

# CERNet: Class-Embedding Predictive-Coding RNN for Unified Robot Motion, Recognition, and Confidence Estimation

Hiroki Sawada<sup>1</sup>, Alexandre Pitti<sup>1</sup> and Mathias Quoy<sup>1</sup>

**Abstract**—Robots interacting with humans must not only generate learned movements in real-time, but also infer the intent behind observed behaviors and estimate the confidence of their own inferences. This paper proposes a unified model that achieves all three capabilities within a single hierarchical predictive-coding recurrent neural network equipped with a class embedding vector, CERNet, which leverages a dynamically updated class embedding vector to unify motor generation and recognition. The model operates in two modes: generation and inference. In the generation mode, the class embedding constrains the hidden state dynamics to a class-specific subspace; in the inference mode, it is optimized online to minimize prediction error, enabling real-time recognition. Validated on a humanoid robot across 26 kinesthetically taught alphabets, our hierarchical model achieves 76% lower trajectory reproduction error than a parameter-matched single-layer baseline, maintains motion fidelity under external perturbations, and infers the demonstrated trajectory class online with 68% Top-1 and 81% Top-2 accuracy. Furthermore, internal prediction errors naturally reflect the model’s confidence in its recognition. This integration of robust generation, real-time recognition, and intrinsic uncertainty estimation within a single neural network framework offers a compact and extensible approach to motor memory in physical robots, with potential applications in intent-sensitive human–robot collaboration.

## I. INTRODUCTION

Robots that share workspaces and interact with humans must be capable of *generating* learnt behaviors, *recognizing* human intentions, and *estimating their own confidence* in uncertain, real-time environments. Predictive–Coding (PC) networks offer a promising framework for achieving such capabilities by minimizing prediction errors through top-down prediction and bottom-up inference. However, existing implementations often treat these functions separately, leaving few PC-based models that unify all three processes within a single closed-loop architecture for robotics. Even among those few, such integration typically relies on complex, multi-module systems; by contrast, our approach implements all functions within a single PC recurrent neural network (RNN).

While various predictive-coding networks have been proposed for visual prediction [1], visuo-motor control [2], or hierarchical perception [3], most still rely on external classifiers or post-hoc thresholding for recognition and confidence estimation. Moreover, emerging methods that embed confidence-related dynamics within inference [4], [5] often

require separate decision modules and remain unvalidated on physical robotic platforms. To concretely evaluate these challenges, we focus on the task of learning alphabet writing trajectories and recognizing the corresponding character classes from observed trajectory segments in real time. Within this setting, a single, parameter-efficient, multi-layer predictive-coding RNN that unifies action generation, label inference, and intrinsic confidence estimation on physical hardware under external disturbances is still lacking. To fill this gap, we introduce the Class-Embedding predictive-coding Recurrent NETwork (CERNet).

To jointly achieve unified action generation, label inference, and confidence estimation under a single model, we combine three design principles in CERNet. First, a multi-layer architecture allows higher layers to maintain abstract motion intentions across longer timescales [6], enabling the model to stably reproduce demonstrated trajectories even on physical hardware. In our experiments, this structure outperforms single-layer networks of equivalent parameter size, which fail to accurately reproduce the demonstrated trajectories under real-world conditions. Second, predictive-coding framework continuously minimizes prediction errors between top-down expectations and sensory input [7]. These dynamics enable the model to adapt to external perturbations (e.g., unforeseen forces or deviations arising in real-world robotic tasks) during execution, recovering towards the target trajectory by updating internal states and correcting future predictions in real time. Third, we introduce a class embedding vector that is updated online via the same prediction-error minimization. As the network observes a motion, the embedding vector gradually drifts toward the corresponding latent subspace, effectively serving as a self-organizing key for label inference. Moreover, by analyzing reconstruction errors of past observations, the network can internally estimate the confidence of its recognition, without requiring any explicit classifier or output module. While each of these principles has been studied in isolation or in simulation, this is the first work to integrate them into a single predictive-coding RNN and validate their combined effectiveness on a physical robotic platform.

Therefore, our experiments reveal four key capabilities of unified architecture of CERNet. First, its hierarchical structure greatly reduces reproduction error, highlighting the benefit of layered temporal abstraction. Second, the predictive-coding framework confers robustness to disturbances: after perturbations, the robot autonomously recovers to the intended trajectory. Third, by tuning the prediction–recognition balance, the model can infer the intended motion class in real time with high accuracy. Fourth, the same internal error signal

This work was supported by the France Excellence scholarship from the French Ministry for Europe and Foreign Affairs.

<sup>1</sup>Hiroki Sawada, Alexandre Pitti and Mathias Quoy are with Equipe Traitement de l’Information et Systèmes Laboratory (ETIS Laboratory), CNRS UMR8051, CY Cergy-Paris Université, ENSEA, Cergy, France [hiroki.sawada@oist.jp](mailto:hiroki.sawada@oist.jp), [alexandre.pitti@cyu.fr](mailto:alexandre.pitti@cyu.fr), [mathias.quoy@cyu.fr](mailto:mathias.quoy@cyu.fr)

that guides inference also indicates confidence, enabling implicit self-evaluation without extra classifiers. These results demonstrate that predictive-coding architectures can serve not only as robust motor generators but also as flexible intention recognizers within a single compact model, providing a foundation for future applications in human–robot collaboration.

The remainder of this paper is organized as follows: Section 2 reviews related work, Section 3 introduces the proposed CERNet architecture, Section 4 describes the experimental setup with the humanoid robot *Reachy*, Section 5 presents results from both simulation and real-world trials, and Section 6 concludes with a summary and future directions.

