signal $x(t)$, such as the motion of a finger when projected onto its main axis of motion, letting $x_{\mathrm{H}}(t)$ be the Hilbert transform of $x(t)$, the phase $\theta^{\mathrm{H}}(t) \in \mathbb{S}$ can be defined as the phase of the complex signal $x(t) + \mathrm{j}x_{\mathrm{H}}(t)$. When $x(t)$ is a signal with a narrow band (i.e., is composed mostly by one prevailing frequency), $\theta^{\mathrm{H}}(t)$ can be approximated to the latent phase (see § 4.2.2) of the signal. Additionally, the module of the complex signal $x(t) + \mathrm{j}x_{\mathrm{H}}(t)$ is the amplitude of the motion.

A different method consists in using Algorithm 1, which is run iteratively on a 1-dimensional discrete-time position signal $(p_1, p_2, p_3, \dots)$, and returns the estimated phase at the time instants, along with the estimated amplitude of motion at time k, say $A(k)$. The Algorithm assumes that the signal p alternates one maximum, one crossing of zero, one minimum, one crossing of zero, and so on. This fact is true if the signal has a narrow band and negligible noise. In fact, if the measured signal has a non-negligible amount of noise, it is necessary to first filter out high frequencies from the signal (this step will cause a delay in the estimation). If necessary, in future implementations of the CA, we might also consider extended versions of the algorithm that do not require this assumption. In Algorithm 1, the normalization at lines 5, 6 is done following (Mörtl et al., 2012).

Algorithm 1: Phase estimation from linear motion

Input: p_{t-1} (position at time $t-1$), p_t (position at time t), v_{t-1} (velocity at time $t-1$), v_t (velocity at time t), $A_{t-1}^{p\geq 0}$ (amplitude of position oscillation, when position is positive, at time $t-1$), $A_{t-1}^{p<0}$ (same, when position is negative), $A_{t-1}^{v\geq 0}$ (amplitude of velocity oscillation, when velocity is positive, at time $t-1$), $A_{t-1}^{v<0}$ (same, when velocity is negative).

Output: θ_t (estimated phase at time t), $A_t^{p\geq 0}$, $A_t^{p<0}$, $A_t^{v\geq 0}$, $A_t^{v<0}$.

1 **if** $p_{t-1} < 0$ *and* $p_t \geq 0$, **then** $A_t^{v\geq 0} = |v_t|$, **else** $A_t^{v\geq 0} = A_{t-1}^{v\geq 0}$;
2 **if** $p_{t-1} \geq 0$ *and* $p_t < 0$, **then** $A_t^{v<0} = |v_t|$, **else** $A_t^{v<0} = A_{t-1}^{v<0}$;
3 **if** $v_{t-1} \geq 0$ *and* $v_t < 0$, **then** $A_t^{p\geq 0} = |p_t|$, **else** $A_t^{p\geq 0} = A_{t-1}^{p\geq 0}$;
4 **if** $v_{t-1} < 0$ *and* $v_t \geq 0$, **then** $A_t^{p<0} = |p_t|$, **else** $A_t^{p<0} = A_{t-1}^{p<0}$;
5 **if** $p_t \geq 0$, **then** $p_t^{\text{norm}} = p_t / A_t^{p\geq 0}$, **else** $p_t^{\text{norm}} = p_t / A_t^{p<0}$;
6 **if** $v_t \geq 0$, **then** $v_t^{\text{norm}} = v_t / A_t^{v\geq 0}$, **else** $v_t^{\text{norm}} = v_t / A_t^{v<0}$;
7 $\theta_t = \text{atan2}(-v_t^{\text{norm}}, p_t^{\text{norm}})$;

Algorithm 1: Estimation of angular phases and amplitudes from linear motions.

4.2.5.2 FROM PHASE TO MOTION

To generate a one-dimensional motion $x(k)$ from a phase $\theta(k)$ it is also necessary to assume the knowledge of an amplitude $A(k)$. In particular, it suffices to take

$$x(\theta(k)) = A(k)\cos\theta(k). \tag{20}$$

Details on how we estimate $A(k)$ are reported in § 4.2.6.2.

4.2.6 DESCRIPTION OF THE COGNITIVE ARCHITECTURE

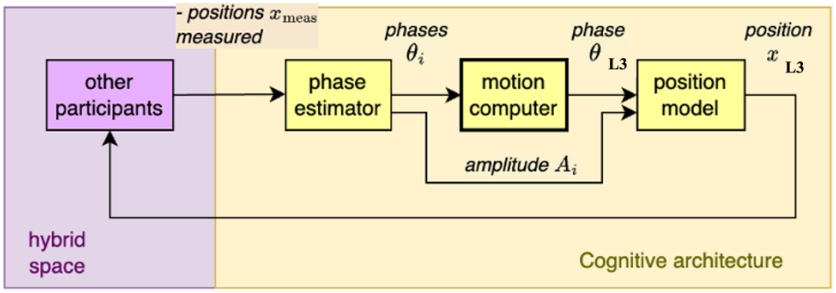

Figure 4: Block scheme of the proposed cognitive architecture for the PoP.

A schematic of the architecture for the L3 VH in the PoP of Social Connectedness is reported in Figure 4. In particular, the inputs to the architecture are the sensed linear motions $x_{i,\text{meas}}$ of the participants (as described in § 4.1), acquired from the hybrid space, while the output is the linear motion of the L3 generated by the CA. Next, we describe each of the blocks of the schematic in Figure 4.

4.2.6.1 "PHASE ESTIMATOR" BLOCK (FROM PHASE TO MOTION)

The "phase estimator" block takes as inputs the discretized motions $x_{i,\text{meas}}(k)\ \forall i$ of the participants and outputs the estimated phases $\theta_i(k)$ of all participants and the estimated amplitudes, say $A_i(k)$, of their motion, according to Algorithm 1.

4.2.6.2 "POSITION MODEL" BLOCK (FROM PHASE TO MOTION)

The phase $\theta_{\text{L3}}(k)$ generated by the "motion computer" block is transformed into a linear motion $x_{\text{L3}}(k)$ by the block "position model" with the transformation (see § 4.2.5.2)

$$x_{\text{L3}}(k) = A_{\text{L3}}(k)\cos\big(\theta_{L3}(k)\big), \tag{21}$$

where

$$A_{\text{L3}}(k) = \begin{cases} \text{mean}_{i\neq \text{L3}}\left(A_i^{p\geq 0}(k)\right), & \text{if } x_{\text{L3}}(k-1) < 0, \\ \text{mean}_{i\neq \text{L3}}\left(A_i^{p<0}(k)\right), & \text{otherwise}, \end{cases} \tag{22}$$

Where $A_i^{p\geq 0}(k)$ and $A_i^{p<0}(k)$ are obtained from Algorithm 1.

4.2.6.3 "MOTION COMPUTER" BLOCK

The "motion computer" in Figure 4 block takes as input the phases $\theta_i(k)$ of all participants and outputs the phase of the L3 VH, denoted by $\theta_{\text{L3}}(k)$.

The objective of the L3 CA is to maximize motor coordination between human participants. To do so, we let the L3 VH change its natural frequency $\omega_{\text{L3}}(k)$ in (3) and connect it to all VHs (i.e., $a_{i,\text{L3}} = a_{\text{L3},i} = 1, \forall i$). Hence, we define the *objective function* $J(k) = r_{\text{tot}}^2(k)$, and have the cognitive architecture solve the following optimization problem

$$\max_{\omega_{\mathrm{L3}}(k)} R = \sum_{k=1}^{\infty} J(k), \tag{23}$$

such that

$$\theta_i(k+1) = \theta_i(k) + \Delta t \left(\omega_i(k) + \frac{c}{N} \sum_{j \in \mathcal{N}_i} a_{ij} \sin\left(\theta_j(k) - \theta_i(k)\right) \right), \quad i \in \{1, \dots, N\}.$$

Addressing this nonlinear optimization problem through classical control theory can be challenging. Consequently, we propose an approach employing reinforcement learning to tune the natural frequency $\omega_{\mathrm{L3}}(k)$.

More specifically, Figure 5 is an inside view of the "motion computer" block, showing the details of how the phase $\theta_{\mathrm{L3}}(k)$ is generated. In particular, the estimated phases $\theta_i(k)$, the L3 phase $\theta_{\mathrm{L3}}(k)$ and the natural frequency $\omega_{\mathrm{L3}}(k)$ are the inputs to the "manipulation block", which has an aggregative function, as it outputs the vector

$$\begin{bmatrix} \mathrm{mean}\left(\theta_i(k) - \theta_{\mathrm{L3}}(k)\right) \\ \mathrm{var}\left(\theta_i(k) - \theta_{\mathrm{L3}}(k)\right) \\ \omega_{\mathrm{L3}}(k) \end{bmatrix}, \tag{24}$$

which is used, in turn, as input to the "DQN" block. Specifically, we use mean and variance to allow the CA to generalize to groups with different numbers of participants.

In the "DQN" block, a Deep Q Network (DQN) is implemented, which is a well-known reinforcement learning algorithm (Mnih et al., 2015). At each time step k, the algorithm outputs the variation of the natural frequency $\Delta\omega_{\mathrm{L3}}$ that maximizes the cumulative reward R in (*23*). $\Delta\omega_{\mathrm{L3}}$ is used as input to the "model of motion" block, in such a way that at each time step k, the natural frequency is updated as:

$$\omega_{\mathrm{L3}} \leftarrow \omega_{\mathrm{L3}} + \Delta\omega_{\mathrm{L3}}. \tag{25}$$

In summary, the *state* of the reinforcement learning agent is the vector (*24*), its *action* is $\Delta\omega_{\mathrm{L3}}(k)$, and its *reward* (objective function) is $J(k)$.

