Dynamic Neural Field-Based Anticipatory Action Selection for Human–Robot Collaboration: A Virtual Reality Experiment

João G. Cunha[1][0000−0003−3392−224X], Estela Bicho[1][0000−0003−2715−4037], Wolfram Erlhagen[2][0000−0002−8150−3587], and Raymond H. Cuijpers[3][0000−0001−5980−9208]

[1] Universidade do Minho, Centre Algoritmi, Guimarães, Portugal
id10787@uminho.pt, estela.bicho@dei.uminho.pt
[2] Universidade do Minho, Centre of Mathematics, Guimarães, Portugal
wolfram.erlhagen@math.uminho.pt
[3] Eindhoven University of Technology, Human-Technology Interaction Group, Eindhoven, Netherlands r.h.cuijpers@tue.nl

Abstract. Neurally inspired mechanisms for anticipatory action selection offer a promising route toward more fluent Human–Robot Collaboration (HRC). In this work, we present a bio-inspired Dynamic Neural Field (DNF) cognitive architecture that generates complementary robot actions through a continuously evolving representation of human intent. The architecture was evaluated in a collaborative Virtual Reality (VR) pick-and-place task, enabling controlled examination of anticipatory coordination dynamics. In a within-subjects experiment, participants interacted with anticipatory and non-anticipatory versions of the robot. Anticipation significantly increased perceived competence, warmth, predictability, and collaboration, while reducing discomfort. Qualitative feedback further described the anticipatory robot as more adaptive and cooperative. Objective measures showed a significant reduction in collisions, indicating smoother division of labour, whereas task duration and idle time did not differ significantly between conditions. These results suggest that neurally inspired DNF architectures can give rise to anticipatory behaviours perceived as more predictable, and fluent in human–robot teamwork, supporting their use as neurally grounded mechanisms for complementary coordination in HRC.

Keywords: Dynamic Neural Fields · Cognitive Robotics · Anticipatory Action Selection · Joint Action · Human–Robot Collaboration · Virtual Reality

1 Introduction

The success of Human-Robot Collaboration (HRC) hinges on the robot's ability to integrate smoothly into human workflows. Central to this integration is the incorporation of anticipatory mechanisms that allow robots to predict and adapt to human actions [12]. Anticipatory behaviours enable robots to act proactively rather than reactively, leading to more fluent interaction, fewer coordination conflicts, and greater user trust [33,14,9].

The concept of anticipatory action selection is deeply rooted in cognitive science and motor control research. Action preparation is understood as a continuous process that leverages prior knowledge and predictive abilities to maximise utility before decisions are made [26]. Neuroscientific insights indicate that joint-action mechanisms rely on shared representations, predictive coding, and motor simulation to enable effective coordination between agents [34,3]. These processes allow us to anticipate others' actions, a capability that can be extended to robotic systems to enhance collaboration [15,16]. Additionally, joint action research emphasises that accurate predictions regarding the 'what,' 'when,' and 'where' of human actions improve coordination and communication, which is crucial for HRC [4].

In robotics, anticipation has been approached through gaze tracking [18], probabilistic planning [38], motion prediction [20], goal inference [23,21], failure anticipation [19], and explicit intent communication [29,35]. These approaches have demonstrated improvements in efficiency, safety, and perceived intelligence [17,24]. However, many existing implementations rely on discrete rule-based or probabilistic formulations that treat intention inference and action selection as sequential or stage-based processes.

Dynamic Neural Fields (DNFs) offer an alternative by modelling anticipation as a continuous dynamical process over neural population activity. In DNF architectures, evolving activation profiles integrate perceptual inputs and contextual constraints over time, producing interpretable intermediate representations that remain stable under noise and uncertainty [32,13,31]. This dynamical perspective provides a neurally grounded account of prediction and decision-making and has previously been applied to goal inference and joint action in robotics [4,11]. Unlike discrete planning, purely probabilistic models, or simple threshold rules, DNFs couple perception and action selection within the same continuous decision dynamics, making them well-suited for modelling online complementary behaviour in HRC.

Virtual Reality (VR) offers a complementary experimental context for such dynamical architectures. It provides continuous streams of behavioural and environmental data (e.g., hand trajectories, object states) that can be naturally mapped onto neural field dynamics. Additionally, VR enables controlled manipulation of interaction conditions and reduces physical safety constraints that may otherwise influence coordination behaviour in HRC [36,7,27,25,28]. This setting provides a well-suited test-bed to isolate and examine the specific effects of anticipatory behaviour within our collaborative task.

In this study, we present a DNF-based cognitive architecture that generates complementary robot actions by continuously inferring human intent from ongoing movement trajectories. The architecture is evaluated in a VR pick-and-place task, comparing anticipatory and non-anticipatory robot behaviour. The contribution of this work is threefold: (i) the implementation of a neurally grounded anticipatory architecture that continuously couples human intent inference with robot action selection to enable complementary behaviour, (ii) its deployment within a controlled VR collaboration setting, and (iii) an exploratory user evaluation of its effects on perceived collaboration and coordination dynamics.

Our findings indicate that participants perceived the anticipatory robot as more competent, predictable, and collaborative, while reporting less discomfort. Anticipation substantially reduced collisions, although task duration and idle time did not differ significantly between conditions. Together, these results suggest that DNF-based decision dynamics can give rise to anticipatory behaviours perceived as more fluent in human–robot teamwork. The remainder of the paper details our methodology (Section 2), results (Section 3), and conclusions (Section 4).

2 Methodology

2.1 Participants

Ten participants (5 female, 5 male, Mage = 27.30, SDage = 5.29) were recruited through Eindhoven University of Technology's participant database. An *a priori* power analysis indicated that with a within-subjects design, this sample was sufficient to detect large effects (Cohen's $d > 0.90$) with 80% power. Given the exploratory aim of the study and the focus on within-subject comparisons, we targeted the detection of robust perceptual effects while minimising between-subject differences.

All participants had normal or corrected-to-normal vision. No prior experience with VR or the Panda Robot was required. Still, participants drawn from this population typically have above-average familiarity with such technologies due to coursework and research activities. Participants were asked to collaborate with a robotic system in a pick-and-place task. Each session lasted approximately 30 minutes, and participants received €6 compensation.