## II. RELATED WORKS

### A. Unifying Recognition, Generation and Confidence

Many recent robot learning models focus either on *recognition*, inferring task identity or intention from sensory observations, or on *generation*, producing learned trajectories from latent representations. Most frameworks treat these two functions separately, using distinct architectures for perception and motion control.

For recognition, common approaches include discriminative RNNs trained with cross-entropy loss [8], latent-variable inference via modular policies [9]. These methods typically require large labeled datasets and assume segmented, unambiguous observations. Generation, on the other hand, is often addressed with autoregressive sequence models such as LSTMs [10], Transformers [11], and more recently, policy-conditioned diffusion models [12]. While these models achieve high-fidelity generation, they generally lack the capacity for real-time task inference. Only a few works attempt to unify these roles within a single architecture. Latent-variable models such as GAIL [13], World Models [14], or MTRNN-based frameworks [6] leverage shared internal dynamics for both recognition and action generation. However, these typically rely on offline clustering or external classifiers and are rarely evaluated on real-time, closed-loop robotic platforms.

While a few studies have explored confidence estimation alongside recognition or generation, they typically rely on separate modules or post-hoc uncertainty measures. While recent studies embed confidence-related dynamics within inference, such as precision-weighted hierarchies [4] and uncertainty-minimizing active sensing [5], they often rely on separate decision stages and remain unvalidated on physical hardware. In contrast, our approach leverages the same internal prediction error that drives generation and recognition, allowing confidence to emerge as an intrinsic property of the network’s ongoing inference process.

Our work addresses this gap by demonstrating that CERNet can serve simultaneously as a generator, a recognizer, and an estimator of its own confidence. The same network dynamics update online through prediction-error minimization, allowing recognition to emerge gradually as motion is observed, confidence to be estimated from reconstruction error, and generation to take over seamlessly without retraining or architectural changes.

### B. Predictive Coding RNNs

Predictive-coding recurrent neural networks (PC-RNNs) maintain internal representations by alternately generating top-down predictions and correcting them using bottom-up inference [7]. This recursive process not only enables online adaptation but also naturally couples generation and recognition: actions are generated by minimizing prediction error, while recognition emerges through internal state updates that best explain sensory observations.

A key insight of the PC framework is that the magnitude and direction of internal prediction errors encode information about the plausibility of the current internal belief [15]. In other words, the network’s confidence can be inferred directly from its internal error dynamics, without requiring separate uncertainty estimation modules. This property makes PC-RNNs particularly attractive for real-time robotics, where decisions must be made under uncertainty and with limited computation. Furthermore, when combined with a class embedding vector that is updated through the same prediction-error minimization loop, the model can unify label inference and motor generation within a single dynamical system. The embedding trajectory gradually converges to a class-specific latent subspace, guiding both inference and generation, while reconstruction error implicitly reflects recognition confidence.

Existing PC-RNN variants, such as PV-RNN [16] or Active Predictive Coding [17], have explored stochastic representations and multi-timescale dynamics, respectively. Class-gated extensions of PC-RNNs have also been proposed to mitigate pattern interference and enable continual sequence learning [18], [19]. Complementary work on Free-Energy-based training schemes, including INFERNO and its gated or spiking variants [20], [21], demonstrates how iterative error minimization can scale to long temporal structures. However, these studies still separate recognition from generation via offline decoding or manual gating, and rarely exploit prediction-error magnitudes as a direct confidence signal. We therefore start from the single-layer class-gated PC-RNN of [18] and expand it in four directions: a three-layer hierarchy, deployment on real robotic hardware, real-time class inference, and intrinsic confidence estimation from prediction errors. These extensions turn a simulation-only proof-of-concept into a unified, closed-loop system capable of robust generation, recognition, and self-evaluation on a physical platform.

## III. METHODS

### A. CERNet: Class-Embedding Predictive-Coding RNN

To realize unified generation, recognition, and confidence estimation under a single dynamical system, we propose the Class-Embedding Predictive-Coding RNN (CERNet). Its architecture extends conventional predictive-coding RNNs by incorporating a class embedding vector that dynamically modulates the internal dynamics via the same prediction-error minimization loop. The model’s overall structure is shown in Fig. 1. Firstly, it is equipped with a one-hot-like class embedding vector  $\mathbf{C} \in \mathbb{R}^K$ , where  $K$  denotes the number

of predefined classes. The vector  $\mathbf{C}^k$  of selected class index  $k \in \{1, \dots, K\}$  is defined as:

$$C_i^k = \begin{cases} 1, & \text{if } i = k \\ 0, & \text{otherwise} \end{cases} \quad \text{for } i = 1, \dots, K. \quad (1)$$

In other words, only the  $k$ -th element of  $\mathbf{C}$  is 1, and all others are 0.

This supports the network in distinguishing different sequential patterns in the training dataset and generating each trained pattern. Secondly, each internal state of CERNet consists of a pair of prior and posterior values, each of which corresponds to the internal state before and after observation, respectively. Throughout the operation of this model, internal states are influenced by the discrepancy between prediction and observation, and updated to minimize those discrepancies.

