Shall We Dance? Resonance of Intentions with an Embodied Agent based on the Free Energy Principle

Takeru Hashimoto
Sony computer science laboratories, inc.
Tokyo, Japan
Okinawa Institute of Science and Technology
Okinawa, Japan
takeru.hashimoto.research@gmail.com

Jun Tani
Okinawa Institute of Science and Technology
Okinawa, Japan
jun.tani@oist.jp

Shunichi Kasahara
Sony computer science laboratories, inc.
Tokyo, Japan
Okinawa Institute of Science and Technology
Okinawa, Japan
kasahara@csl.sony.co.jp

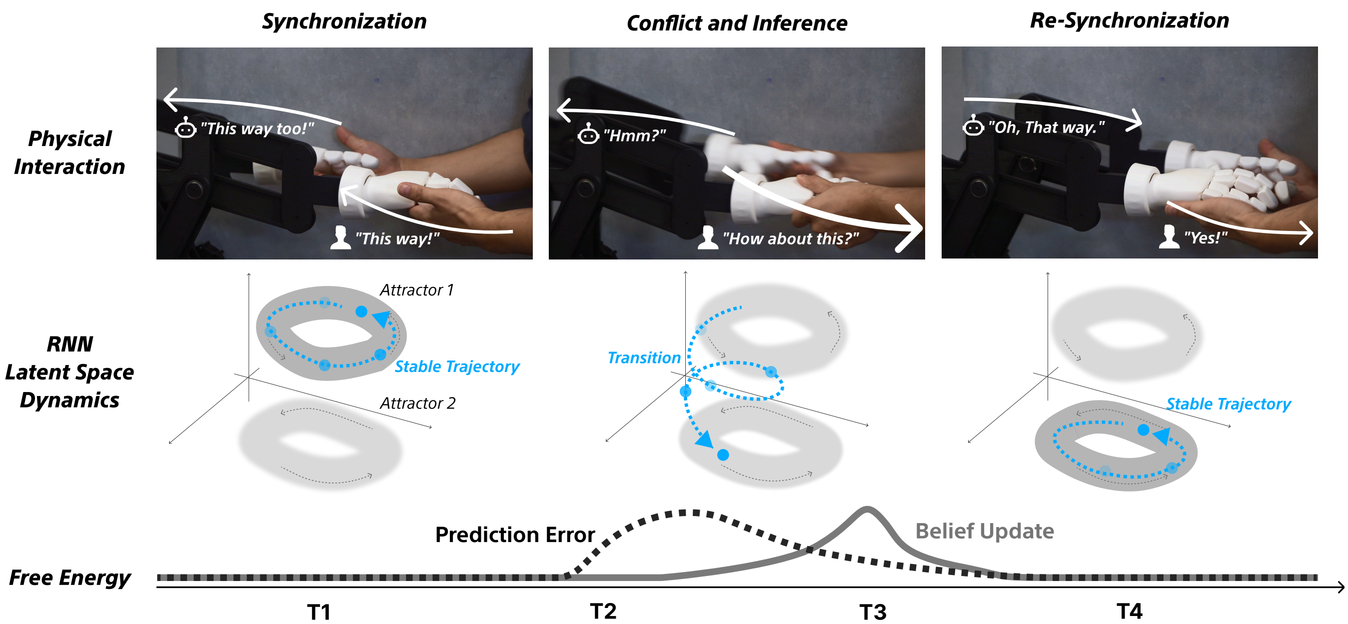

Figure 1: Real-time adaptive physical coordination via latent attractor dynamics driven by free energy minimization. The robot agent and human perform a continuous hand-in-hand dance. The process illustrates the agent's adaptation to a human partner's changing intention through four phases: (T1) Synchronization: They share a synchronized rhythm, represented as a stable limit cycle (Attractor 1) in the RNN latent space. (T2) Conflict: When the human changes intention, the physical mismatch generates a spike in Prediction Error in the agent's brain (dashed line). (T3) Inference: To minimize this prediction error, the agent infers how to update their belief (i.e., Belief Update, solid line). This corresponds to driving the internal state to transition from the current attractor to a new one (Attractor 1 → 2). (T4) Re-synchronization: The internal state settles into the new attractor, allowing the agent to physically adapt to the human's new intention.