2.2 Design

We employed a within-subjects design with two conditions: anticipatory and non-anticipatory robot behaviour. The primary independent variable was whether the robot predicted the participant's next action. In the anticipatory condition, the robot adapted its actions based on inferred human intent, ensuring complementary rather than competing choices. Whereas, in the non-anticipatory condition, the human-intention representation layer was disabled in the architecture, so the robot selected from the remaining available objects without considering the participant's intent. In both conditions, if the participant grasped an object that the robot was also targeting before the robot reached it, the robot retreated and re-selected from the remaining objects. All other robot behaviours (motion planning and grasping) were identical across conditions, ensuring that any observed differences could be attributed to anticipation.

A within-subjects design was chosen to control for individual variability, with each participant serving as their own control. This approach is particularly appropriate given the small sample size, as it eliminates between-subject variance in response styles. The order of conditions was counterbalanced: participants with even IDs experienced the non-anticipatory condition first, and those with odd IDs began with the anticipatory condition. No significant order effects were observed.

Each participant completed two sets of 20 trials (one per condition). In every trial, the participant and robot worked simultaneously to place three objects into a container, each agent independently selecting and placing objects until none remained. This simple pick-and-place task allowed us to isolate the specific contribution of anticipation without interfering factors from more complex task logistics (e.g., sequencing, tool use). After completing each condition, participants filled out two questionnaires: the RoSAS, measuring perceived competence, warmth, and discomfort, and a custom Likert-scale instrument assessing perceived collaboration. Following the final set, participants also answered open-ended questions to provide qualitative feedback on their experience.

2.3 Experimental setup

The experiment took place in a VR environment built in CoppeliaSim, a widely used robotics simulator [30]. The scene replicated the layout of the VR Laboratory at Eindhoven University of Technology, consisting of a table (50 cm along the axis where objects were placed) with three objects, a container, and a 3D model of the Franka Emika Panda robot positioned opposite the participant (see Figure 1).

Participants interacted using an HTC Vive headset and handheld controller, grasping objects by pressing and holding a button and releasing it to place them. The VR interface was developed with the CoppeliaSim VR Toolbox [6], and a custom C++ application coordinated real-time interactions between the participant, the VR environment, and the robot. This integration ensured synchronous, low-latency interactions between human and robot, enabling reliable measurement of both subjective and objective collaboration outcomes.

2.4 Robot control architecture

Model overview The cognitive architecture employed in this work is based on the framework proposed by Bicho *et al.* [4] and further developed in subsequent studies [5,37]. It employs coupled DNFs to map observed human movements onto complementary robot actions in a context-dependent manner. For the present study, the architecture was designed to focus on its anticipatory capabilities.

The key function of the architecture is to infer the participant's intended action and to generate robot behaviours that complement, rather than duplicate, those actions. This reflects core mechanisms of human joint action, in which partners continuously monitor each other's movements, predict goals, and adapt accordingly [3].

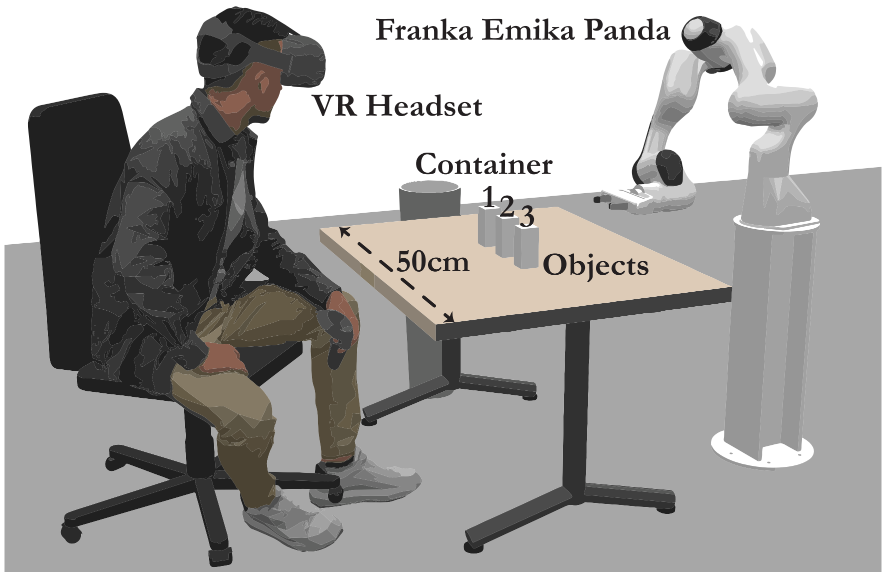

Fig. 1: VR experimental setup. Participants interacted with a 3D model of the Franka Emika Panda robot in a pick-and-place task. Using an HTC Vive headset and controller, participants collaborated with the robot to place three objects into a shared container. A video demonstration is available https://youtu.be/oyNoqGCwlLg.

As illustrated in Figure 2, the architecture consists of four neural fields: i) the **Action Observation Layer** (AOL) encodes the human hand trajectory, mapped from the HTC Vive controller onto a one-dimensional abstraction of the workspace (50 cm wide table), supporting continuous perceptual monitoring of the partner's movements; ii) the **Object Representation Layer** (ORL) encodes the positions of all objects present in the workspace; iii) the **Action Simulation Layer** (ASL) integrates information from AOL and ORL to infer which object the human intends to act upon, generating an interpretable intermediate activation representation, and mimicking the predictive simulation mechanisms associated with joint action coordination. Supra-threshold ($u_{ASL} > 0$) activation occurs only when both the object's presence and the participant's reaching motion align; and iv) the **Action Execution Layer** (AEL), which selects the robot's target grasp position, from which the robot's reaching, grasping, and placing actions are generated. Importantly, the AEL is inhibited by the ASL. That is, once the hand trajectory provides sufficient evidence for a specific target, the corresponding robot action plan is suppressed, and the system commits to an alternative object. This biasing occurs continuously during the human's movement, before human–object contact, enabling the robot to gradually accumulate evidence for the intended target and re-plan smoothly if the human changes direction mid-reach.