During forward computation of CERNet, the prior  $p$  internal state  $h_t^{n,p}$  at time  $t$  and layer  $n$  is updated based on the recurrent dynamics, top-down inputs (except for the top layer), and the cause input (for the top layer only), as follows:

$$\begin{aligned} h_t^{n,p} = & \left(1 - \frac{1}{\tau_h^n}\right) h_{t-1}^{n,q} \\ & + \frac{1}{\tau_h^n} \left( \tanh(h_{t-1}^{n,q}) W_r^{n\top} + \delta_{n=N} C W_c^\top \right. \\ & \left. + \delta_{n < N} \tanh(h_t^{n+1,q}) W_{hh}^{n\top} + b_r^n \right), \end{aligned} \quad (2)$$

where  $\tau_h^n$  is a time constant controlling the update rate at layer  $n$ .  $W_r^n$  and  $W_{hh}^n$  are the recurrent and top-down weight matrices, respectively, and  $b_r^n$  is a bias term.  $\delta_{n < N}$  and  $\delta_{n=N}$  are indicator functions that activate the corresponding term only when the condition is satisfied.  $C$  is the class-embedding vector, applied only in the top layer ( $n = N$ ).

Prediction at time  $t$ ,  $x_t^0$ , is then calculated in the bottom layer as follows:

$$x_t^0 = \tanh(h_t^{0,p}) W_o^\top + b_o, \quad (3)$$

where  $h_t^{0,p}$  is the corresponding prior hidden state obtained from Eq. 2.  $W_o$  and  $b_o$  denote the output-layer weight matrix and bias vector, respectively.

After the model computed the prediction for time  $t$ , it is fed with the observation  $\tilde{x}_t^o$ . Utilizing the bottom-up error between the prediction and the observation  $\varepsilon_t (= x_t^o - \tilde{x}_t^o)$ , the posterior  $q$  internal state  $h_t^{n,q}$  deviates from the prior internal state to minimize the bottom-up error:

$$S_t^n = \begin{cases} \varepsilon_t W_o, & n = 0, \\ \varepsilon_t^{n-1} W_{hh}^{n-1}, & n > 0. \end{cases} \quad (4)$$

$$h_t^{n,q} = h_t^{n,p} - \alpha_h^n [1 - \tanh^2(h_t^{n,p})] S_t^n, \quad (5)$$

$$\varepsilon_t^n = h_t^{n,p} - h_t^{n,q}, \quad \varepsilon_t = x_t^0 - x_t^{\text{obs}}, \quad (6)$$

where  $\alpha_h^n$  represents the learning-rate parameter for each layer, and  $S_t^n$  provides the appropriate error-driven—sensory prediction error at  $n=0$  (bottom layer) or hierarchical prediction

error from the layer below for  $n > 0$  (top layer). Layer-wise prediction errors  $\varepsilon_t^n$  are recursively fed upward via  $S_t^{n+1}$ , while  $\varepsilon_t$  denotes the sensory prediction error at the bottom layer.

The proposed CERNet operates under a unified predictive-coding framework across three functional phases: 1) a *training phase* for learning from labeled demonstrations, 2) a *generative phase* for autonomous reproduction of motion, and 3) an *inference phase* for online recognition of task identity. All three phases share the same network architecture and internal dynamics, relying on the minimization of prediction error to update both hidden states and, when applicable, the class embedding vector  $\mathbf{C}$ .

During the *training phase*, the network receives labeled demonstrations: the observation sequence (end-effector trajectories) together with the corresponding class identifier; the model learns its parameters. After training, in the *generation phase*, a target class identifier is specified and the network autonomously reproduces the motion in closed loop on the robot; internal states are updated online during execution, but the learned parameters and the class representation remain fixed. In the *inference phase*, the observed end-effector trajectory is provided incrementally from the robot; the model updates its internal states online and *infers* the class identity from the partial observation.

During the *training phase*, after computing the forward pass, the model parameters,  $\theta = \{W_o, b_o, W_c, W_{hh}^n, W_r^n, b_r^n\}$  are optimised by back-propagation to minimise the *prediction-error loss* (equivalent, up to a constant factor, to the variational free energy):

$$\mathcal{L}_{\text{PE}} = \frac{1}{2} \sum_t \left( \|\varepsilon_t\|_2^2 + \sum_{n=0}^N \alpha_n \|\varepsilon_t^n\|_2^2 \right), \quad (7)$$

while the variational free-energy (VFE) is computed in parallel only for logging and analysis. All layer-specific gradients are concatenated into a single vector

$$g_t = \left[ \frac{\partial \mathcal{L}}{\partial W_o}, \frac{\partial \mathcal{L}}{\partial b_o}, \frac{\partial \mathcal{L}}{\partial W_c}, \frac{\partial \mathcal{L}}{\partial W_{hh}^{0:N-1}}, \frac{\partial \mathcal{L}}{\partial W_r^{0:N}}, \frac{\partial \mathcal{L}}{\partial b_r^{0:N}} \right], \quad (8)$$

which is subjected to global  $L_2$ -norm clipping ( $G_{\text{max}} = 1.0$ ) and updated with Adam,

$$\begin{aligned} \hat{g}_t &= g_t \cdot \min\left(1, \frac{G_{\text{max}}}{\|g_t\|_2 + \varepsilon}\right), m_t = \beta_1 m_{t-1} + (1 - \beta_1) \hat{g}_t, \\ v_t &= \beta_2 v_{t-1} + (1 - \beta_2) \hat{g}_t^2, \theta_t = \theta_{t-1} - \eta \frac{m_t / (1 - \beta_1^t)}{\sqrt{v_t / (1 - \beta_2^t)} + \varepsilon}, \end{aligned} \quad (9)$$

with  $\beta_1=0.9$ ,  $\beta_2=0.999$ , and  $\varepsilon=10^{-6}$ .

During both the *generation phase* and the *inference phase*, the model no longer updates its network weights. Instead, the internal hidden states  $h_t^n$  are dynamically updated at each timestep to minimize prediction errors, using the same error-driven update rules as in the training phase. Specifically, Eqs. (4)–(6) govern the propagation of prediction errors and the corresponding updates to each layer's internal state, thereby enabling the model to adapt its predictions online without parameter changes.