Finally, the "model of motion" block computes the output phase $\theta_{\mathrm{L3}}(k)$ of the L3 VH using (*3*).

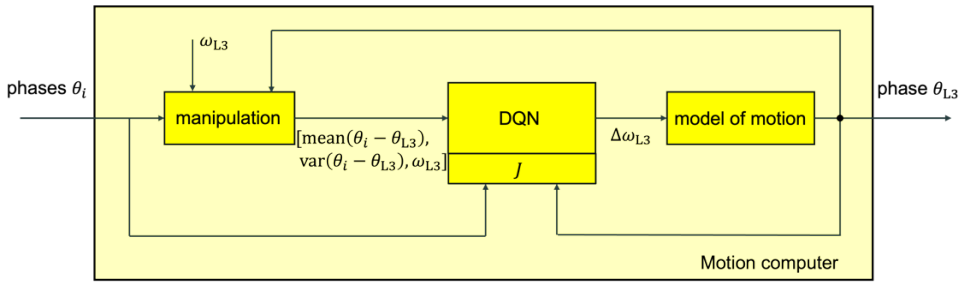

Figure 5: Schematic of the generation process of the ideal phase θ_{L3}.

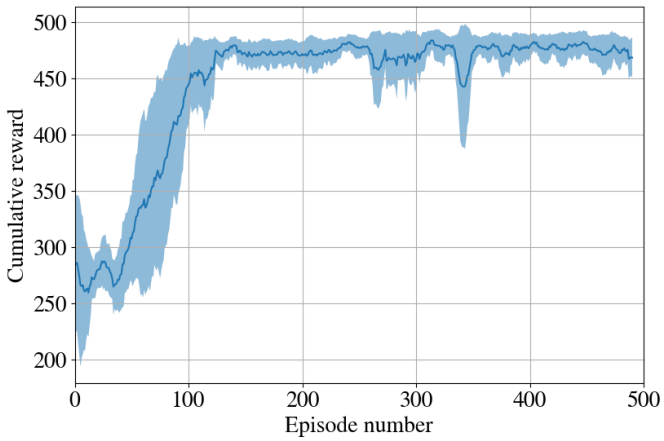

Figure 6: Cumulative reward R.

4.3 Training and validation of the L3 CA in the Proof of Principle of Social Connectedness

In reinforcement learning algorithms such as DQN, two distinct phases are identified: *training* and *validation*. In the training phase, the algorithm interacts with the environment, varying its behavior policy to seek the best one, whereas in the validation phase the learnt behavior is tested, possibly in conditions never exactly experienced during training, to assess the capabilities of the algorithm to generalize. The validation of the L3 CA for the PoP of Social Connectedness was carried out both numerically and experimentally. Below, we provide a detailed description of these two phases.

4.3.1 Description of the training phase

To train the "motion computer" block with the proposed architecture, we would need a huge amount of data, capturing participants' responses to different L3 actions under various conditions. Nonetheless, collecting this data is unfeasible, as it would necessitate thousands of experiments. Consequently, the CA is trained through interaction with synthetic agents, that is using a sim-to-real approach, which simulate human behavior utilizing the Kuramoto model in (*2*). First, we note that the input to the L3 DQN block in Figure 5 is chosen as mean and variance of the phases of

other VHs (rather than directly all the phases), allowing the trained CA to be connected to any number of neighbors, rather than just a number that was used during training. Furthermore, as only mean and variance are extrapolated from the input phases, we train the CA only in a configuration where it is connected to two other nodes.

Training is composed of a series of repeated simulations, called *episodes*, where the CA explores different policies, and gradually learn which is the strategy leading to the best cumulative reward (*23*) over time (Mnih et al. 2015, Sutton and Barto, 2018). To expose the CA to a broad spectrum of scenarios, both the natural frequencies and the initial conditions of the agents to whom the CA is connected to are varied in each episode. In particular, the natural frequencies are extracted, at each episode, from a uniform distribution in the interval $\left[\omega_{\mathrm{m}} - \frac{\Delta\omega_{\mathrm{m}}}{2}, \omega_{\mathrm{m}} + \frac{\Delta\omega_{\mathrm{m}}}{2}\right]$, where ω_{m} is the mean frequency of a group of people performing the oscillatory task in isolation; $\Delta\omega_{\mathrm{m}}$ and the initial conditions are chosen so as to allow the possibility of at least frequency synchronization to occur, according to what is reported in Bullo (2022). In particular, following the experimental results in Alderisio et al. (2017), we set $\omega_{\mathrm{m}} = 4$ rad/s and $c = 1.25$. According to Bullo (2022), to allow for frequency synchronization, it must hold that $\Delta\omega_{\mathrm{m}} < c$: hence, we set $\Delta\omega_{\mathrm{m}} = 1.2$; moreover, for the same reason the initial conditions are drawn uniformly at random in the interval $\left[\frac{\pi}{4}, \frac{3\pi}{4}\right]$.

We note that the state space has dimension equal to 3 (which is the size of the vector in (*24*)); the action space (from which $\Delta\omega_{\mathrm{L3}}$ is chosen) is set to $\{-0.5, -0.4, \dots, 0.5\}$. The NN architecture in the DQN was chosen heuristically and is composed of one input layer with 3 nodes, one output layer with 11 nodes and two hidden layers of 128 and 64 nodes, respectively. The number of training episodes is set to 500; the simulation time in each episode is T = 5 time units, with a step size of $\Delta t = 0.01$. Training is then repeated in 3 independent sessions to evaluate the dispersion of the training process while in numerical and experimental validation only the policy resulting from the first learning session is deployed.

Figure 6 shows the results of the training phase in terms of the cumulative reward R in each episode.

4.3.2 NUMERICAL VALIDATION

4.3.2.1 COMPARISON OF GRAPH TOPOLOGIES

To restrict the number of trials to be carried out in the numerical and experimental validations, we need to select the kind of interaction pattern between group members to analyze. As a guiding criterion, we seek a topology where (i) permutations of the order of the participants do not qualitatively affect synchronization (because it will not be feasible to test for all possible permutations) and (ii) the graph topology does not show almost perfect synchronization in the absence of the L3 (so that there is margin for improvement, in order to motivate the use of L3).

In particular, we compare numerous simulations of Kuramoto networks with 5 agents, in all-to-all, path, ring, and star graphs topologies. For each graph topology, we permutate randomly the ordering of the agents in the graph 50 times. For each simulation, we compute the mean in time of the order parameter r_{tot} and the group synchronization index ρ_{tot} (see § 4.2.4; note that since in this case there is no L3, $r_{\mathrm{tot}} = r_{\mathrm{net}}$ and $\rho_{\mathrm{tot}} = \rho_{\mathrm{net}}$). The values of the natural frequencies used for the agents are drawn uniformly at random from the same interval used during the training phase. Moreover, to assess

the margin for improvement present in each case, and as preliminary evidence of the capability of the L3 agent to improve synchronization, we perform the same simulations also in the case that an L3 agent is added to the network, and connected to all the nodes; in that case, we compute the mean in time of the net order parameter r_{net} and the net group synchronization index ρ_{net} (see § 4.2.4).

We report the results of these simulations in Figure 7. We notice that predictably in the case of the all-to-all graph, no variability is observed, as the graph is invariant to permutation of the vertices. This would in principle be ideal for the validation; however, in this case the network already achieves a very good level of synchronization, leaving practically no margin for improvement, and hence is unsuitable to validate the L3 CA. Conversely, the path and star topologies leave significantly more margin for improvement of synchronization, but they also display too large variability (see the vertical size of the box plots) when the ordering of the VHs is changed. A good compromise is found in the ring topology, which shows rather small variability together with sufficient margin for improvement.

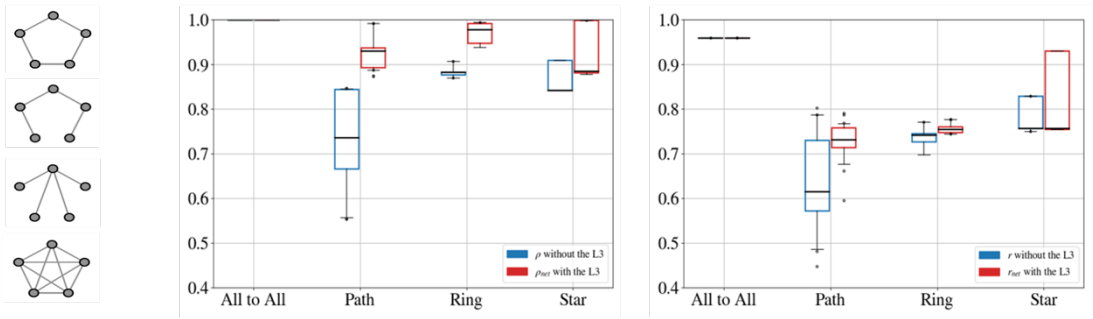

Figure 7: Box plots to evaluate the interconnection topology to use in the experiments. On the left, the four interaction topologies tested: from top to bottom, the ring, path, star, and all-to-all graphs. In the center, the box plots related to ρ_{net}. On the right, the box plots related to r_{net}.

4.3.2.2 *RESEARCH QUESTIONS AND EXPERIMENTAL CONDITIONS*

The aim of the numerical and experimental validation is to provide an answer to the following research questions:

Q1. Can an L3 agent added to a group improve its synchronization?
Q2. When replacing a human participant in the group, can an L3 agent improve synchronization?
Q3. Do L3 agents behave in a human-like manner? (this could be seen as a simple example of a motor Turing test).