Abstract

Physical joint action shapes a sense of partnership by enabling individuals to negotiate intentions through mutual coordination. However, conventional physical agents typically rely on pre-learned transition rules, which limits their ability to adapt to a partner's improvisational interventions. To address this, we propose a physical agent grounded in the Free Energy Principle that negotiates intentions through physical interaction. The agent uses a predictive coding-based RNN to dynamically update their internal intention (μ) and time constant (τ), minimizing real-time prediction errors caused by physical contact. This mechanism allows the agent to synchronize with the human's motion pattern and tempo. In the demonstration "Shall We Dance?", visitors hold hands with the agent and haptically perceive the resistance at the hands disappear as the agent infers the human's new intent.

SIGGRAPH Emerging Technologies '26, Los Angeles, CA, USA

ACM ISBN 979-8-4007-2549-4/26/07
https://doi.org/10.1145/3799826.3812424

Keywords

Physical Human-Robot Interaction, Joint Action, Free Energy Principle, Predictive Coding, Active Inference

ACM Reference Format:
Takeru Hashimoto, Jun Tani, and Shunichi Kasahara. 2026. Shall We Dance? Resonance of Intentions with an Embodied Agent based on the Free Energy Principle. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Emerging Technologies (SIGGRAPH Emerging Technologies '26), July 19–23, 2026, Los Angeles, CA, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3799826.3812424

1 Introduction

Human joint action involves continuous negotiation, where partners maintain their own intentions (priors) while inferring those of others. This coordination can enhance motor performance beyond individual capabilities in terms of precision and efficiency [Ganesh et al. 2014]. However, most conventional physical agents treat interaction errors merely as control problems to be resolved through motor output. Consequently, these systems lack the cognitive capacity to lead with intent or yield with understanding.

In cognitive robotics, predictive coding-inspired Variational RNNs have been proposed to realize such cognitive give-and-take based on free energy minimization [Ahmadi and Tani 2019]. While these models enable robots to mimic behaviors like turn-taking [Wirkuttis et al. 2023] and mutual yielding [Sawada et al. 2025], they typically rely on transitions learned as statistical regularities from training data. As a result, these agents struggle to adapt to improvisational interventions that occur at arbitrary moments.

To address this, we propose a system grounded in the Free Energy Principle that eliminates the need for learned switching rules. Instead, intention switching emerges dynamically through real-time prediction error minimization during physical contact. This allows the agent to follow a partner's changes from any state, regardless of prior training. We demonstrate this resonance using a hand-in-hand dance setup between a human and a physical agent.

2 System Architecture

2.1 Predictive Coding-based RNN Model

The core of our system is a continuous-time RNN (CTRNN) based on the concept of predictive coding [Friston 2005]. To enable the agent to autonomously generate multiple movement patterns as prior beliefs, we embed multiple limit cycles as attractors in the high-dimensional latent space through training (Fig. 2 top). The internal state $\boldsymbol{\mu}_t \in \mathbb{R}^N$ ($N = 64$) represents the expected state of the hidden dynamics, updated as:

$$\begin{aligned} \boldsymbol{\mu}_t &= \left(1 - \frac{1}{\tau}\right)\boldsymbol{\mu}_{t-1} + \frac{1}{\tau}(\mathbf{W}_h\mathbf{h}_{t-1} + \mathbf{b}_h) \\ \mathbf{h}_t &= \tanh(\boldsymbol{\mu}_t) \end{aligned} \tag{1}$$

where τ is the time constant, $\mathbf{W}_h$ is the recurrent weight matrix, $\mathbf{b}_h$ is the bias vector, and $\mathbf{h}_t$ is the activation value. Under the Free Energy Principle, we treat the observed movement $\mathbf{x}_t$ and state $\boldsymbol{\mu}_t$ as Gaussian random variables. During training, the network minimizes the Variational Free Energy $\mathcal{F}$, approximated by the sum of squared prediction errors:

$$\mathcal{F} \approx \underbrace{\frac{\pi_x}{2}||\mathbf{x}_t - g(\boldsymbol{\mu}_t)||^2}_{\text{Prediction Error } E_x} + \underbrace{\frac{\pi_\mu}{2}||\boldsymbol{\mu}_t - f(\boldsymbol{\mu}_{t-1})||^2}_{\text{State Error } E_\mu} \tag{2}$$

where $g(\cdot)$ represents the decoder mapping latent states to physical coordinates, and $f(\cdot)$ denotes the transition dynamics defined in Eq. (1).