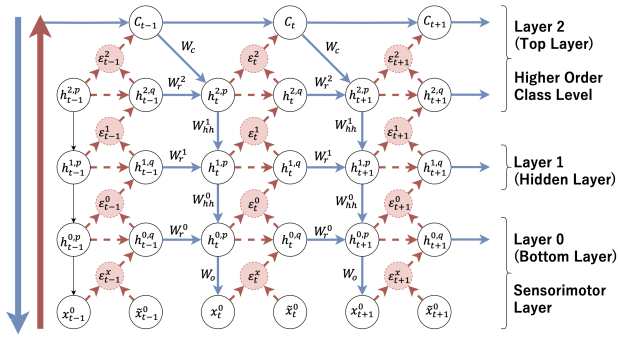


Fig. 1. Schematic illustration of CERNet. The model integrates top-down predictions and bottom-up errors across hidden states in multiple layers. Blue and red arrows indicate computation for forward propagation and error propagation, respectively. Refer to Eqs. 1–9 for detailed explanation of each variable.

In addition to updating  $h$ , the *inference phase* further refines the class embedding vector  $\mathbf{C}$ . This is achieved through a procedure known as **past reconstruction error minimization**, inspired by inference techniques in stochastic PC models such as PV-RNN [16], where internal states are optimized over past trajectories. Rather than relying solely on the current timestep,  $\mathbf{C}$  is iteratively updated to minimize the accumulated prediction error over the full observed sequence up to time  $t$ . While the network weights remain fixed, the embedding vector is refined using gradient descent:

$$\mathbf{C} \leftarrow \mathbf{C} - \alpha_c \cdot \frac{1}{t - (t - W)_+ + 1} \sum_{\tau=(t-W)_+}^t \nabla_{\mathbf{C}} \mathcal{L}_{\text{PE}}(\tau), \quad (10)$$

with the gradients  $\nabla_{\mathbf{C}} \mathcal{L}_{\text{PE}}(\tau)$  back-propagated through the hierarchical PC layers using the same equations as above. To reduce computational cost, the gradients are averaged only over a sliding window of length  $W$ , i.e.,  $\tau$  runs from  $(t - W)_+$  to  $t$  with normalization by the number of terms, where  $(\cdot)_+ = \max(\cdot, 0)$ .

This inference–update cycle is repeated for  $n_{\text{iter}}$  iterations at each timestep, where the current class embedding  $\mathbf{C}$  is used to regenerate a prediction sequence  $x_{0:t}^{\text{pred}}$ , and then updated to reduce its reconstruction error. Through this iterative process,  $\mathbf{C}$  gradually drifts toward the latent subspace most consistent with the observed sequence. We refer to this approach as **past reconstruction**, which unifies recognition and generation within the same predictive-coding framework.

### B. Humanoid Robot: Reachy

All physical experiments were conducted on Reachy 2021<sup>1</sup>, a humanoid robotic platform developed by Pollen Robotics. In this study, we used the left arm (7 DoF), which includes Dynamixel motors controlled via the Reachy SDK over a UDP interface. Joint commands were sent using the minimum-jerk interpolation mode provided by the `reachy-sdk`, with a control frequency of 20 Hz. The end-effector position

(3-dimensional Cartesian) was computed using the built-in forward/inverse kinematics solvers, and all training trajectories were recorded through kinaesthetic teaching. During the interaction phase, the predicted end-effector trajectories were converted into joint angles via inverse kinematics and executed directly on the robot. To ensure consistent execution, each trajectory was preceded by a reset-to-home motion. No external sensor (e.g., camera or IMU) was used in this study.

## IV. EXPERIMENT

All experiments were conducted as described below. To promote transparency and reproducibility, **all raw trajectory data, model weights, training scripts, and analysis code are publicly available**.<sup>2</sup>

### A. Experimental Setup

1) *Train Data-Set Preparation*: To collect training data, we physically guided the left arm of the humanoid robot Reachy to draw each letter of the English alphabet on a table. The corresponding end-effector positions were estimated using Reachy’s built-in inverse kinematics and recorded at a sampling rate of 20 Hz. Each sequence was recorded for 100 time steps without interpolation. After completing each character, we held the end-effector stationary for the remaining steps to preserve a consistent sequence length. The final dataset thus comprises 26 distinct trajectory sequences, each representing one alphabet character, and each sequence consists of 100 time steps.

2) *Network Configuration*: We prepared six CERNet variants with varying hidden sizes and depths, as summarized in Table I. These include three single-layered models (A: **SingleMini**, B: **SingleStandard**, C: **SingleLarge**) and three hierarchical models (D: **MultiMini**, E: **MultiStandard**, F: **MultiLarge**), each differing in total parameter count and node allocation.

This setup allows us to compare the effectiveness of the hierarchical architecture against single-layered models under similar model capacities. To ensure fair comparison, the number of parameters in each pair of models—Mini, Standard, and Large—has been adjusted so that the single-layered and hierarchical models have approximately the same total number of parameters. To guarantee statistical reliability in the analysis of experimental results, we have trained each model 10 times with identical parameters (Table I). # of nodes,  $\tau_h$ ,  $\alpha_h$ , # of params indicate the number of nodes in each hidden layer (with larger sizes in the top layer and smaller ones in the bottom layer for hierarchical models), the time constant, the update rates of posterior hidden state in each hidden layer and the number of trainable parameters, respectively.

Each model was trained for 10,000 epochs to minimize the *prediction-error loss* shown in Eq. 7. The network weights were initialized randomly, using a different random seed for each training run.