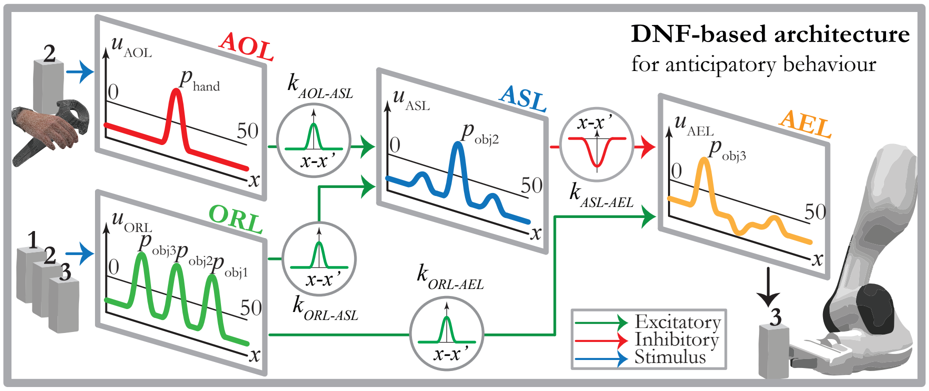

Fig. 2: Cognitive architecture for anticipatory collaboration. The system maps the participant's observed hand trajectory (AOL) and the positions of objects in the workspace (ORL) onto an internal simulation of the participant's intended action (ASL). This inferred intention modulates the robot's action execution layer (AEL), which selects a target grasp position for the robot. A suppression mechanism prevents the robot from selecting the same object as the human, ensuring complementary rather than competing behaviour. The activation profiles illustrate a scenario where the participant is reaching toward object 2, leading the AEL to select object 3.

This inhibition–excitation dynamic ensures that the robot does not compete for the same object as the human, but instead coordinates its behaviour to achieve efficient division of labour. In this way, the architecture operationalises neurally inspired mechanisms of prediction and joint action within a collaborative VR task. In the non-anticipatory condition, this predictive coupling was disabled; the robot could still resolve conflicts, but only reactively, *i.e.*, retreating and re-selecting after the participant grasped the contested object.

Model equations In DNFs, neural populations are represented as continuous activation fields, where each point in the field corresponds to the activation level of a specific group of neurons. These activation fields evolve based on interactions between neighbouring populations, which can be excitatory or inhibitory. The following equation describes the population dynamics in each field [1]:

$$\tau \dot{u}(x,t) = -u(x,t) + h + s(x,t) + \int k_{uu}(x-x')g(u(x',t))dx' + q\xi(x,t) \quad (1)$$

where $u(x,t)$ represents the average membrane potential at time t at position x, see implementation details in [13]. $h(<0)$ is the resting potential (a measure of the baseline activation level), $s(x,t)$ represents time-dependent localised input at site x from an external source or interconnected fields. The constant τ defines the time scale of the field dynamics. The output function $g(u)$ dictates that only neurons surpassing a predetermined threshold of activation contribute to the interaction (sigmoid function). The interaction kernel $k_{uu}(x, x')$ defines the coupling strength between any two neural positions in the field. ξ is added Gaussian white noise, with strength q, which models neuronal variability and facilitates selection [32]. Parameters were tuned to ensure stable activation dynamics and consistent complementary selection across trials. Time constants were adjusted to match human movement speed, interaction kernels were tuned for stable activation dynamics, and noise strength was calibrated to enable selection. All neural fields had a feature dimension of 50, corresponding directly to the 50 cm table where the objects were distributed. The spatial discretisation step was set to $d_x = 0.5$.

The different layers of the architecture apply this general equation with self-excitation interaction kernels $k_{uu}(x-x')$ tailored to their function. The ORL uses a Mexican-hat kernel, with intentionally weak lateral inhibition, which allows for multiple bumps to co-exist simultaneously (corresponding to the multiple objects):

$$k_{uu}(x - x') = A_{exc} \exp\left(-\frac{|x - x'|^2}{\sigma_{exc}^2}\right) - A_{inh} \exp\left(-\frac{|x - x'|^2}{\sigma_{\text{inh}}^2}\right) \tag{2}$$

where A_{exc} defines the amplitude of the excitatory Gaussian component, σ_{exc} the standard deviation of the excitatory Gaussian, A_{inh} the amplitude of the inhibitory Gaussian component, and σ_{inh} the standard deviation of the inhibitory Gaussian. Conversely, for fields where at most one bump should evolve at any time as a form of decision-making (AOL, ASL, and AEL), we implement a Gaussian kernel with global inhibition:

$$k_{uu}(x - x') = A \exp\left(-\frac{|x - x'|^2}{\sigma^2}\right) - A_{glob} \tag{3}$$

where A and σ define the amplitude and standard deviation, respectively, and A_{glob} defines the approximately constant inhibition for distant neurons.

External inputs to the AOL (hand position) and ORL (object positions) are modelled as Gaussian stimuli centred on the observed positions. The ORL receives one Gaussian per object in the workspace:

$$s_{\text{ORL}}(x, t) = \sum_i A \exp\left(-\frac{(x - p_i)^2}{2\sigma^2}\right), i = 3 \tag{4}$$

The AOL receives a single Gaussian stimulus encoding the hand position p with variable amplitude A_h that increases as the hand approaches an object:

$$s_{\text{AOL}}(x, t) = A_h \exp\left(-\frac{(x - p)^2}{2\sigma^2}\right), A_h = \sigma(-d) \tag{5}$$

where $\sigma(d)$ is a sigmoid function of the Euclidean distance d between the hand and the centre of the table (where the objects are aligned).

Beyond external inputs, DNFs interact via structured couplings. When activity in one field u projects to another field v the output from u provides spatially structured input to v via convolution [32]:

$$s_v(x,t) = \int k_{uv}(x-x')g(u(x',t))dx' \tag{6}$$

where $s_v(x,t)$ is the input to field v at position x and time t, $g(u(x',t))$ is the thresholded output of field u, and $k_{uv}(x-x')$ is a Gaussian interaction kernel (Equation 3 with $A_{glob} = 0$).

In our architecture, these couplings implement the flow of information between layers: the AOL and ORL jointly activate the ASL to infer human intent ($k_{AOL\to ASL}, k_{ORL\to ASL}$), while the ASL in turn inhibits the AEL to suppress competing actions ($k_{ASL\to AEL}$). Meanwhile, excitatory input from the ORL preshapes the AEL toward available objects ($k_{ORL\to AEL}$). Together, these interactions ensure that the robot selects complementary grasp targets rather than duplicating the participant's choices.

To support reproducibility, the code for the DNF architecture and the VR experimental setup is openly available at Jgocunha/vr-hr-joint-task. Field-specific parameters are listed in Table 1.