To answer these questions, we designed the following experiments:

- E1 (“No L3”). In this scenario the group is composed only of human participants or simulated people (according to the Kuramoto model). This condition will function as a baseline to be compared with the conditions where an L3 VH is present.
- E2 (“L3 added to the group”). In this case, an L3 VH is added to a group of real/simulated people. The purpose of this scenario is to answer to Q1.
- E3 (“L3 replacing a person”). In this scenario, an L3 VH replaces a person in the group, in order to leave the graph topology unaltered. This case can be used to provide an answer to Q2.

4.3.2.3 NUMERICAL SETUP

We validate numerically the cognitive architecture for L3 avatars by simulating the experimental conditions E1, E2, E3 in § 4.3.2.2. We assume groups of 7 agents for E1, and to simulate the motion of people, we use oscillators with time-varying natural frequencies. In particular, following (Alderisio et al., 2017), at each time step, the natural frequency of agent i, say $\omega_i(k)$, is extracted from a Gaussian distribution identified from real data; we will consider 2 sets of Gaussian distributions, termed as *group 1* and *group 2*. The initial conditions are set to $\theta_i(0) = \pi/2, \forall i$.

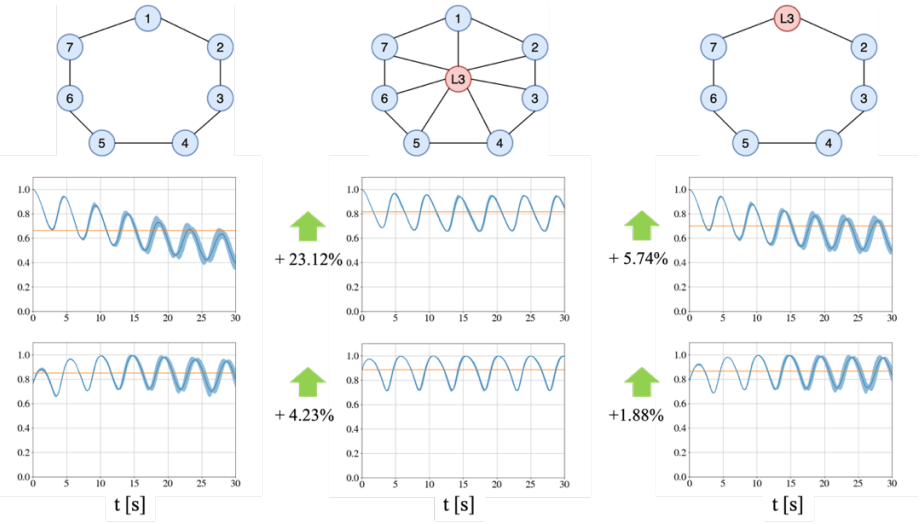

Figure 8: Numerical results of the three experiments for group 1 in terms of ρ_{net}, ρ_{tot}, r_{net} and r_{tot}. The top row reports the interaction topologies used. Panels in the middle row report in blue the order parameter r_{tot} , r_{net} and r_{tot} respectively over time and in orange the mean value of these quantities. Panels in the bottom row depict in blue the group synchronization index ρ_{tot}, ρ_{net} and ρ_{tot} respectively over time and in orange the mean values of these quantities. Panels on the left, middle and right portray the results in conditions E1, E2, E3, respectively.

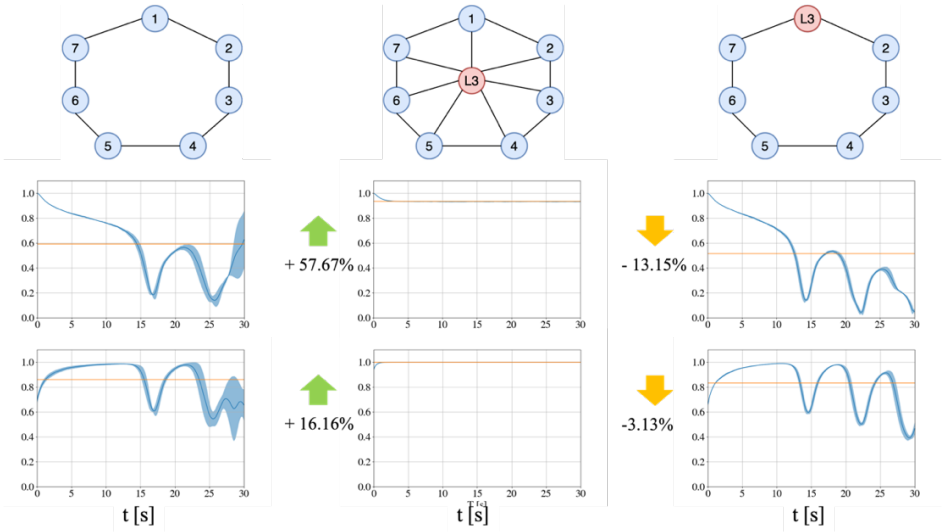

Figure 9: Numerical results of the three experiments for group 2 in terms of ρ_{net}, ρ_{tot}, r_{net} and r_{tot}. The top row reports the interaction topologies used. Panels in the middle row report in blue the order parameter r_{tot} , r_{net} and r_{tot} respectively over time and in orange the mean value of these quantities. Panels in the bottom row depict in blue the group synchronization

index ρ_{tot}, ρ_{net} and ρ_{tot} respectively over time and in orange the mean values of these quantities. Panels on the left, middle and right portray the results in conditions E1, E2, E3, respectively.

4.3.2.4 RESULTS OF NUMERICAL VALIDATION

In Figure 8, we show the results of the three experimental conditions E1, E2, E3, each repeated 5 times (each time the realization of the natural frequencies of the simulated participants is drawn randomly). In group 1, by comparing the results of E2 with those of E1, we see a significant improvement in synchronization, especially concerning the metrics r_{net}. In E3, we selected a random participant to be replaced by an L3 VH. When comparing the results of E3 with those in E1, we still see a minor increase in the mean values of the designated metrics. The enhancement in synchronization demonstrates the advantageous effect of employing L3 agents within the experimental framework, especially when an L3 VH is added to the ensemble. The metrics obtained are also reported in Table 3. Concerning group 2, again a significant benefit is obtained when a L3 VH is added to the group, although synchronization performance is not improved when the L3 VH replaces one of the simulated participants.

	E1 (tot)	E2 (net)		E3 (tot)	
$\text{mean}_k(r(k))$	0.662	0.815	+ 23.12%	0.700	+ 5.74%
$\text{std}_k(r(k))$	0.141	0.102	- 27.65%	0.121	- 14.18%
$\text{mean}_k(\rho(k))$	0.850	0.886	+ 4.23%	0.866	+ 1.88%
$\text{std}_k(\rho(k))$	0.092	0.095	+ 3.26%	0.090	- 2.17%
ρ_{L3}	\	0.995		\	
ρ_1	0.938	0.978	+ 4.26%	0.972	+ 3.62%
ρ_2	0.986	0.987	+ 0.01%	0.990	+ 0.40%
ρ_3	0.946	0.981	+ 3.69%	0.961	+ 1.58%
ρ_4	0.966	0.986	+ 2.07%	0.981	+ 1.55%
ρ_5	0.971	0.986	+ 1.54%	0.985	+ 1.44%
ρ_6	0.174	0.228	+ 31.03%	0.175	+ 0.57%
ρ_7	0.901	0.972	+ 7.88%	0.926	+ 2.77%

	E1 (tot)	E2 (net)		E3 (tot)	
$\text{mean}_k(r(k))$	0.593	0.935	+ 57.67%	0.515	- 13.15%
$\text{std}_k(r(k))$	0.238	0.009	- 96.21%	0.275	+ 15.54%
$\text{mean}_k(\rho(k))$	0.860	0.999	+ 16.16%	0.833	- 3.13%
$\text{std}_k(\rho(k))$	0.139	0.005	- 96.40%	0.169	+ 21.58%
ρ_{L3}	\	1		\	
ρ_1	0.936	0.999	+ 6.73%	0.807	- 13.78%
ρ_2	0.930	0.999	+ 7.41%	0.812	- 12.68%
ρ_3	0.944	1	+ 5.93%	0.875	- 7.30%
ρ_4	0.683	0.999	+ 46.26%	0.733	+ 7.32%
ρ_5	0.627	0.999	+ 59.33%	0.585	- 6.69%
ρ_6	0.757	0.998	+ 31.83%	0.699	- 7.66%
ρ_7	0.954	1	+ 4.82%	0.794	- 16.77%

Table 3: Summary tables of the metric computed during the numerical experiments. On the left, the table related to group 1, On the right, the table related to group 2. The percentages refer to comparisons with the appropriate values in E1.

4.3.2.5 VERIFICATION OF OPTIMALITY OF THE L3 CA STRATEGY

Next, we aim to verify that L3 CA actually selects the optimal natural frequency for the L3 VH. To assess this, we compare the net order parameter r_{net} obtained by adding to the group of simulated participants an L3 VH with non-adaptive fixed natural frequency with the value of r_{net} obtained by using the RL-based trained L3 CA. The range of frequencies analyzed for the L3 with fixed frequency is $[3, 6]$ rad/s. We report the results of this analysis in Figure 10, for group 1, where we portray the mean and standard deviation value of r_{net} for each fixed value of natural frequency of the L3 agent. We observe the existence of an optimal natural frequency for synchronization in terms of r_{net} (≈ 4.5 rad/s); the L3 agent based on RL is able to achieve a very similar value (≈ 4.2 rad/s).