<sup>1</sup>Pollen Robotics: <https://pollen-robotics.github.io/reachy-2021-docs/>

<sup>2</sup>[https://git.cyu.fr/hiroki/cernet\\_paperresultreproduction](https://git.cyu.fr/hiroki/cernet_paperresultreproduction) (eCILL license).

TABLE I  
MODEL CONFIGURATIONS FOR EACH CERNETS

Type	ID	Nodes	$\tau_h$	$\alpha_h (\times 10^{-2})$	Params
Single	A)	[50]	[10]	[1]	3.9k
	B)	[150]	[10]	[1]	22.9k
	C)	[300]	[10]	[1]	90.9k
Multi	D)	[50,15,7]	[10,20,40]	[1, 0.05, 0.0005]	4.6k
	E)	[120,40,20]	[10,20,40]	[1, 0.05, 0.0005]	22.9k
	F)	[250,70,20]	[10,20,40]	[1, 0.05, 0.0005]	90.9k

3) *Alphabet Drawing Experiment*: After training, the model parameters loaded for evaluation were taken from the epoch with the lowest loss recorded during training. All 60 trained network instances were first evaluated in simulation to verify closed-loop prediction stability across the full set of 26 alphabet characters. Based on this screening, four models—**MultiStandard**, **MultiLarge**, **SingleStandard**, and **SingleLarge**—were selected for physical experiments on the humanoid robot *Reachy*, as they demonstrated sufficiently stable and safe behavior in the simulated loop.

In the real-world closed-loop interaction setting, the network produced a 3D Cartesian prediction of the left-arm end-effector position at each control step. This prediction was forwarded to *Reachy*'s built-in impedance controller, which converted it to joint angle commands via inverse kinematics and moved the arm accordingly. The controller then returned the measured end-effector position to the network, closing the prediction–observation loop. This cycle was repeated for 100 consecutive timesteps while the robot traced each of the 26 alphabet characters learned during training. The class embedding vector was initialized to the one-hot vector corresponding to the correct character for each trial. The drawn trajectory was then evaluated by calculating the Dynamic Time Warping (DTW) score, a distance-based measure that quantifies similarity between two temporal sequences by optimally aligning them even if they vary in speed or timing, with respect to the training data. Throughout the experiment, the round-trip latency between the host computer and the robot remained below 30 ms.

4) *Perturbation Resistance Experiment Using Reachy*: In addition to the closed-loop interaction setting, we have programatically inserted an external perturbation into the observation of the end-effector position during the experiment for robustness evaluation. Specifically, while *Reachy* was operated by a trained network, we applied an external disturbance from time step 40 to 45, which forced its arm to drift from the predicted trajectory. We then evaluated the ability of the model to recover and return to the learned trajectory. This perturbation experiment was conducted only on model F: **MultiLarge**, as other models did not consistently demonstrate stable closed-loop behavior on the physical platform and were deemed unsuitable for safe evaluation, based on the results of the alphabet drawing experiment in Section IV-A.3.

5) *Class Inference & Confidence Estimation Experiment*: In a separate experiment, we evaluated the ability of the model to recognize the intended class based solely on observed robot

motion. For this test, we used the same physical platform as in the generation experiment, but reversed the interaction direction: the robot arm was manually moved along one of the 26 training trajectories, while the model passively observed the resulting end-effector positions in real time. To simulate uncertainty in initial prior knowledge, the 26-dimensional class embedding  $\mathbf{C}$  was initialized with  $\mathcal{N}(0, 0.1^2)$  noise. At each control step, the observed trajectory up to the current timestep was fed into the model, and the class embedding vector  $\mathbf{C}$  was updated by minimizing the accumulated past prediction error, as described in Sec. III-A. We used a learning rate  $\alpha_c = 2.5 \times 10^{-2}$ , a window size of 50 timesteps and performed  $n_{\text{iter}} = 15$  past-reconstruction iterations per timestep. The computation time per timestep did not exceed approximately 130 ms.

Recognition was performed using the model that achieved the lowest DTW error in alphabet drawing experiment (i.e., the best-performing **MultiLarge** instance). Following each update, the model also predicted motion of the future timesteps using the updated  $\mathbf{C}$ , enabling real-time class inference and future prediction.

Each of the 26 motion classes was tested in 10 independent trials with different random initializations of  $\mathbf{C}$ , resulting in a total of **260 recognition trials**. At the end of each trial, the predicted class was defined as the index of the maximum value in  $\mathbf{C}$  (Top-1), and the second-highest index was recorded as the Top-2 prediction. Accuracy was evaluated across all trials using Top-1 and Top-2 classification rates. In addition, the final reconstruction error of internal states was recorded for each trial, allowing to compare the error distributions across three outcome groups: Top-1 correct, Top-2 correct, and Incorrect.

## V. RESULT AND DISCUSSION

### A. Training Result

As a baseline comparison, we begin by summarizing the training performance of all 60 CERNet networks (10 runs per model across 6 types). Table II shows the average minimum prediction-error loss per model type, corresponding to the configurations listed in Table I.

Model performance improved with size across both architectures. For example, among single-layer CERNets, Mini vs. Standard ( $p = 1.1 \times 10^{-6}$ ) and Standard vs. Large ( $p = 0.015$ ) were both significant. Multi-layer variants showed the same trend. Importantly, at every model size, the multi-layer version consistently outperformed its single-layer counterpart, with particularly strong effects in the Standard ( $p = 3.8 \times 10^{-9}$ ) and Large ( $p = 5.1 \times 10^{-13}$ ) settings. Even in the Mini configuration, MultiMini showed a measurable advantage ( $p = 0.0041$ ). These results clearly indicate that hierarchical structure improves prediction accuracy, especially at higher capacities, but even small models benefit from layering.