Table 1: Field-specific parameters of the DNF architecture. All fields shared $\tau = 100$, $h = -5$, a sigmoid output function (threshold 0, steepness 4), and noise strength $q = 0.05$.

Field	**Self-interaction kernel**	**Outgoing inter-field couplings**
AOL	Eq. 3: $A = 1.5$, $\sigma = 1$, $A_{\text{glob}} = 0.05$	AOL→ASL Eq. 3: $A = 0.755$, $\sigma = 3$
ORL	Eq. 2: $A_{\text{exc}} = 2$, $\sigma_{\text{exc}} = 1$,	ORL→ASL Eq. 3: $A = 0.7$, $\sigma = 1.9$
	$A_{\text{inh}} = 1.5$, $\sigma_{\text{inh}} = 0.5$	ORL→AEL Eq. 3: $A = 1.5$, $\sigma = 3$
ASL	Eq. 3: $A = 1.5$, $\sigma = 1$, $A_{\text{glob}} = 0.05$	ASL→AEL Eq. 3: $A = -2.5$, $\sigma = 4$
AEL	Eq. 3: $A = 2.8$, $\sigma = 2.5$, $A_{\text{glob}} = 1.1$	—

2.5 Measures

We collected both subjective and objective measures to evaluate the impact of anticipation on user perceptions and task performance.

Subjective metrics The RoSAS [10] was used to assess how the robot's anticipatory behaviour influenced participants' perceptions of its social attributes. This scale measured three key dimensions: perceived warmth, competence, and discomfort. Internal consistency was high across dimensions (Cronbach's $\alpha = 0.847, 0.805, 0.909$, respectively).

To capture perceptions of collaboration quality, we administered an 8-item Likert-scale questionnaire ($\alpha = 0.903$). Items addressed perceived control, predictability, helpfulness, trust, and preference for future collaboration. All items were rated on a 5-point Likert scale (1 = strongly disagree, 5 = strongly agree). Item 8 (perceived competition) was reverse-coded, and responses were averaged to obtain a single collaboration score:

1. I felt in control during the task.
2. The robot seemed to understand my actions.
3. The robot's behaviour was predictable.
4. The robot helped me complete the task more efficiently.
5. I would prefer to perform the task with the robot rather than alone.
6. I trusted the robot's choices during the task.
7. I would like to work with this robot on other tasks.
8. I felt the robot was competing with me at some points during the task.

Finally, after completing both conditions, participants answered open-ended questions probing their impressions of the robot's role, helpfulness, and influence on task strategy, as well as their preferences between the two robot behaviours:

1. How did you perceive the robot's role in the task?
2. What aspects of the robot's behaviour were helpful or unhelpful?
3. How could the interaction between you and the robot be improved?
4. In what ways did the robot influence your strategy for completing the task?
5. Which robot was the easiest to work with?
6. Which robot would you prefer to collaborate with?

Objective metrics We hypothesised that anticipatory behaviour would improve task performance by reducing conflicts and inefficiencies. To test this, we quantified three objective metrics. **Collisions** were defined as instances where the participant's VR controller and the robot's end-effector overlapped within a predefined spatial boundary while both agents targeted the same object (whichever agent grasped the contested object first claimed it, requiring the other to select a different target). Each boundary was modelled as a sphere with a radius of 5 cm, centred on the tip of the robot's end-effector and on the VR controller, respectively. This radius was initially set to 10 cm (approximately half the average adult hand length) but was reduced to 5 cm after pilot testing revealed false collision detections due to the geometry of the VR controller. **Idle time** captured the cumulative duration during which either agent remained inactive while waiting for the partner to act. Finally, **task completion time** measured the total duration required for both agents to place all objects in the container.

Together, these measures allow us to assess whether anticipation improves not only subjective perceptions of the robot but also observable efficiency in joint action.

2.6 Procedure

Upon arrival, participants provided informed consent and received a briefing on the task and experimental setup. They were then seated and equipped with an HTC Vive headset and controller. To familiarise participants with the VR environment and establish a baseline measure of task performance, each began with 10 practice trials performed without robot assistance.

Participants then completed two experimental sets: one with anticipatory robot behaviour and one without. Each set consisted of 20 trials, with the order of conditions counterbalanced across participants to mitigate order effects. After each set, participants completed two questionnaires: the RoSAS and the custom-designed questionnaire. At the end of the session, participants also responded to open-ended questions to provide qualitative feedback. Finally, participants were debriefed and compensated. Each session lasted approximately 30 minutes.

Video demonstrations of both experimental conditions, along with visualisations of the underlying neural dynamics, are available in the following video https://youtu.be/oyNoqGCwlLg.

3 Results

3.1 Perception of the robot

Figure 3 shows mean ratings on the RoSAS across the two conditions. Normality was assessed using the Shapiro–Wilk test; the results for Competence met the normality assumptions, while those for Warmth and Discomfort did not. Accordingly, paired-samples t-tests were used for normally distributed variables (Competence), and Wilcoxon signed-rank tests[4] were used for non-normally distributed variables (Warmth and Discomfort).

Participants rated the robot as significantly **more competent** when anticipatory behaviour was enabled ($t(9) = 4.586$, $p = 0.001$, $d = 2.06$). They also perceived the robot as significantly **less discomforting** ($V = 0$, $p = 0.009$, $r = 0.88$), and significantly **warmer** ($V = 41$, $p = 0.032$, $r = 0.71$).

These results suggest that anticipation was associated with more positive social perceptions of the robot. Even in a simple pick-and-place task, anticipation increased trust and perceived collaboration while reducing unease, consistent with the interpretation that participants perceived the robot as more socially intelligent when it proactively adapted to their actions.

3.2 Collaboration questionnaire

Participants' perceptions of collaboration quality were assessed using an eight-item Likert-scale questionnaire addressing control, predictability, efficiency, trust, and willingness to collaborate in the future. Figure 4 presents mean scores for both conditions.

[4] For the Wilcoxon tests, we report the test statistic V as returned by the R implementation.