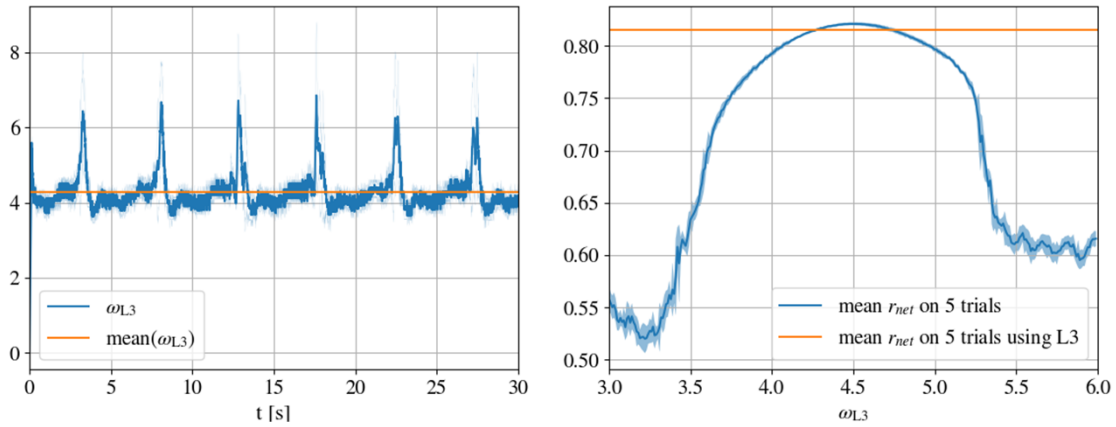

Figure 10: On the left, the natural frequency ω_{L3} of L3 (blue) and its mean in time (orange). On the right, mean with standard deviation of r_{net} obtained with an L3 with fixed frequency ω_{L3} (blue) and mean of r_{net} obtained with the RL-based L3 (orange).

Results obtained for group 2 are similar and thus omitted for the sake of brevity.

4.3.3 EXPERIMENTAL VALIDATION - CHRONOS

Figure 11: On the left, a photo of the experiment. On the right, a screenshot of the Chronos interface.

4.3.3.1 DESCRIPTION OF THE SETUP

To validate our strategy with real people, we extended and employed a pre-existing software called "Chronos" (Alderisio et al., 2017b) (https://dibernardogroup.github.io/Chronos/index.html). This platform allows to carry out experiments involving multiple human participants, with the possibility of also inserting artificial agents. Specifically, each participant is visually and acoustically isolated from the other participants and has in front of them a computer on which the software is installed; see Figure 11. Each person maneuvers a ball shown on the screen with a mouse or a "leap motion" tracking device (Potter et al., 2013), and sees the balls of some other participants, depending on the interaction structure decided before the experiment (see again Figure 11). All participants are instructed to move in a smooth oscillatory fashion and to synchronize their motion with the others. One of the balls might be driven by the L3 CA but this is unknown to the participants.

We carried out experiments E1, E2, E3 (see § 4.3.2.2) with 2 groups of 5 people each. Each experiment was repeated 5 times, and each iteration lasted 30 s. Furthermore, to provide an answer to Q3 in §

4.3.2.2, at the end of each trial, we asked the participants whether they believed they had interacted with an autonomous agent or not, and, if so, which ball (displayed on their screen) was the one belonging to the L3 VH.

The position data of the players was collected at a rate of 10 Hz. This data was then interpolated using a piecewise cubic Hermite interpolating polynomial (PCHIP) interpolation to increase the sampling rate to 100 Hz. Next, we applied a Butterworth filter, with a cutoff frequency set at 6 Hz, which is twice the approximate standard frequency associated to natural human motion, that is 3 Hz (Calabrese et al., 2022). The Hilbert transform (Kralemann et al., 2008) was used to reconstruct the phase associated to the motion of each agent from the time series of its position. Finally, in E3, the participant to be substituted was again selected randomly.

4.3.3.2 RESULTS

We show the results of the experiments for the two groups in Figure 12 and Figure 13, respectively. The data suggest that L3 is able to improve synchronization, as measured by the mean values of r_{net} and ρ_{net}. The only exception is contained in E3 for group 2, where L3 is not able to improve synchronization: this might also be due to the fact that the in that case (see E1 for group 2) synchronization was already high. The values of the metrics obtained in the experiments are reported in Table 4.

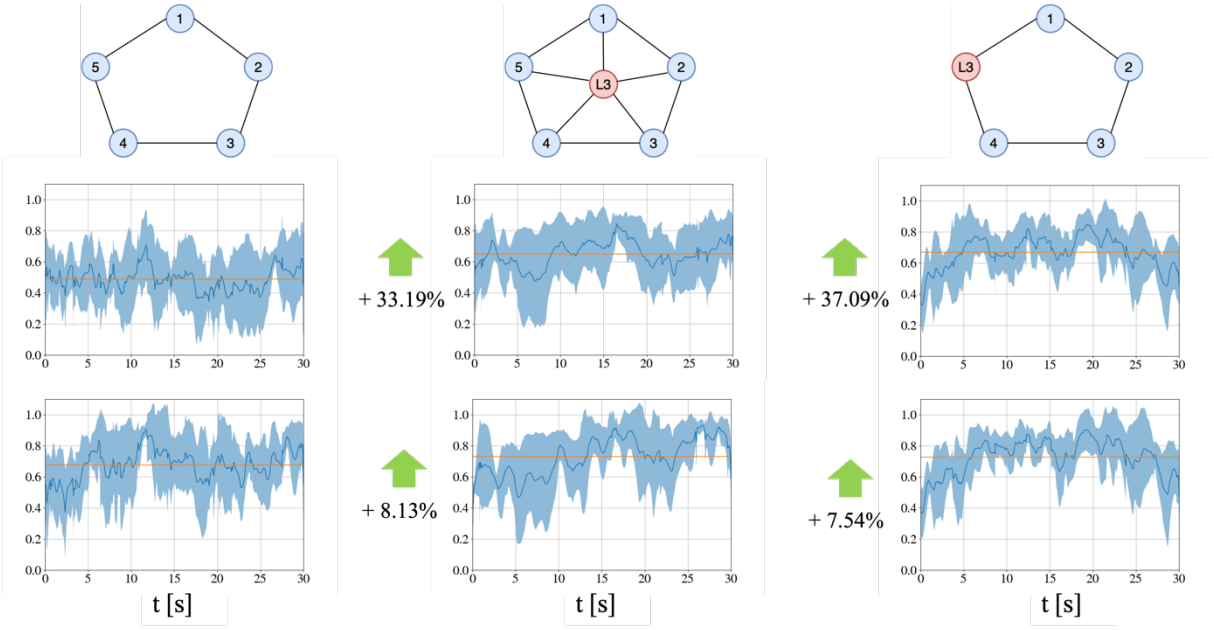

Figure 12: Results of the three real experiments for group 1 in terms of ρ_{net}, ρ_{tot}, r_{net} and r_{tot}. The top row reports the used interaction topologies. Panels in the middle row report in blue the order parameter r_{tot} , r_{net} and r_{tot} respectively over time and in orange the mean value of these quantities. Panels in the bottom row depict in blue the group synchronization index ρ_{tot}, ρ_{net} and ρ_{tot} respectively over time and in orange the mean values of these quantities. Panels on the left, middle and right portray the results in conditions E1, E2, E3, respectively.

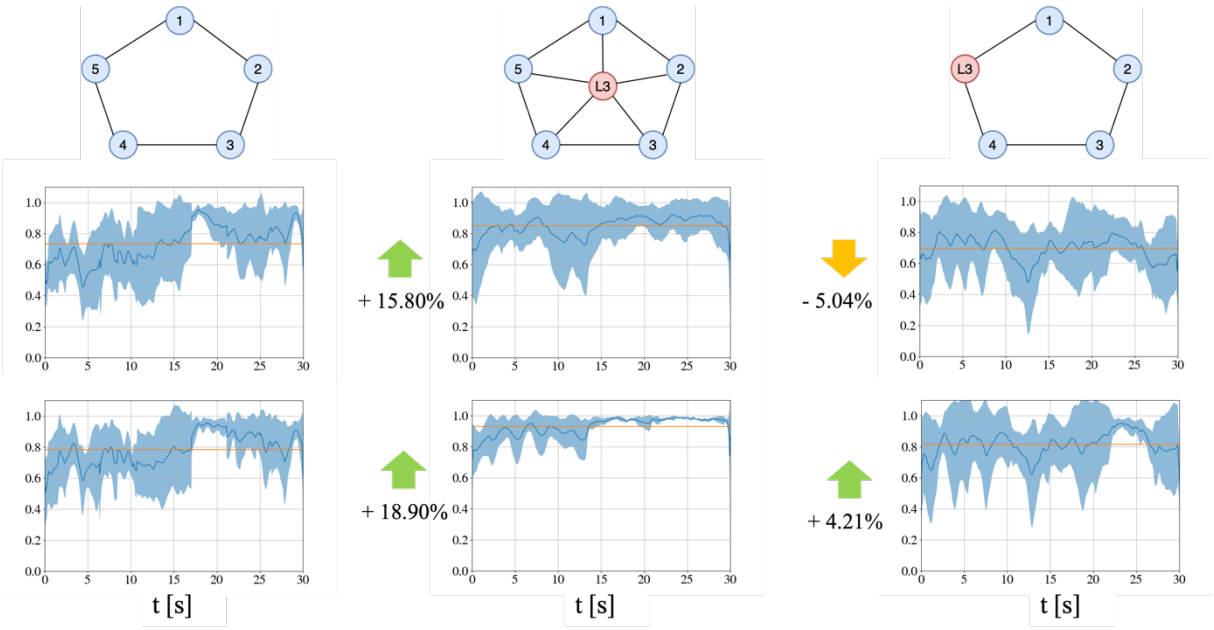

Figure 13: Results of the three real experiments for group 2 in terms of ρ_{net}, ρ_{tot}, r_{net} and r_{tot}. The top row reports the used interaction topologies. Panels in the middle row report in blue the order parameter r_{tot} , r_{net} and r_{tot} respectively over time and in orange the mean value of these quantities. Panels in the bottom row depict in blue the group synchronization index ρ_{tot}, ρ_{net} and ρ_{tot} respectively over time and in orange the mean values of these quantities. Panels on the left, middle and right portray the results in conditions E1, E2, E3, respectively.