### B. Alphabet Drawing Experiment

In this section, we present the results of the alphabet drawing experiment in simulation and real-robot. First, we

TABLE II  
MINIMUM PREDICTION ERROR LOSS DURING TRAINING ( $\times 10^{-3}$ )

	Mini	Standard	Large
Single	A) $7.51 \pm 0.81$	B) $4.60 \pm 0.91$	C) $3.66 \pm 0.55$
Multi	D) $4.53 \pm 2.59$	E) $0.90 \pm 0.54$	F) <b><math>0.26 \pm 0.10</math></b>

show the Dynamic Time Warping (DTW) score to the ground-truth training trajectories of each model as a quantitative evaluation. Next, we illustrate the actual alphabet drawings both in simulation and real-robot as a qualitative evaluation.

1) *Quantitative Evaluation: Dynamic Time Warping score:* Table III summarizes DTW scores between generated and training trajectories.

TABLE III  
DTW SCORE ( $\times 10^{-1}$ )

Model ID	DTW in Simulation	DTW in Robot Exp.
A) SingleMini	$1.70 \pm 0.84$	-
B) SingleStandard	$1.23 \pm 0.78$	$1.90 \pm 0.87$
C) SingleLarge	$1.03 \pm 0.70$	$2.32 \pm 1.52$
D) MultiMini	$1.07 \pm 0.73$	-
E) MultiStandard	$0.44 \pm 0.28$	$1.13 \pm 0.48$
F) <b>MultiLarge</b>	<b><math>0.25 \pm 0.11</math></b>	<b><math>0.95 \pm 0.39</math></b>

Table III further illustrates the impact of hierarchical structure in simulation. The three-layer *MultiLarge* networks achieved the lowest DTW score ( $0.25 \pm 0.11$ ), reducing alignment error by 76% compared to *SingleLarge* ( $1.03 \pm 0.70$ ). Although differences at smaller scales were less pronounced (e.g., MultiMini vs. SingleMini), the multi-layer models consistently outperformed their single-layer counterparts at all sizes.

Additionally, while all models exhibit some performance degradation when transferred from simulation to the physical robot experiment, this effect is particularly visible in the single-layer architectures. As shown in Table III, the average DTW scores increase across all models, yet the scores alone do not fully capture the severity of the degradation in the single-layer models. In particular, qualitative inspection of the reproduced motions (detailed in the next section) reveals that multi-layer models maintain legible trajectories that still resemble the intended characters, whereas single-layer models often fail to generate discernible shapes, despite their average DTW scores remaining within a similar range.

This discrepancy highlights the limitation of purely quantitative metrics in capturing perceptual fidelity and further supports the practical advantage of hierarchical architectures in real-world deployment.

2) *Qualitative Analysis: Actual Alphabet Drawing:* To qualitatively evaluate the legibility of reproduced trajectories, we selected five representative characters (b, e, k, l, m) and visualized the motions generated by the best-performing instance of each model. The results are shown in Fig. 2, illustrating how model architecture affects stroke structure and spatial coherence.

Multi-layered CERNets produce visually recognizable characters that align well with the training data, consistent

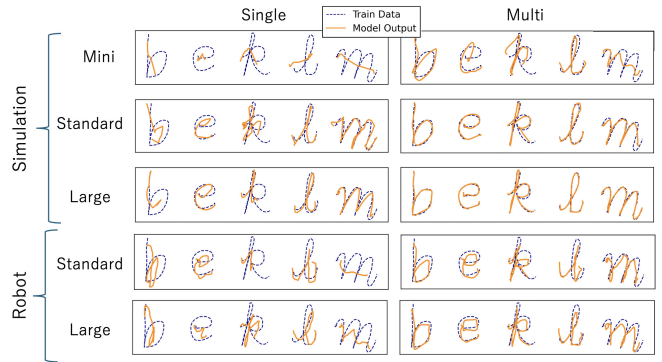


Fig. 2. Reproduction of the letters *b*, *e*, *k*, *l*, *m* by the best-performing networks of each model type. Each row corresponds to a different model scale and experimental condition (simulation or real robot), while columns compare single-layer and multi-layer architectures. Dotted lines indicate the original training trajectories, and solid lines represent the generated motions by the models.

with their low DTW scores. In contrast, single-layer models show greater distortion, especially in complex characters like 'k' and 'b', reflecting their higher DTW scores. Despite having fewer parameters, MultiMini (D) matches SingleLarge (C) in output quality, highlighting the efficiency of hierarchical structures.

When transferred to the real robot, all models exhibit some performance degradation. However, while single-layer models often produce illegible outputs, multi-layer models still retain clearly recognizable character shapes. This further supports the practical robustness of hierarchical architectures in noisy, real-world settings. These observations imply that hierarchical architectures improve not only numerical alignment with training data, but also the qualitative correctness of generated motion.

### C. Perturbation Resistance Experiment

In this section, we present the result of the perturbation experiment conducted on Reachy using the proposed multi-layered CERNets. A demonstration video of a representative trial is available,<sup>3</sup> which shows the robot drawing the letter *p* while an external disturbance is applied between timestep 40 and 45 (see Fig. 3). Fig. 3 shows a representative trial where the robot was drawing the letter *p* and an external disturbance was applied between timestep 40 and 45.