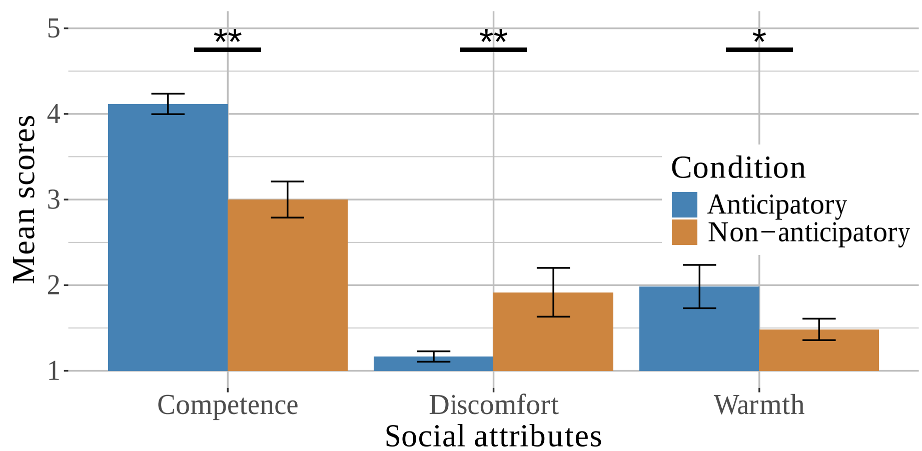

Fig. 3: Mean Likert scores on the RoSAS across the anticipatory and non-anticipatory conditions. Error bars show the standard error of the mean. Significant differences are denoted with * ($p < 0.05$) and ** ($p < 0.01$). Participants rated the anticipatory robot as significantly more **competent** and **warm**, and significantly **less discomforting**.

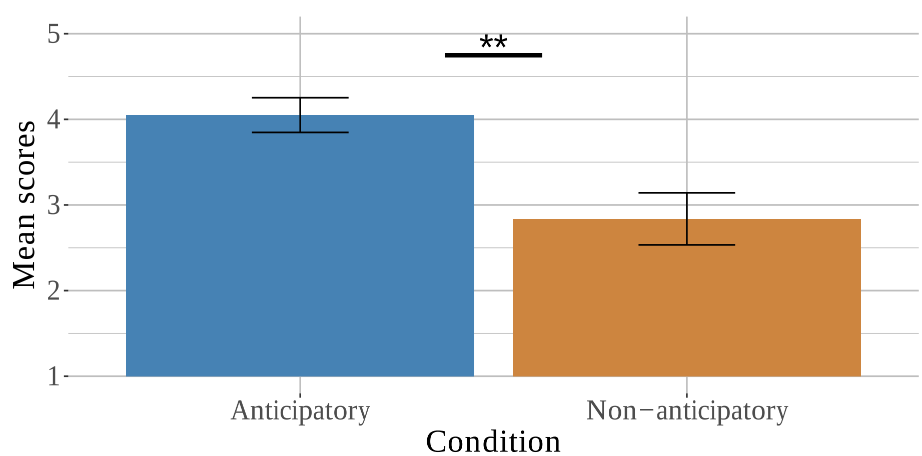

Fig. 4: Mean scores on the collaboration questionnaire across anticipatory and non-anticipatory conditions. Error bars show the standard error of the mean. Significant differences are denoted with ** ($p < 0.01$). Participants rated the anticipatory robot as a significantly more collaborative partner, reflecting greater predictability, helpfulness, trust, and willingness to work together again.

Data from both conditions met normality assumptions (anticipatory: $W = 0.911$, $p = 0.286$; non-anticipatory: $W = 0.948$, $p = 0.648$). A paired t-test revealed a significant effect of anticipation, with higher collaboration ratings in the anticipatory condition ($t(9) = 4.309$, $p = 0.002$, $d = 1.44$). The largest per-item mean differences were observed for perceived understanding of the participant's actions (M = 4.3 vs. 2.3), willingness to collaborate again (M = 4.5 vs. 2.8), and trust (M = 4.2 vs. 2.9).

These findings suggest that participants experienced the anticipatory robot as a **more collaborative partner**. Thus, beyond competence and warmth, anticipation also enhanced perceptions of mutual coordination and trust, reinforcing the role of proactive behaviour in shaping positive team dynamics.

3.3 Objective metrics

We analysed collision frequency, idle time, and task duration to assess whether anticipation improved task efficiency.

Anticipation substantially reduced **collisions** between human participants and the robot. Collisions are reported as the mean number of collisions per trial per participant (*i.e.*, total collisions divided by the number of trials in each condition), which allows for the comparison across conditions while accounting for repeated trials within participants. The number of collisions was significantly lower in the anticipatory condition compared to the non-anticipatory condition ($V = 822$, $p < 0.001$, $r = 0.309$), see Figure 5a. Since both conditions shared the same reactive fallback, this reduction reflects the benefit of prediction during the reach rather than reactive re-selection after the participant had already grasped the contested object.

Idle time was analysed using a linear mixed-effects model [2], which accounts for the repeated-measures structure of the data (multiple trials nested within participants) and allows modelling both participant-level variability and condition effects without aggregating across trials. We constructed the model with idle time as the dependent variable, agent type (human vs. robot), and condition (anticipation vs. no anticipation) as fixed effects, including their interaction, and random slopes for condition and agent type nested within participant ID. Robots spent more time idle overall than humans ($t(30.97) = 5.54, p < 0.0001$); however, there was no significant main effect of anticipation ($t(26.35) = 1.01, p = 0.32$) and no interaction between condition and agent type ($t(777.96) = -0.45, p = 0.65$). Although the mean human idle time was lower in the anticipatory condition, this difference was not statistically significant, suggesting that any effects of anticipation on idle time were small and may not be detectable with the present sample size (Figure 5b, 5c).

Task duration did not differ significantly between conditions ($V = 5422$, $p = 0.10$, $r = 0.12$), see Figure 5d. Mean task duration was slightly lower in the anticipatory condition, but the effect was small and non-significant.

Taken together, these results indicate that anticipation reliably **reduced coordination conflicts** during collaboration, even if overall efficiency measures were unaffected. This likely reflects the simplicity of the pick-and-place task; with only three objects per trial, there was limited opportunity for anticipation to shorten overall execution time.