	E1 (tot)	E2 (net)		E3 (tot)	
$\text{mean}_k(r(k))$	0.488	0.650	+ 33.19%	0.669	+ 37.09%
$\text{std}_k(r(k))$	0.069	0.079	+ 14.49%	0.099	+ 43.47%
$\text{mean}_k(\rho(k))$	0.676	0.731	+ 8.13%	0.727	+ 7.54%
$\text{std}_k(\rho(k))$	0.095	0.118	+ 24.21%	0.111	+ 16.84%
ρ_{L3}	\	0.717		\	
ρ_1	0.615	0.722	+ 17.39%	0.743	+ 20.18%
ρ_2	0.643	0.796	+ 23.79%	0.851	+ 32.34%
ρ_3	0.639	0.674	+ 5.47%	0.670	+ 4.85%
ρ_4	0.637	0.709	+ 11.30%	0.577	- 9.41%
ρ_5	0.559	0.620	+ 10.91%	0.765	+ 36.85%

	E1 (tot)	E2 (net)		E3 (tot)	
$\text{mean}_k(r(k))$	0.734	0.850	+ 15.80%	0.697	- 5.04%
$\text{std}_k(r(k))$	0.111	0.057	- 48.64%	0.071	+ 36.03%
$\text{mean}_k(\rho(k))$	0.783	0.931	+ 18.90%	0.816	+ 4.21%
$\text{std}_k(\rho(k))$	0.098	0.055	- 43.87%	0.071	- 27.55%
ρ_{L3}	\	0.921		\	
ρ_1	0.795	0.927	+ 16.60%	0.897	+ 12.83%
ρ_2	0.803	0.932	+ 16.06%	0.780	- 2.86%
ρ_3	0.853	0.942	+ 10.43%	0.719	+ 15.70%
ρ_4	0.664	0.940	+ 41.56%	0.841	+ 26.65%
ρ_5	0.741	0.846	+ 13.36%	0.875	+ 18.08%

Table 4: Summary tables of the metrics computed during the real experiments. On the left, the table related to group 1, On the right, the table related to group 2.

To assess the capability of the L3 VH to improve synchronization, we conducted a Wilcoxon signed-rank test to validate the hypotheses that the order parameter r_{net} and the group synchronization index ρ_{net} are larger when the L3 is used, both when adding a L3 VH to the group (thus comparing E2 to E1) and when the L3 VH is replacing a person (comparing E3 to E1). For the first comparison (E2 vs. E1), concerning the hypotheses that r_{net} and ρ_{net} improve, respectively, we obtained p-values of 0.062 and 0.219 for group 1, and 0.223 and 0.003 for group 2. For the second comparison (E3 vs. E1), we obtained p-values of 0.031 and 0.15 for group 1, and 0.7 and 0.71 for group 2.

These results are very preliminary as the dataset is not large enough to conduct a proper statistical test; however, they provide early indication that L3 improves coordination when added to the group,

and when replacing a person in the group. In this latter case, we see that the L3 VH is able to improve coordination in group 1, but not in group 2, pointing out that selecting the right person to replace is an important factor.

In Table 5, we report the average percentage of correct answers obtained in each group, when participants were asked to identify L3 (without specifying if it was present or not).[2] The baseline percentages in the table are computed by considering that in E1 and E3 each person is connected to 2 agents (hence, three possible answers, including the absence of L3), while in E2 each person is connected to 3 agents (thus, 4 possible answers). In experimental condition E1, group 1 obtained a low percentage of correctness, while group 2 did better. In E2, both groups did not do better than chance, while in E3 both groups did better than chance. Hence, in conclusion, we have three instances where L3 was not well recognized as an artificial agent, and three where it was recognized more than by guessing at random.

Experiment	Baseline (random choice)	Group 1	Group 2
E1 (“No L3”)	33%	10 %	55 %
E2 (“L3 added to the group”)	25%	25 %	25 %
E3 (“L3 replacing a person”)	33%	57 %	54 %

Table 5: Percentages of right answers to the questions of identifying the presence of the L3 agent.

[2] Data for Player 2 is unavailable due to a software error during data collection; hence, the percentages are computed excluding player 2’s answers.

4.3.4 EXPERIMENTAL VALIDATION - CRANKS

Figure 14: On the top left image, users are engaged in the third test. On the right the users are involved in the first test. On the bottom left, a single user and a motorized node.

4.3.4.1 DESCRIPTION OF THE SETUP

The proposed strategy was further validated through a different practical setup involving real people. In this platform, illustrated in Figure 14, users are tasked to smoothly manipulating a crank–to which a stick is attached–from left to right and vice versa, while attempting to synchronize with the sticks they see. Specific intercommunication topologies are obtained by obstructing sight of the participant through plastic barriers. Additionally, the system allows the integration of artificial agents through nodes without cranks, but with a stick, operated by stepper motors. During the trials, users wear passive noise-canceling headphones.

The platform comprises three main components:

- the *central unit*, (i.e., an Arduino Due);
- the *user nodes*, consisting of cranks with visible vertical sticks, equipped with an angular positioning sensor, and that can be operated by human participants;
- the *motorized nodes*, which are sticks driven by stepper motors. The nodes are linked to the central unit using standard ten-way IDC connectors, enabling the arrangement of different configurations.

The central unit is responsible for controlling the dynamics of the motorized nodes and acquiring the state of the human nodes, at a rate of 200 Hz. Participants are instructed to operate the crank by performing complete semi-rotations between 0 and 180 degrees, aiming to execute the movement smoothly. Data from all the trajectories are saved on a personal computer, at a rate of 100 Hz.

For each participant, the angular position of their crank is mapped to phases in $\mathbb{S}$ as follows. While moving in the clockwise direction, angular positions in $[0, \pi]$ are mapped to $[0, \pi]$; conversely, when moving counterclockwise, angular positions in $[\pi, 0]$ are mapped to $[-\pi, 0]$. Consequently, each phase is in $(-\pi, \pi]$.

With this platform, we carried out 2 experiments involving one group of 5 people. Each experiment was repeated 4 times with a duration of 30 s each. The objective of the experiments is to determine whether the autonomous agent can enhance synchronization when acting as an intermediary between a group of people and a reference motion they must follow (e.g., as could happen in a gym or rehabilitation class). In the first experiment, people can see the reference directly and are instructed to follow it. In the second one, people see an L3 agent, which is connected to all people and to the reference motion. The reference motion is created beforehand in MATLAB, and is the output of a harmonic oscillator with a time-variant angular velocity, being equal to $\omega_{\mathrm{ref}} = 5.5 + \sin(2\pi w(k)k)$, where $w(t)$ is changed every second, being drawn from a uniform probability distribution between 2 and 5. To assess how well the group is synchronized in frequency with respect to the reference motion, we define the two further metrics

$$\rho_{\mathrm{tot}}^{\mathrm{ref}}(k) \coloneqq \frac{1}{N}\left|\sum_{i=1}^{N} e^{\mathrm{j}\left(\phi_{\mathrm{tot},i}(k)-\mathrm{ref}(k)\right)}\right| \in [0,1], \tag{26}$$

$$\rho_{\mathrm{net}}^{\mathrm{ref}}(k) \coloneqq \frac{1}{N-1}\left|\sum_{i=1}^{N} e^{\mathrm{j}\left(\phi_{\mathrm{net},i}(k)-\mathrm{ref}(k)\right)} - e^{\mathrm{j}\left(\phi_{\mathrm{net,L3}}(k)-\mathrm{ref}(k)\right)}\right| \in [0,1]. \tag{27}$$

4.3.4.2 RESULTS

We report the results of the experiments in **Fehler! Verweisquelle konnte nicht gefunden werden.**. Remarkably, the L3 CA is able to increase both group synchronization (r_{net} and ρ_{net}, which exclude both L3 and the reference) and the capability of the group to follow the reference ($\rho_{\mathrm{net}}^{\mathrm{ref}}$).