2.2 Real-time Intention and Tempo Negotiation

In the inference phase, the model weights are fixed, and the system optimizes $\boldsymbol{\mu}$ and τ in real-time against sequential observations $\mathbf{x}$. By minimizing $\mathcal{F}$, the agent calculates the optimal $\boldsymbol{\mu}$ that explains the partner's movement while maintaining their current prior. This allows the agent to adjust their phase or switch between attractors by crossing the saddle when the partner's movement deviates from the current prediction. Simultaneously, by updating the time constant τ, the agent adapts their execution speed to match the partner's tempo.

2.3 Action-Perception Loop

The interaction is realized through the simultaneous processes of action and perception (Fig. 2 bottom). In perception, the agent updates their latent state $\boldsymbol{\mu}_t$ to infer the partner's intention. In action, the agent's next-step prediction is directly used as the target position, from which motor commands are computed and sent to a custom-made 2-DOF force-feedback robot arm. The inference loop, including observation, state update, and motor command generation, runs at 30 Hz on a laptop CPU. The motor commands are low-pass filtered and sent to the actuators at 100 Hz. This in-situ loop enables the agent to transition from leading to yielding based on real-time prediction errors.

3 Technical Validation

We validated the active inference capability by testing whether the agent could update their intention through external intervention. While the agent initially generated learned motions, a human partner applied force to guide it toward target trajectories, alternating between pattern 1 and 2. Fig. 3 illustrates the adaptive process through physical interaction. When a human applies force (blue areas in the graph), the Prediction Error E_x initially rises due to the sensory-motor discrepancy. To minimize this, the agent updates their internal state, causing a transient increase in State Error E_μ. As the agent adapts to the new pattern, both errors decrease, and the system re-synchronizes. This sequence, from external force to state update and re-settling, confirms that the inference mechanism adapts the agent's intention in real-time.

To achieve temporal coordination, we treat motion speed as an inferable variable by introducing a speed ratio α that modulates the time constant τ in Equation 1. By updating the effective time constant as τ/α, the agent can accelerate or decelerate their internal dynamics to match the partner's tempo. Fig. 4 demonstrates this adaptation as the partner changes the interaction frequency. The agent infers the optimal α to minimize prediction errors. As shown in the bottom plot, α dynamically tracks the partner's speed, so that the agent synchronizes their tempo while the original attractor

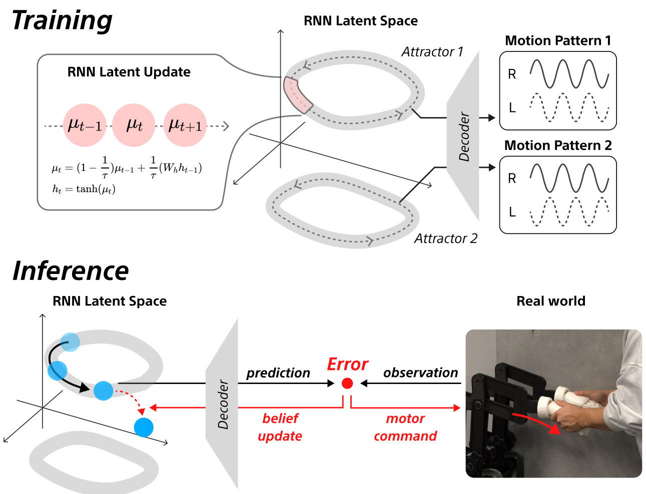

Figure 2: Schematic of the proposed RNN-based active inference model. Training: The network learns multiple motion primitives as distinct attractors in the latent space. Inference: When a movement gap occurs, the resulting prediction error triggers a belief update. By shifting between attractors to minimize this error, the agent synchronizes in real-time.

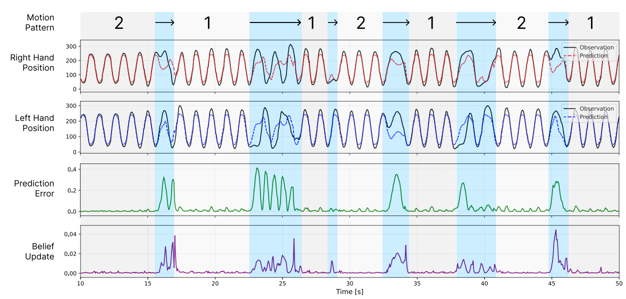

Figure 3: Time-series of motion pattern switching. The top label shows the motion pattern (in-phase or anti-phase). The middle two plots show the trajectories of the right and left hands . The bottom two plots show the prediction error and belief update. Blue shaded regions indicate the periods where the human applied force to switch the motion pattern.

structure is maintained. These results show that the model adapts to varying frequency through real-time inference of the time constant.