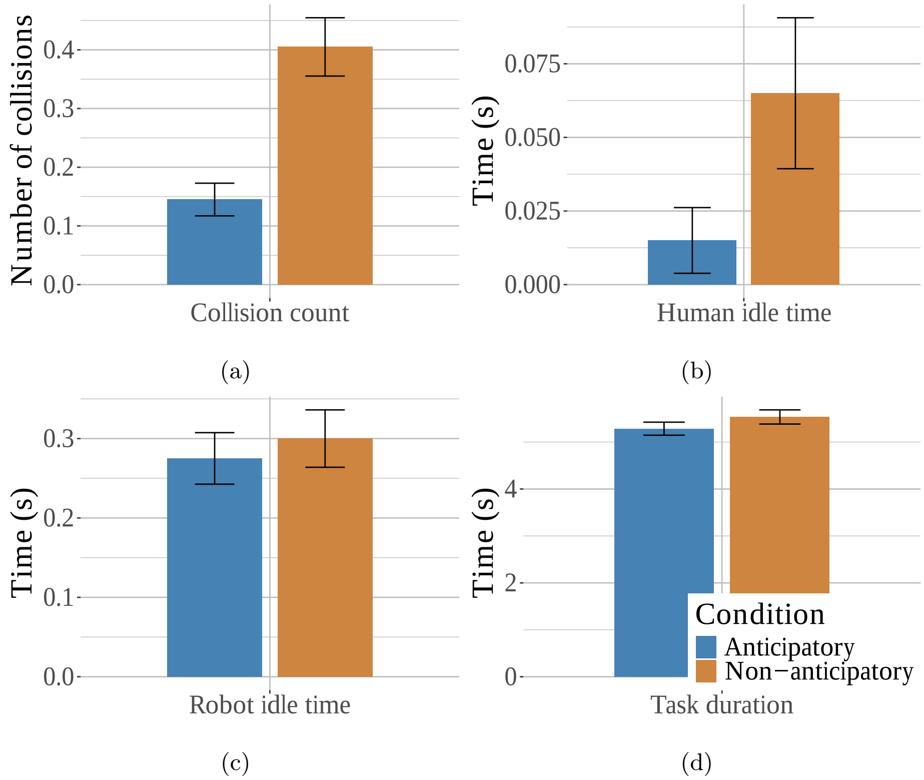

Fig. 5: Objective performance metrics across anticipatory and non-anticipatory conditions. (a) Collisions were significantly reduced when the robot anticipated the participant's actions. (b) Human idle time, (c) Robot idle time, and (d) Task duration showed no significant differences between conditions. Error bars indicate the standard error of the mean.

3.4 Open-ended questionnaire

Qualitative feedback reinforced the quantitative findings, underscoring the contrast between the two conditions.

In the anticipatory condition, participants consistently described the robot as a **helpful and competent collaborator**. They highlighted its adaptability and predictability as key benefits. For instance, one participant noted that the robot's behaviour "saved time and avoided hindrance" by "changing its mind" when reaching for the same object. Additionally, the robot's behaviours fostered a sense of cooperation, with one participant describing it as a "better team player".

In the non-anticipatory condition, the robot was perceived as competitive and obstructive. Participants described the interaction as inefficient and even adversarial. For example, one participant noted, "I had to consider the robot's movements before choosing the object myself," while another mentioned, "I felt like I had to fight for every block I grabbed". This behaviour often led participants to adopt reactive strategies, with participants waiting to see the robot's actions before making their own moves.

Overall, 9 out of 10 participants preferred collaborating with the anticipatory robot, and 8 out of 10 rated it easier to work with. They emphasised its **predictability**, **adaptability**, and **supportiveness** as qualities that made it a more desirable teammate.

4 Discussion and Conclusions

This study examined whether anticipatory action selection, implemented through a Dynamic Neural Field (DNF) architecture, influences user experience and coordination dynamics in a collaborative Virtual Reality (VR) setting. The results suggest that anticipatory behaviour enhanced perceived collaboration quality and reduced coordination conflicts.

Across the Robotic Social Attributes Scale (RoSAS), the collaboration questionnaire, and open-ended feedback, participants consistently reported a more positive interaction experience with the anticipatory robot. Anticipation enhanced perceived competence and warmth while reducing discomfort, and was associated with higher ratings of predictability, trust, and collaboration. Notably, nine out of ten participants preferred the anticipatory robot as a teammate. Anticipation altered social perception beyond coordination outcomes, with participants attributing greater warmth and competence to a robot that differed only in when it adjusted its actions. These findings reinforce the view that proactive, complementary behaviour shapes not only task execution but also the perceived quality of joint action [17,9]. In contrast, the non-anticipatory condition was perceived as obstructive and even competitive, which forced participants into reactive strategies. This also aligns with previous literature that suggests that a lack of coordination in human-robot teams can lead to frustration and a sense of inefficiency [22,8].

The objective metrics showed that anticipation significantly reduced collisions, indicating fewer instances of competing target selection and smoother division of labour. This finding is meaningful beyond its immediate efficiency implications. In HRC research, interaction fluency is not synonymous with task speed [17]. Collisions represent coordination failures, *i.e.* moments where the shared action model broke down, and their reduction is direct evidence that the robot was genuinely tracking human intent rather than merely acting in parallel. Furthermore, in physical deployments beyond VR, competing reaches carry injury risk; a robot that reliably avoids contesting the same object as its human partner is therefore safer by design, regardless of whether total task duration changes.

Idle time and task duration did not differ significantly between conditions. This reflects the structure of the task rather than the limits of the architecture: with only three objects per trial and fully parallel action possibilities, there was virtually no structural opportunity for anticipation to shorten overall duration.

Collision reduction is, therefore, the only efficiency signal the task design was sensitive enough to detect. Under these conditions, the anticipatory mechanisms primarily affected *how* the interaction unfolded, reducing conflicts and improving fluency, rather than shortening total duration. Tasks involving tighter temporal dependencies, shared object manipulation, or sequential coordination demands would provide a richer substrate for detecting performance-related effects of anticipation [39].

Several limitations should be considered. First, the study relied on a small within-subject sample (N = 10). The within-subjects design was chosen specifically to maximise sensitivity given this constraint, as it eliminates between-subject variance and allows each participant to serve as their own control. The large effect sizes observed on subjective measures provide direct evidence that the perceptual effects detected are robust. Second, evaluation was conducted in a virtual setting, which provided precise, occlusion-free hand trajectories ideal for validating the architecture's intent inference, but differs from physical collaboration in important ways. The absence of physical risk allowed participants to reach freely without concern for contact with the robot arm, whereas in a real setting, physical proximity may induce more cautious reaching behaviour, potentially amplifying coordination conflicts and increasing the relative benefit of anticipation. Additionally, the massless virtual objects could be placed effortlessly; in more physically demanding tasks, where manipulation carries a real cost, people may rely more heavily on a collaborative partner, suggesting that the benefits of anticipation observed here could be even more pronounced in embodied settings. Investigating how the architecture performs in real-world physical collaboration, therefore, remains an important direction for future work.