Fig. 3 a) and b) reveal that the prediction error, used to update internal states, increased sharply during the perturbation period, indicating that the model identified the deviation between prediction and observation. Fig. 3 c) demonstrates that, once the external disturbance ceased, the robot's trajectory gradually converged back to the original path, confirming the self-correcting capability of the model. Moreover, as shown in Fig. 3 d) (for a different trial drawing the letter *g*), the predicted trajectory after the disturbance

<sup>3</sup>Video available at: [https://git.cyu.fr/hiroki/cernet\\_paperresultreproduction/-/blob/main/perturbation\\_youtube.mp4](https://git.cyu.fr/hiroki/cernet_paperresultreproduction/-/blob/main/perturbation_youtube.mp4)

gradually shifted, illustrating how future predictions are updated through internal error minimization.

Overall, this experiment highlights the ability of CERNet to withstand external disturbances by dynamically adjusting its internal states, spontaneously recovering the learnt trajectory.

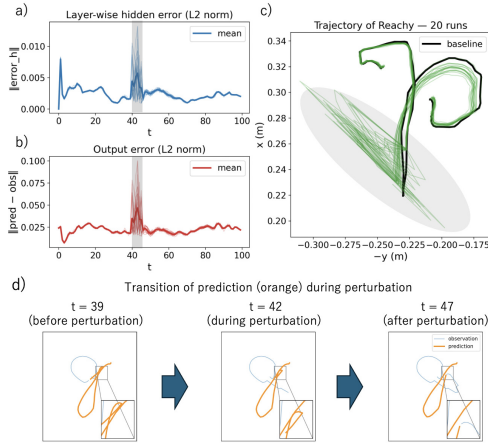


Fig. 3. Perturbation recovery during alphabet reproduction on Reachy using multi-layered CERNet model. (a) and (b): Layer-wise prediction errors with respect to time while reproducing letter *p*. The grey area indicates the timesteps where perturbation was injected. (c) End-effector trajectory while reproducing letter *p*. (d) Future trajectory prediction while writing *g* under an injected perturbation: future prediction of the trajectory realigns as internal states adapt.

#### D. Class Inference & Confidence Estimation Experiment

In this section, we present the result of the class inference and confidence-level estimation experiment conducted on Reachy using multi-layered CERNet model. A demonstration video of a representative trial is available,<sup>4</sup> which illustrates how the model infers the intended character "b" in real time given the trajectory. The goal of this experiment is: (i) to evaluate the model's ability to infer the intended trajectory class in real time from observed motions, and (ii) to assess whether the internal prediction error reflects the model's confidence in its own inference.

To evaluate class inference, we analyzed the 260 physical trials described in the experiment section, in which Reachy's arm was manually moved along one of the 26 training trajectories while the model observed the end-effector positions. The class embedding vector  $\mathbf{C}$  was updated online at each timestep by minimizing past prediction error, and the final inferred class was determined by the index with the highest value in  $\mathbf{C}$  (Top-1), with the second-highest taken as Top-2.

Across 260 trials (10 trials per class), the model achieved a Top-1 classification accuracy of **68%** and a Top-2 accuracy of **81%**. These results demonstrate that the model is capable of performing real-time recognition over a large set of motion classes in a noisy physical environment.

To evaluate the model's implicit confidence estimation, we compared the final mean squared error (MSE) of the

<sup>4</sup>Video available at: [https://git.cyu.fr/hiroki/cernet\\_paperresultreproduction/-/blob/main/Experiment3\\_InferLetter.mp4](https://git.cyu.fr/hiroki/cernet_paperresultreproduction/-/blob/main/Experiment3_InferLetter.mp4)

past reconstruction loss  $L_{\text{past}}$  among three groups: correctly recognized at Top-1, correctly recognized at Top-2, and incorrect. The results are summarized below:

To assess whether the model's internal reconstruction error reflects recognition accuracy, we compared the final past reconstruction MSE across three categories: correctly recognized at Top-1, recognized at Top-2, and incorrect predictions. Statistical comparisons using the Mann-Whitney U test revealed that the MSE in the Top-1 group was significantly lower than that in the Top-2 group ( $U = 1723.0$ ,  $p = 0.00019$ ), and also significantly lower than the Incorrect group ( $U = 2038.0$ ,  $p < 10^{-8}$ ). Although the Top-2 group tended to show lower error than the Incorrect group, the difference was not statistically significant ( $U = 622.0$ ,  $p = 0.060$ ). It should be noted that the variance of the errors within each group was relatively large, suggesting that further analyses will be required to fully characterize the relationship between error magnitude and confidence. Overall, these results suggest that the model's internal prediction error serves as an implicit confidence measure: the lower the error, the more likely the inferred class is correct.

These findings suggest that the model is capable of internally differentiating between correct and incorrect inferences based on its own prediction error. In particular, lower MSE is correlated with more accurate recognition, implying an emergent internal confidence indicator.

As an illustrative example, Fig. 4 shows a representative trial in which the model infers the character "b". As time progresses, the updated class embedding gradually drifts toward the subspace corresponding to "b", and both past reconstruction and prediction align more closely with the ground-truth trajectory. Furthermore, Fig. 5 displays a boxplot of the final past reconstruction MSE for each recognition category. The trend confirms that the model exhibits higher prediction error when its recognition is incorrect, and lower error when it is confident and correct, supporting our claim of self-monitoring capability.




Fig. 4. Time development of the class embedding vector of CERNet in inference mode. The top row illustrates the time development of the prediction and the observation from the robot. Whereas the bottom row illustrates the time development of the intrinsic prediction of the observed motion.