4 Demonstration: Shall We Dance?

In this installation, the visitor grasps the hand of the robot with both hands to dance together. The agent begins generating one of their learned motion patterns, and the visitor is invited to follow along. Once synchronized, the visitor is encouraged to change the motion pattern. At the moment of transition, the agent's inference process is directly felt through the palms: a stiffening when intentions diverge, followed by a sudden lightness as the agent grasps your intent.

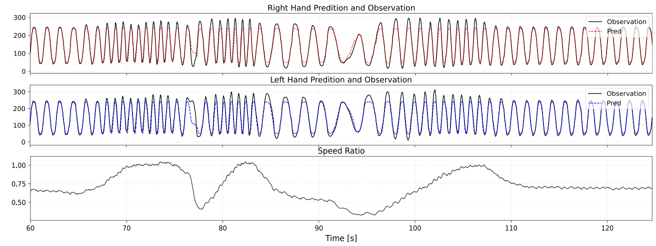

Figure 4: Time-series log of tempo adaptation. The top two plots show the trajectories of the Right and Left hands. The bottom plot shows the transition of the inferred Speed Ratio (α). As the partner changes the frequency of the pattern, the agent updates α to synchronize the shared tempo.

5 Future Work

In the current system, both the cognitive and mechanical impedance of the robot are fixed. Consequently, during the negotiation from an intention conflict to the next stable state, the agent does not ease their force to yield to the human's intention, so reaction forces persist throughout the transition. A future direction is to dynamically adjust the cognitive and mechanical impedance so as to minimize free energy. Changes in motor impedance affect the perceived properties of objects and of one's own body [Hashimoto et al. 2022, 2025]. Force intervention during movement can also reconfigure the user's body perception itself [Hashimoto et al. 2023]. Incorporating such impedance modulation into the active inference framework could allow haptic interaction to serve as a communication channel through which humans infer the internal state of the robot more directly.

Acknowledgments

This work was supported by JST Grant Number JPMJCR2552, JPMJPF2205 and JPMJMS2013.

References

Ahmadreza Ahmadi and Jun Tani. 2019. A novel predictive-coding-inspired variational RNN model for online prediction and recognition. *Neural Comput.* 31, 11 (Nov. 2019), 2025–2074. doi:10.1162/neco_a_01228

Karl Friston. 2005. A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 360, 1456 (April 2005), 815–836. doi:10.1098/rstb.2005.1622

G Ganesh, A Takagi, R Osu, T Yoshioka, M Kawato, and E Burdet. 2014. Two is better than one: physical interactions improve motor performance in humans. *Sci. Rep.* 4 (Jan. 2014), 3824. doi:10.1038/srep03824

Takeru Hashimoto, Shigeo Yoshida, and Takuji Narumi. 2022. MetamorphX: An Ungrounded 3-DoF Moment Display that Changes its Physical Properties through Rotational Impedance Control. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Article 72, 14 pages. doi:10.1145/3526113.3545650

Takeru Hashimoto, Shigeo Yoshida, and Takuji Narumi. 2023. SomatoShift: A Wearable Haptic Display for Somatomotor Reconfiguration via Modifying Acceleration of Body Movement. In *ACM SIGGRAPH 2023 Emerging Technologies (SIGGRAPH '23)*. Article 17, 2 pages. doi:10.1145/3588037.3595390

Takeru Hashimoto, Shigeo Yoshida, and Takuji Narumi. 2025. Effects of Impedance Control on Haptic Perception of Object Properties: A Qualitative and Quantitative Approach. *IEEE Access* 13 (2025), 52432–52445. doi:10.1109/ACCESS.2025.3553042

Hiroki Sawada, Wataru Ohata, and Jun Tani. 2025. Human-robot kinaesthetic interactions based on the free-energy principle. *IEEE Trans. Syst. Man Cybern. Syst.* 55, 1 (Jan. 2025), 366–379. doi:10.1109/tsmc.2024.3481251

Nadine Wirkuttis, Wataru Ohata, and Jun Tani. 2023. Turn-taking mechanisms in imitative interaction: Robotic social interaction based on the free energy principle. *Entropy (Basel)* 25, 2 (Jan. 2023), 263. doi:10.3390/e25020263