In sum, this exploratory study suggests that a neurally inspired DNF architecture can generate complementary anticipatory behaviour that substantially shapes perceived collaboration in HRC. Anticipation reduced coordination conflicts and enhanced perceived fluency, predictability, and trust. The use of VR provided a controlled context in which continuous movement trajectories could be directly coupled to neural field dynamics, allowing anticipatory decision processes to be examined in isolation. These findings support the potential of continuous, neurally grounded decision architectures to contribute to socially attuned human–robot teamwork.

Acknowledgments. This work was financed by FCT (Portuguese Foundation for Science and Technology) through the R&D Unit Project Scope UID/00319/2025 - Centro ALGORITMI (ALGORITMI/UM) and PhD Fellowship under grant 2022.13948.BD.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Amari, S.i.: Dynamics of pattern formation in lateral-inhibition type neural fields. Biol. Cybern. **27**(2), 77–87 (1977). https://doi.org/10.1007/BF00337259

2. Bates, D., Mächler, M., Bolker, B., Walker, S.: Fitting Linear Mixed-Effects Models Using **lme4**. J. Stat. Soft. **67**(1), 51 (2015). https://doi.org/10.18637/jss.v067.i01
3. Bekkering, H., De Bruijn, E.R.A., Cuijpers, R.H., Newman-Norlund, R., Van Schie, H.T., Meulenbroek, R.: Joint Action: Neurocognitive Mechanisms Supporting Human Interaction. Topics in Cognitive Science **1**(2), 340–352 (Apr 2009). https://doi.org/10.1111/j.1756-8765.2009.01023.x
4. Bicho, E., Erlhagen, W., Louro, L., Costa E Silva, E.: Neuro-cognitive mechanisms of decision making in joint action: A human–robot interaction study. Human Movement Science **30**(5), 846–868 (Oct 2011). https://doi.org/10.1016/j.humov.2010.08.012
5. Bicho, E., Erlhagen, W., Louro, L., Silva, E.C.E., Silva, R., Hipolito, N.: A dynamic field approach to goal inference, error detection and anticipatory action selection in human-robot collaboration. In: Dautenhahn, K., Saunders, J. (eds.) Advances in Interaction Studies, vol. 2, pp. 135–164. John Benjamins Publishing Company, Amsterdam (Dec 2011). https://doi.org/10.1075/ais.2.10bic
6. Bogaerts, B., Sels, S., Vanlanduit, S., Penne, R.: Connecting the CoppeliaSim robotics simulator to virtual reality. SoftwareX **11**, 100426 (Jan 2020). https://doi.org/10.1016/j.softx.2020.100426
7. Boschetti, G., Bottin, M., Faccio, M., Maretto, L., Minto, R.: The influence of collision avoidance strategies on human-robot collaborative systems. IFAC-PapersOnLine **55**(2), 301–306 (2022). https://doi.org/10.1016/j.ifacol.2022.04.210
8. Campagna, G., Rehm, M.: Trust Assessment with EEG Signals in Social Human-Robot Interaction. In: Ali, A.A., Cabibihan, J.J., Meskin, N., Rossi, S., Jiang, W., He, H., Ge, S.S. (eds.) Social Robotics, vol. 14453, pp. 33–42. Springer Nature Singapore, Singapore (2024). https://doi.org/10.1007/978-981-99-8715-3_4, series Title: Lecture Notes in Computer Science
9. Cantucci, F., Marini, M., Falcone, R.: The Role of Robot Competence, Autonomy, and Personality on Trust Formation in Human-Robot Interaction. J. Hum.-Robot Interact. p. 3786770 (Dec 2025). https://doi.org/10.1145/3786770
10. Carpinella, C.M., Wyman, A.B., Perez, M.A., Stroessner, S.J.: The Robotic Social Attributes Scale (RoSAS): Development and Validation. In: Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction. pp. 254–262. ACM, Vienna Austria (Mar 2017). https://doi.org/10.1145/2909824.3020208
11. Chatziparaschis, D., Zhong, S., Christopoulos, V., Karydis, K.: Adaptive Environment-Aware Robotic Arm Reaching Based on a Bio-Inspired Neurodynamical Computational Framework. In: 2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN). pp. 510–515. IEEE, Pasadena, CA, USA (Aug 2024). https://doi.org/10.1109/RO-MAN60168.2024.10731226
12. Duarte, N.F., Rakovic, M., Tasevski, J., Coco, M.I., Billard, A., Santos-Victor, J.: Action Anticipation: Reading the Intentions of Humans and Robots. IEEE Robot. Autom. Lett. **3**(4), 4132–4139 (Oct 2018). https://doi.org/10.1109/LRA.2018.2861569
13. Erlhagen, W., Bicho, E.: The dynamic neural field approach to cognitive robotics. J. Neural Eng. **3**(3), R36–R54 (Sep 2006). https://doi.org/10.1088/1741-2560/3/3/R02