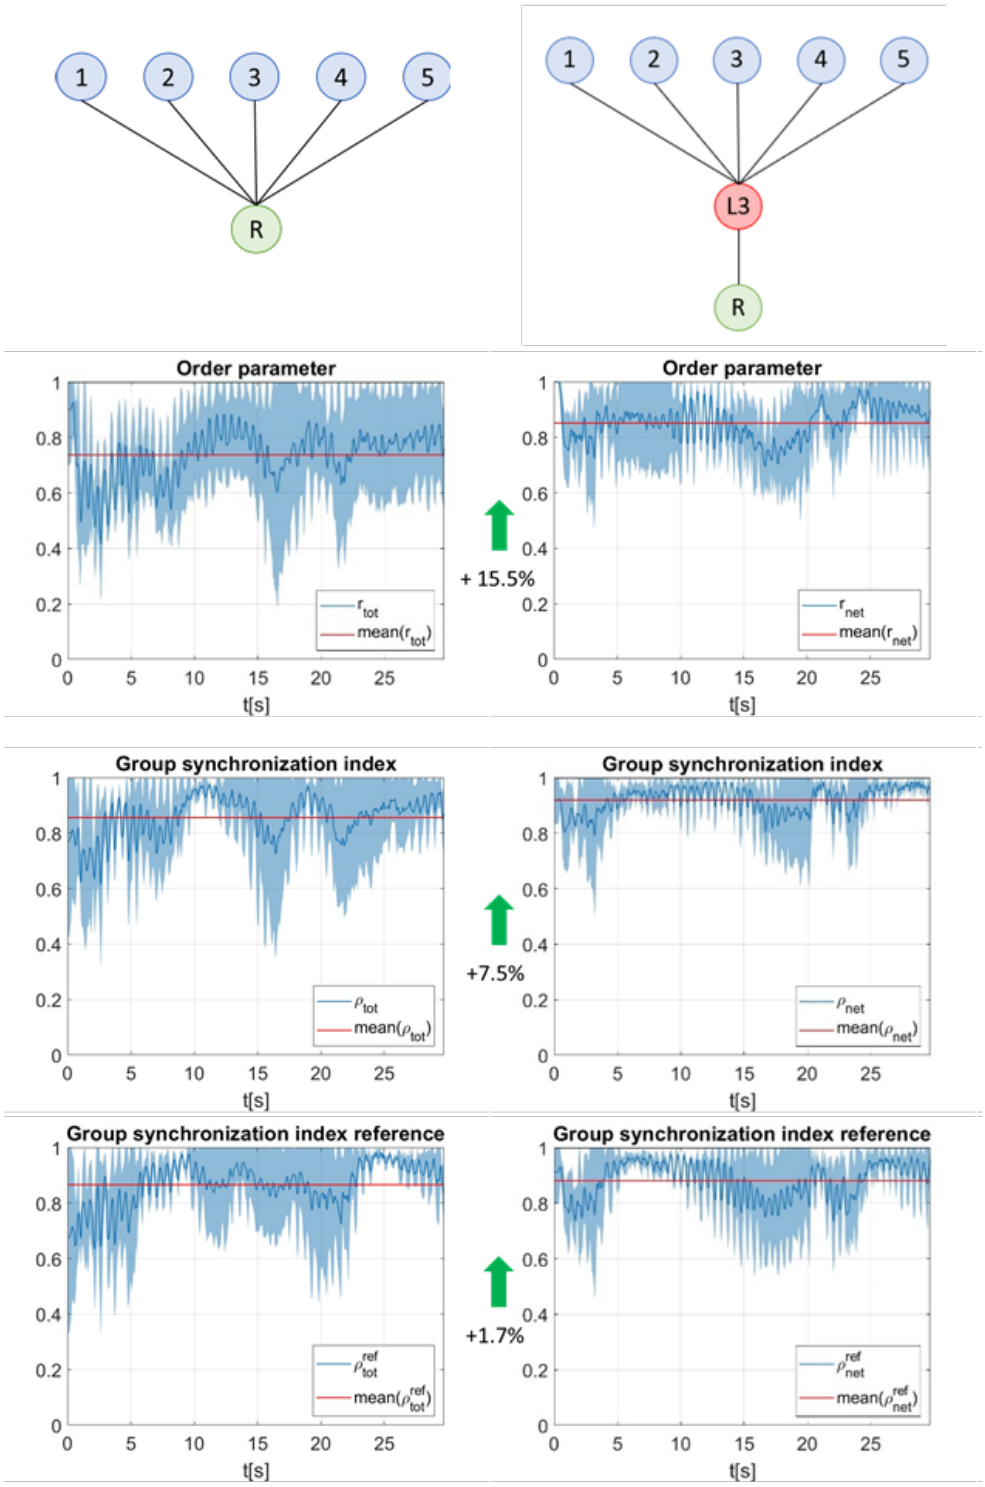

Figure 15: Results of the two real experiments in terms of ρ_{net} ,ρ_{tot}, r_{net},r_{tot} ,ρ_{net}^{ref} and ρ_{tot}^{ref}. On the top row, the interaction topologies used in the experiments, with node R being driven by the reference motion ref(k). In the second row, the results in terms of r_{net} and r_{tot}. On the third row, the results in terms of ρ_{net} and ρ_{tot} of the group of humans. In the bottom row, the results in terms of ρ_{net}^{ref} and ρ_{tot}^{ref} of the group of humans relative to the reference signal ref(k). Panels on the left (resp. right) are related to the experimental setup portrayed in the top left (resp. right) panel.

5 L3 CA IN THE PROOFS OF PRINCIPLE OF AMPLIFICATION

5.1 ROLE AND REQUIREMENTS OF THE L3 CA IN THE PROOF OF PRINCIPLE OF AMPLIFICATION

In this Proof of Principle, participants are asked to pass an object to each other to study how information propagates through body kinematics. An in-depth description of this Proof of Principle is provided in deliverable D1.2. However, in this document, we consider the case where some of the virtual humans in the chain pass the object and transmit information (e.g., fear) encoded in their kinematics (see definition 5.11.4 of deliverable D1.1). In this context, the objective of the L3 CA is to start the chain with different types of information encoded in the movement.

The objective of these experiments is to check whether the L3 architecture can influence the human group by starting the kinematic chain to start a flow of information transfer among the participants.

5.2 IMPLEMENTATION OF L3 CA IN THE PROOF OF PRINCIPLE OF AMPLIFICATION

For this Proof of Principle, the L3 Cognitive Architecture leverages a *movement primitives library* as that of the L2 CA described in deliverable D5.1 and whose structure is depicted in Figure 17. The library serves as a database of examples of speed profiles that *kinematically encode* the designated information and speed profiles that do not. In what follows, in accordance with D1.2, we assume that the encoded information is fear. In general, the library required by the L3 CA is a collection of signals that define the movement primitives. In such a collection of examples, each sample is a set of signals (e.g. position, velocity, acceleration etc.; see D2.1) that are labeled to identify the following properties:

- the type of motion (e.g., “reach to grasp”, “pass the object”, “return to rest”).
- the presence of socially relevant encoded information (e.g., fear of an object).

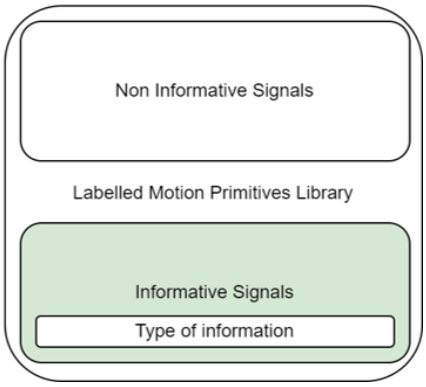

Figure 16: Block scheme of the movement primitives library.

Starting from such a database of movement primitives, the L3 CA can generate movements with some target encoded information. Specifically, in so doing, the L3 CA, is based on a schematic similar to the one previously reported for the L2 Cognitive Architecture as described in deliverable D5.1., In particular, we can use the blending coefficients (obtained as a result of the L2 CA training) to blend

two samples from the movement primitives library so as to obtain new generated movements which encode the reference information.

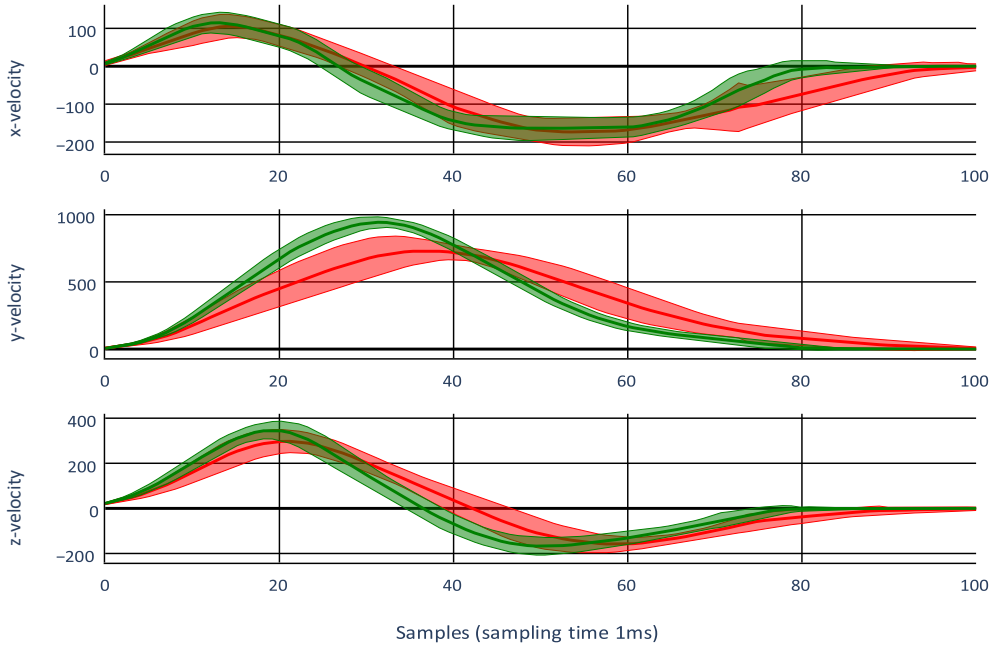

Figure 17: Speed profiles projection on x, y, and z planes of a test user performing reach to grasp movement primitive with fear (red) and without fear (green). Measurements refer to the average wrist marker in multiple sessions. The solid lines represent the average, while the shaded areas correspond to two times the standard deviation.