14. Fishman, A., Paxton, C., Yang, W., Fox, D., Boots, B., Ratliff, N.: Collaborative Interaction Models for Optimized Human-Robot Teamwork. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 11221–11228. IEEE, Las Vegas, NV, USA (Oct 2020). https://doi.org/10.1109/IROS45743.2020.9341369
15. Frith, C.D., Frith, U.: How we predict what other people are going to do. Brain Research **1079**(1), 36–46 (Mar 2006). https://doi.org/10.1016/j.brainres.2005.12.126
16. Gallese, V.: Mirror neurons and the simulation theory of mind-reading. Trends in Cognitive Sciences **2**(12), 493–501 (Dec 1998). https://doi.org/10.1016/S1364-6613(98)01262-5
17. Hoffman, G., Breazeal, C.: Cost-Based Anticipatory Action Selection for Human–Robot Fluency. IEEE Trans. Robot. **23**(5), 952–961 (Oct 2007). https://doi.org/10.1109/TRO.2007.907483
18. Huang, C.M., Mutlu, B.: Anticipatory robot control for efficient human-robot collaboration. In: 2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI). pp. 83–90. IEEE, Christchurch, New Zealand (Mar 2016). https://doi.org/10.1109/HRI.2016.7451737
19. Izquierdo-Badiola, S., Canal, G., Rizzo, C., Alenya, G.: Improved Task Planning through Failure Anticipation in Human-Robot Collaboration. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 7875–7880. IEEE, Philadelphia, PA, USA (May 2022). https://doi.org/10.1109/ICRA46639.2022.9812236
20. Koppula, H.S., Saxena, A.: Anticipating Human Activities Using Object Affordances for Reactive Robotic Response. IEEE Trans. Pattern Anal. Mach. Intell. **38**(1), 14–29 (Jan 2016). https://doi.org/10.1109/TPAMI.2015.2430335
21. Laplaza, J., Moreno, F., Sanfeliu, A.: Enhancing Robotic Collaborative Tasks Through Contextual Human Motion Prediction and Intention Inference. Int J of Soc Robotics **17**(10), 2077–2096 (Oct 2025). https://doi.org/10.1007/s12369-024-01140-2
22. Lasota, P.A., Fong, T., Shah, J.A.: A Survey of Methods for Safe Human-Robot Interaction. FNT in Robotics **5**(4), 261–349 (2017). https://doi.org/10.1561/2300000052
23. Liu, C., Hamrick, J.B., Fisac, J.F., Dragan, A.D., Hedrick, J.K., Sastry, S.S., Griffiths, T.L.: Goal Inference Improves Objective and Perceived Performance in Human-Robot Collaboration. In: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems. pp. 940 – 948. AAMAS '16, International Foundation for Autonomous Agents and Multiagent Systems, Singapore, Singapore (2016). https://doi.org/10.48550/arXiv.1802.01780
24. Mainprice, J., Berenson, D.: Human-robot collaborative manipulation planning using early prediction of human motion. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 299–306. IEEE, Tokyo (Nov 2013). https://doi.org/10.1109/IROS.2013.6696368
25. Menezes, J.C.: MuModaR: Multi-modal Framework for Human-Robot Collaboration in Cyber-physical Systems. In: Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction. pp. 755–759. ACM, Boulder CO USA (Mar 2024). https://doi.org/10.1145/3610978.3640688
26. Pezzulo, G., Ognibene, D.: Proactive Action Preparation: Seeing Action Preparation as a Continuous and Proactive Process. Motor Control **16**(3), 386–424 (Jul 2012). https://doi.org/10.1123/mcj.16.3.386

27. Plomin, J., Schweidler, P., Oehme, A.: Virtual reality check: a comparison of virtual reality, screen-based, and real world settings as research methods for HRI. Front. Robot. AI **10**, 1156715 (Jun 2023). https://doi.org/10.3389/frobt.2023.1156715
28. Pons, P., Soler-Dominguez, J.L., Navas-Medrano, S.: Designing Gesture-Based Interfaces for Cobots: Insights from a Mixed Reality Elicitation Study. J. Hum.-Robot Interact. p. 3786760 (Dec 2025). https://doi.org/10.1145/3786760
29. Psarakis, L., Nathanael, D., Marmaras, N.: Communicating robots' intent through visual cues enhances human anticipatory behavior in human–dual robot collaboration. Robotics and Computer-Integrated Manufacturing **92**, 102886 (Apr 2025). https://doi.org/10.1016/j.rcim.2024.102886
30. Rohmer, E., Singh, S.P.N., Freese, M.: V-REP: A versatile and scalable robot simulation framework. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 1321–1326. IEEE, Tokyo (Nov 2013). https://doi.org/10.1109/IROS.2013.6696520
31. Sandamirskaya, Y., Zibner, S.K., Schneegans, S., Schöner, G.: Using Dynamic Field Theory to extend the embodiment stance toward higher cognition. New Ideas in Psychology **31**(3), 322–339 (Dec 2013). https://doi.org/10.1016/j.newideapsych.2013.01.002
32. Schöner, G., Spencer, J., Research Group, D.: Dynamic Thinking: A Primer on Dynamic Field Theory. Oxford University Press, Oxford (Nov 2015). https://doi.org/10.1093/acprof:oso/9780199300563.001.0001
33. Sciutti, A., Mara, M., Tagliasco, V., Sandini, G.: Humanizing Human-Robot Interaction: On the Importance of Mutual Understanding. IEEE Technol. Soc. Mag. **37**(1), 22–29 (Mar 2018). https://doi.org/10.1109/MTS.2018.2795095
34. Sebanz, N., Bekkering, H., Knoblich, G.: Joint action: bodies and minds moving together. Trends in Cognitive Sciences **10**(2), 70–76 (Feb 2006). https://doi.org/10.1016/j.tics.2005.12.009
35. Seok, S., Jeon, T.H., Chae, Y.J., Kim, C., Lim, Y.: Explorative Study on the Nonverbal Backchannel Prediction Model for Human-Robot Interaction. In: Ali, A.A., Cabibihan, J.J., Meskin, N., Rossi, S., Jiang, W., He, H., Ge, S.S. (eds.) Social Robotics, vol. 14453, pp. 264–275. Springer Nature Singapore, Singapore (2024). https://doi.org/10.1007/978-981-99-8715-3_23, series Title: Lecture Notes in Computer Science
36. Silva, G., Olivier, A.H., Crétual, A., Pettré, J., Fraichard, T.: Effective Human-Robot Collaboration in near symmetry collision scenarios. In: 2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN). pp. 1–8. IEEE, New Delhi, India (Oct 2019). https://doi.org/10.1109/RO-MAN46459.2019.8956360
37. Silva, R., Louro, L., Malheiro, T., Erlhagen, W., Bicho, E.: Combining intention and emotional state inference in a dynamic neural field architecture for human-robot joint action. Adaptive Behavior **24**(5), 350–372 (Oct 2016). https://doi.org/10.1177/1059712316665451
38. Wang, Z., Boularias, A., Mülling, K., Schölkopf, B., Peters, J.: Anticipatory action selection for human–robot table tennis. Artificial Intelligence **247**, 399–414 (Jun 2017). https://doi.org/10.1016/j.artint.2014.11.007
39. Wojtak, W., Ferreira, F., Louro, L., Bicho, E., Erlhagen, W.: Adaptive timing in a dynamic field architecture for natural human–robot interactions. Cognitive Systems Research **82**, 101148 (Dec 2023). https://doi.org/10.1016/j.cogsys.2023.101148