Let us consider the reach-to-grasp movement primitive as an illustrative example, wherein the L3 VH is tasked with moving the hand to pick up an object. The initiation of this process involves acquiring a set of samples that encapsulate crucial information derived from the movement primitives library, as illustrated in Figure 18. In this context, the L3 cognitive architecture is responsible for generating the end-effector speed trajectory with reference-encoded information. This process is schematically depicted in Figure 19, and, in particular, we can summarize the key functionalities of the cognitive architecture as follows:

- *Task-specific requirements*: In this block, we specify the requirements of the motion in terms of a sequence of movement primitives and reference-encoded information. This piece of information is used to start the kinematic chain with the intent of transmitting the reference-encoded information to the other human participants. This block is also responsible for checking the current execution of the L3 VH in terms of its end-effector position to determine the correct movement primitive in the specified sequence.
- *Movement generation*: This block is responsible for generating a new motion profile for the L3 end-effector. The motion generation process is trained offline using the movement

primitives library in the same way as described in D5.1 and uses two velocity profiles to be blended to encode reference information as specified in the task-specific requirements.

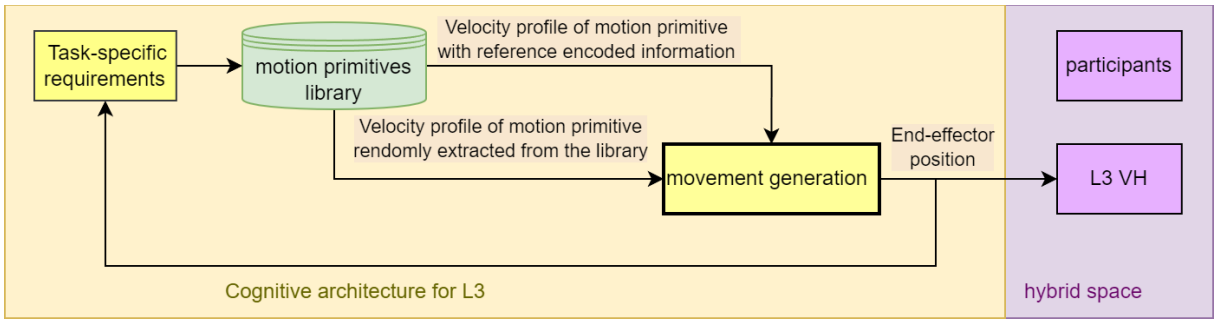

Figure 18: Block scheme of L3 cognitive architecture in PoP of Amplification

For simplicity, considering the sequence of a single movement primitive (e.g. reach-to-grasp), the L3 CA begins by selecting two samples from the library. The former encodes the reference information and a second is a movement of the same type which we can extract randomly from the library. Both samples are then used in the blending process described in § 5.1 of D5.1.

However, depending on the requirements of the experiments, the L3 CA could also use leaning algorithms to select the two samples from the library to be used in the blending process. As a matter of fact, as described in D5.1, the movement generator is trained to encode information in the blended signal. So, the choice of the samples should not impact on fulfilments of the task specific requirements in terms of encoded information. However, as pointed in D2.4, each human participant *reads out* information in different ways and, in this sense, the combination of samples used for the blending process could impact on the readability of the encoded information. In future iterations, the L3 CA could use *reinforcement learning from human feedback* (RLHF) to select the two samples from the *movement primitives library* and further generate motion with the intent of improve read out of information by individual human participants (MacGlashan et al., 2017).

In essence, the L3 involves a dynamic process to maximize its efficiency in replicating human-like trajectories with reference encoded information and individual read out properties. This process allows to leverage the properties of the trained *movement generator* to encode information in a human movement while giving to the L3 avatar a plethora of possible alternatives to prevent L3 VH from replaying the same movements.

The implementation and validation of the L3 CA in this PoP is currently under development as it relies on the availability of an experimentally validated movement primitives library currently being constructed at UKE. The results will be presented in the next iteration of this deliverable due on month 19.

6 INTEGRATION WITH THE SCENARIOS

In this section, we describe how the CAs presented in Sections 4 and 5 could be applied to drive L3 agents in the application scenarios described in deliverable D1.2. A more specific description will be presented in the next deliverables.

6.1 THE L3 COGNITIVE ARCHITECTURE IN SHARESPACE FOR HEALTH

Let us start by recalling the scenario "Sharespace for health: social low back pain exergame". In this scenario, patients coping with chronic low back pain, will participate in physical therapy sessions in a hybrid space, on the SHARESPACE platform, either from their home or in a hospital room, guided by their therapist. Additional non-human patients might be represented by L3 VHs, with the aim to help group interaction and foster a sense of shared goal. Given the importance of coordinated movement, in this scenario it is possible to use the CA presented in Section 4 (that is related to the PoP on Social Connectedness) to drive the L3 VHs.

In this scenario, patients move many parts (if not all) of their body, rather than mainly just a single arm (as in the PoP on Social Connectedness). Therefore, it will be needed to adapt the "phase estimator" block and the "position model" blocks introduced in Section 0, to transform movements in phases and vice-versa. In particular, for each physical exercise, we make the following assumptions:

- We have access to the shape of the correct exercise, say a curve $\Gamma \in \mathbb{R}^{3d}$, where d is the number of captured points in space.
- We have access to a trajectory representing the correct execution of the movement, say $m_{\mathrm{r}}(t) : [0, T] \to \Gamma$, T being the duration of the exercise.
- There are no points on the trajectory that have the same position and velocity (i.e., $\nexists\, t_1, t_2$ such that $t_1 \neq t_2$ and $m_{\mathrm{r}}(t_1) = m_{\mathrm{r}}(t_2)$ and $\dot{m}_{\mathrm{r}}(t_1) = \dot{m}_{\mathrm{r}}(t_2)$).

Then, we define the phase of the movement as the result of a transformation ($f_\theta : (m_{\mathrm{r}}, \dot{m}_{\mathrm{r}}) \mapsto \theta$) that takes in input the position and velocity of the correct movement, and outputs the current phase ($f_\theta : \mathbb{R}^{6d} \to \mathbb{S}$), with the requirements that $\frac{\mathrm{d}\theta(m_{\mathrm{r}}(t), \dot{m}_{\mathrm{r}}(t))}{\mathrm{d}t} = \frac{2\pi}{T} = \omega$, i.e., the phase grows linearly between time 0 and time T. Note that this phase might be understood as the measure of completion of the movement within the exercise, transforming a complex spatial motion into a simple one-dimensional signal.

Furthermore, to obtain the transformation that yields the (latent) phase of any movement (that is close to the correct one), we will have to define a projection transformation $P : (m, \dot{m}) \mapsto (m_{\mathrm{r}}, \dot{m}_{\mathrm{r}})$ that projects any configuration of the body into the closest body configuration on the correct exercise execution. A simple implementation of P could be $P(m) = \arg\min_{\mu \in \Gamma} \|\mu - m\|$; however more sophisticated implementations will be tested if needed (for example, using the velocity $\dot{m}$ to disambiguate between different points in the trajectory). Hence, given a generic motion with position and velocity $(m, \dot{m})$, its latent phase is given as $\theta = f_\theta(P(m, \dot{m}))$.

Once a phase is available for all VHs, the CA can compute the phase of the L3 VH it is driving. After that, to obtain a body configuration for the L3 VH from its phase, a further transformation is needed

that maps phases to sets of points in the three-dimensional space. This transformation can either be obtained from the inversion of the correct trajectory m_r, or be restricted to part of it, computing only the position of a few important points (e.g., the back in an exercise to mitigate low back pain), and the rest of the points might be decided in a different way, e.g., to exhibit a specific style, while still ensuring consistency of the skeleton of the VH.

6.2 THE L3 COGNITIVE ARCHITECTURE IN SHARESPACE FOR SPORT

Next, we consider the scenario "SHARESPACE for Sport: Peloton Cycling". In that case, the situation can occur that three bikers are riding in a line, each aiming to win a race. Say B1 the biker in front of the group, B2 the one in the middle, and B3 the one behind. At some point, B2 starts a so-called *attack*, i.e., it tries to pass B1. Then, it is in the best interest of B3 to immediately detach from the line and follow B2 in his attack, so to keep drafting it.[3] This situation is recreated to train athletes or enthusiasts in a hybrid space on the SHARESPACE platform. In particular, B2 might be an L3 VH, while B3 is a real person, represented by a L1 VH. The role of the L3 VH (here, B2) is to emphasize its movement just before attacking, in order to train the biker B3 in rapidly responding to the attack. Since in this situation one of the most important aspects is the transmission of information, we plan to use the CA described in Section 5 (that is related to the PoP of Amplification) to drive the motion of the L3 VH in this scenario.

To carry out the task, the L3 CA needs to have access to an exhaustive movement primitive library, that includes both motion that clearly telegraphs the intention to attack, and motion that does so much more subtly. Let us remark that the presence of this information is akin to the information of fear being present or not in the PoP of Amplification. Then, to generate original motion, the L3 will generate a new trajectory by blending together multiple instances of motion that together have the desired level of informativity concerning the display of telling features of the attack. To help the learning biker (B3, in the above description) the CA might progressively reduce the amount of informativity of its motion.

In the L3 VH's motion generation process, the motion of the virtual biker consists of two major parts: that from the waist up, and that from the waist down. The former might be more related to the transmission of information (although this is yet to be verified in the context of the Research Project) and the latter influences the position of the bike on the road. In principle, the L3 CA will use both parts of motion to generate the L3 VH's motion, although it might be necessary to detach the two to ease training. In particular, we will consider the possibility that showing the attack motion on the trunk first and then on the legs (which also modifies the position of the bike) makes it easier to teach the attack pattern to the learning biker.

6.3 THE L3 COGNITIVE ARCHITECTURE IN SHARESPACE FOR ART

Finally, we consider the scenario "SHARESPACE for Art: Shared Creativity"; this will see the realization of artworks designed by artists recruited by the project and by the Ars Electronica team. Artworks

[3] A bike (or vehicle) is said to *draft* another when both are moving in line, and the second is close behind the first. This reduces the amount of drag experienced by the bike on the back.

might feature a combination of interactions of motion, music, and light to co-create a new aesthetic form. Some of the participants will be present in the Deep Space 8K arena, while others will be located remotely; all of them will interact through the SHARESPACE platform. The artworks might involve different situations and goals: then, the role and behavior of the CA will depend on the specific artworks, which are yet to be finalized. Still, two main themes will be the coordination of motion and the contagion of social information. Both the implementations of the L3 CA presented in this document can be used to achieve these goals. Specifically, the CA presented in Section 4 is better suited to drive VHs that coordinate their motion with that of human VHs, while the CA presented in Section 5 is better used to ease or hamper the transmission of social information, to create interesting social and movement patterns.

7 CONCLUSIONS AND FUTURE WORK

In this deliverable, we introduced the first implementation and validation of the AI-driven Cognitive Architectures (CA) to drive fully autonomous L3 agents in the SHARESPACE project. After discussing the state of the art, we recalled the meaning of the levels of autonomy and outlined the main general assumptions on the L3 CA. Then, we presented the implementations for the L3 CA to be used in the PoP of Social Connectedness and in the PoP of Amplification. The former focuses on improving the level of motion coordination, whereas the latter focuses on passing or blocking certain information via motor clues. In particular, we explained how the motion of L3 for the first PoP is generated using the combination of a Kuramoto model and a Deep Q-Networks algorithm. Conversely the motion of an L3 VH for the second PoP is generated by blending signals of the movement primitives library. We also reported extensive numerical validation and experimental validation with two setups for the CA concerning the PoP of Social Connectedness. Current results show the ability of the CA to improve synchronization when it is added to a group of people, and its ability to maintain coordination when it is used to replace a human participant. When people were asked to identify L3 VHs in a group, results were mixed.

The next steps of the project concerning the development of the L3 CA are the following:

- Experimental validation of the L3 CA for the PoP of Social Connectedness through the trials to be carried out as described in D1.2.
- Experimental validation of the L3 CA for the PoP of Amplification. A crucial step is to finalize the CA to assess if the current architecture can effectively encode information in movements. To do so, in collaboration with WP2 ("Sensorimotor primitives of social interaction") and WP4 ("Rendering"), we will assess the capabilities of the CA using pre-recorded movement.
- Integration of the CA with the System Architecture, involving a detailed analysis of the interaction of CA's functionalities with the architectural components to ensure seamless communication and data exchange.
- Online experimental validation on the SHARESPACE platform. Namely, we will integrate the CA with the System Architecture and the multi-stream platform (currently under development in task T5.3). These results will inform the refinement of the CA for the three scenarios of the SHARESPACE project.

REFERENCES

N. Ahmad, R. A. R. Ghazilla, N. M. Khairi, and V. Kasi, “Reviews on various inertial measurement unit (IMU) sensor applications”, International Journal of Signal Processing Systems, vol. 1, no. 2, pp. 256–262, 2013.

F. Alderisio, G. Fiore, R. N. Salesse, B. G. Bardy, and M. di Bernardo, “Interaction patterns and individual dynamics shape the way we move in synchrony”, Scientific Reports, vol. 7, no. 1, 2017.

F. Alderisio, M. Lombardi, G. Fiore, and Mario di Bernardo, “A novel computer-based set-up to study movement coordination in human ensembles”, Frontiers in psychology, vol. 8, pp. 967, 2017b.

D. Babajanyan, G. Patil, M. Lamb, R. W. Kallen, and M. J. Richardson, “I know your next move: Action decisions in dyadic pick and place tasks”, in 44th Annual Meeting of the Cognitive Science Society: Cognitive Diversity, pp. 563–570, 2022.

B. G. Bardy et al., “Moving in unison after perceptual interruption”, Scientific Reports, vol. 10, no. 1, Art. no. 1, 2020.

M. M. N. Bieńkiewicz, S. Janaqi, P. Jean, and B. G. Bardy, “Impact of emotion-laden acoustic stimuli on group synchronisation performance”, Scientific Reports, vol. 13, no. 1, p. 7094, 2023.

F. Bullo, “Lectures on network systems”, ed. 1.6. Kindle Direct Publishing, 2022.

C. Calabrese, M. Lombardi, E. Bollt, P. De Lellis, B. G. Bardy, and M. di Bernardo, “Spontaneous emergence of leadership patterns drives synchronization in complex human networks”, Scientific Reports, vol. 11, no. 1, p. 18379, 2021.

C. Calabrese, B. G. Bardy, P. De Lellis, and M. di Bernardo, “Modeling frequency reduction in human groups performing a joint oscillatory task”, Frontiers in Psychology, vol. 12, 2022.

H. Haken, J. A. S. Kelso, and H. Bunz, “A theoretical model of phase transitions in human hand movements”, Biological Cybernetics, vol. 51, no. 5, pp. 347–356, 1985.

B. Kralemann, L. Cimponeriu, M. Rosenblum, A. Pikovsky, and R. Mrowka, “Phase dynamics of coupled oscillators reconstructed from data”, Physical Review E vol. 77, iss. 8, 2008.

M. Lombardi, D. Liuzza, and M. di Bernardo, “Deep learning control of artificial avatars in group coordination tasks”, in Proc. of 2019 IEEE International Conference on Systems, Man and Cybernetics (SMC), pp. 714–719, 2019.

M. Lombardi, D. Liuzza, and M. di Bernardo, “Dynamic input deep learning control of artificial avatars in a multi-agent joint motor task”, Frontiers in Robotics and AI, vol. 8, 2021a.

M. Lombardi, D. Liuzza, and M. di Bernardo, “Using learning to control artificial avatars in human motor coordination tasks”, IEEE Transactions on Robotics, vol. 37, no. 6, pp. 2067–2082, 2021b.

K. R. McKee et al., “Scaffolding cooperation in human groups with deep reinforcement learning”, Nature Human Behaviour, vol. 7, no. 10, Art. no. 10, 2023.

V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, “Human-level control through deep reinforcement learning”, Nature, vol. 518, no. 7540, pp. 529–533, 2015.

A. Mörtl, T. Lorenz, B. N. S. Vlaskamp, A. Gusrialdi, A. Schubö, and S. Hirche, “Modeling inter-human movement coordination: synchronization governs joint task dynamics”, Biolocial Cybernetics, vol. 106, no. 4, pp. 241–259, 2012.

L. Noy, E. Dekel, and U. Alon, “The mirror game as a paradigm for studying the dynamics of two people improvising motion together”, Proceedings of the National Academy of Sciences, vol. 108, no. 52, pp. 20947–20952, 2011.

P. Panter, “Modulation, noise, and spectral analysis”. McGraw-Hill, New York, 1965.

G. Patil, P. Nalepka, H. F. Stenning, R. W. Kallen, and M. J. Richardson, “Scaffolding deep reinforcement learning agents using dynamical perceptual-motor primitives”, Proceedings of the 45th Annual Conference of the Cognitive Science Society, 2023.

A. Pikovskij, M. Rosenblum, and J. Kurths, “Synchronization: A universal concept in nonlinear sciences”, Cambridge: Cambridge Univ. Press, 2003.

L. E. Potter, J. Araullo, L. Carter, “The Leap Motion controller: a view on sign language”, OzCHI '13: Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration, pp. 175–178, 2013.

E. Scaliti, K. Pullar, G. Borghini, A. Cavallo, S. Panzeri, and C. Becchio, “Kinematic priming of action predictions”, Current Biology, 2023.

S. Shahal et al., “Synchronization of complex human networks”, Nature Communications, vol. 11, no. 1, Art. no. 1, 2020.

S. H. Strogatz, Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering, Second edition. CRC Press, 2015.

R. Sutton, A. Burto, “Reinforcement learning: An introduction”, The MIT Press, Second edition, (2018).

G. Turri, A. Cavallo, L. Romeo, M. Pontil, A. Sanfey, S. Panzeri, and C. Becchio, “Decoding social decisions from movement kinematics”, iScience, vol. 25, no. 12, p. 105550, 2022.

C. Zhai, F. Alderisio, K. Tsaneva-Atanasova, and M. di Bernardo, “A novel cognitive architecture for a human-like virtual player in the mirror game”, in 2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 754–759, 2014.

C. Zhai, F. Alderisio, P. Słowiński, K. Tsaneva-Atanasova, and M. di Bernardo, “Design of a virtual player for joint improvisation with humans in the mirror game”, PLOS ONE, vol. 11, no. 4, p. e0154361, 2016.

C. Zhai, F. Alderisio, P. Słowiński, K. Tsaneva-Atanasova, and M. di Bernardo, “Design and Validation of a Virtual Player for Studying Interpersonal Coordination in the Mirror Game”, IEEE Transactions on Cybernetics, vol. 48, no. 3, pp. 1018–1029, 2018a.

C. Zhai, M. Z. Q. Chen, F. Alderisio, A. Yu. Uteshev, and M. di Bernardo, “An interactive control architecture for interpersonal coordination in mirror game”, Control Engineering Practice, vol. 80, pp. 36–48, 2018b

M. Varlet and M. J. Richardson, “Computation of continuous relative phase and modulation of frequency of human movement”, Journal of Biomechanics, vol. 44, no. 6, pp. 1200–1204, 2011.

J. MacGlashan, M. K. Ho, R. Loftin, B. Peng, G. Wang, D. L. Roberts, M. E. Taylor, and M. L. Littman. "Interactive learning from policy-dependent human feedback.", International conference on machine learning, Proceedings of Machine Learning Research, pp. 2285-2294, 2017.

Y. Kuramoto, “Lecture notes in physics”, International Symposium on Mathematical Problems in Theoretical Physics”, Springer-Verlag, vol. 39, p. 420, 1975.