## VI. CONCLUSION

This study presented a unified neural architecture, CERNet, Class-Embedding Predictive-Coding RNN, that integrates

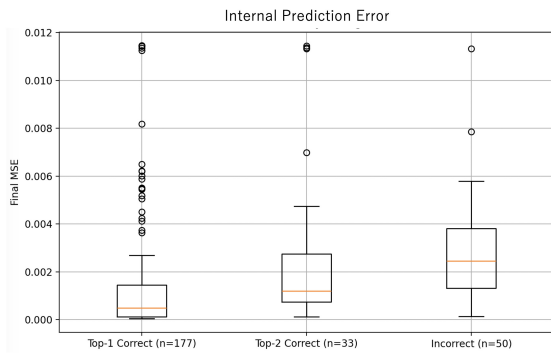


Fig. 5. Final mean squared error (MSE) of the internal prediction error, averaged over time, grouped by recognition outcome. Each boxplot shows the distribution of reconstruction error across trials where the correct class was identified at Top-1 (left), at Top-2 (middle), or was not among the top two predictions (right). Lower errors correspond to higher inference accuracy, indicating that the model’s prediction error implicitly reflects its confidence.

recognition, generation, and confidence estimation within a single model and operates in real time on a physical humanoid robot. By combining a multi-layer predictive-coding framework with a shared class embedding vector, the proposed model enables scalable learning and regeneration of complex motor repertoires while minimizing interference between patterns. We demonstrated that this architecture successfully stores and reproduces 26 handwritten alphabet trajectories on the humanoid robot, Reachy, achieving a 76% reduction in reproduction error compared to a single-layer baseline with almost an identical number of parameters. The hierarchical structure not only demonstrated accuracy but also enhanced robustness, as shown in Experiment 2 by spontaneous recovery from physical perturbations through online error correction.

Furthermore, we showed that the same model can infer the demonstrated motion class in real time, achieving 68% Top-1 and 81% Top-2 classification accuracy without retraining. Importantly, the past-reconstruction error serves as an indicator of confidence: trials with correct recognition exhibited significantly lower past-reconstruction error than incorrect ones, suggesting that the model can internally assess the certainty of its inference.

Together, these findings position CERNet as a lightweight, real-time-capable control model that integrates trajectory generation, class recognition, and confidence estimation within a single framework. Its ability to encode a large repertoire of motion patterns, maintain robustness to perturbations through internal state adaptation, and infer motion class with measurable confidence makes it particularly well-suited for embodied interaction scenarios. The main contribution of this study is to propose an integrated framework and demonstrate its feasibility on a humanoid robot. The emphasis of the present work is therefore on the architectural unification of generation, online recognition, and intrinsic confidence estimation within a single predictive-coding dynamical system. More extensive empirical comparisons with alternative sequence models remain as future work.

We believe this architecture provides a promising founda-

tion for future research in human–robot interaction, where smooth and natural collaboration requires real-time understanding of intent, execution of adaptive behavior, and intrinsic evaluation of inference reliability. Future work will explore its extension to online learning and integration with multimodal sensory inputs such as vision, enabling more complex and context-aware interaction behaviors.

## REFERENCES

- [1] W. Lotter, G. Kreiman, and D. Cox, “Deep predictive coding networks for video prediction and unsupervised learning,” *arXiv preprint arXiv:1605.08104*, 2016.
- [2] J. Tani, “Accounting for the minimal self and the narrative self: Robotics experiments using predictive coding,” in *CEUR workshop proceedings*, vol. 2287. Okinawa Institute of Science and Technology Graduate University, 2019.
- [3] S. Qiu, S. Bhattacharyya, D. Coyle, and S. Dora, “Deep predictive coding with bi-directional propagation for classification and reconstruction,” *Neural Networks*, p. 107785, 2025.
- [4] A. Granier, M. A. Petrovici, W. Senn, and K. A. Wilmes, “Confidence and second-order errors in cortical circuits,” *PNAS nexus*, vol. 3, no. 9, p. pgae404, 2024.
- [5] A. Sharafeldin, N. Imam, and H. Choi, “Active sensing with predictive coding and uncertainty minimization,” *Patterns*, vol. 5, no. 6, 2024.
- [6] Y. Yamashita and J. Tani, “Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment,” *PLoS computational biology*, vol. 4, no. 11, p. e1000220, 2008.
- [7] R. P. Rao and D. H. Ballard, “Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects,” *Nature neuroscience*, vol. 2, no. 1, pp. 79–87, 1999.
- [8] Y. Duan, M. Andrychowicz, B. Stadie, O. Jonathan Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba, “One-shot imitation learning,” *Advances in neural information processing systems*, vol. 30, 2017.
- [9] C. Devin, A. Gupta, T. Darrell, P. Abbeel, and S. Levine, “Learning modular neural network policies for multi-task and multi-robot transfer,” in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 2169–2176.
- [10] J. Merel, L. Hasenclever, A. Galashov, A. Ahuja, V. Pham, G. Wayne, Y. W. Teh, and N. Heess, “Neural probabilistic motor primitives for humanoid control,” *arXiv preprint arXiv:1811.11711*, 2018.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, E. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [12] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The Inter. Journal of Robotics Research*, p. 02783649241273668, 2023.
- [13] J. Ho and S. Ermon, “Generative adversarial imitation learning,” *Advances in neural information processing systems*, vol. 29, 2016.
- [14] D. Ha and J. Schmidhuber, “Recurrent world models facilitate policy evolution,” *Advances in neural information processing systems*, vol. 31, 2018.
- [15] K. Friston and S. Kiebel, “Predictive coding under the free-energy principle,” *Philosophical transactions of the Royal Society B: Biological sciences*, vol. 364, no. 1521, pp. 1211–1221, 2009.
- [16] A. Ahmadi and J. Tani, “A novel predictive-coding-inspired variational rnn model for online prediction and recognition,” *Neural computation*, vol. 31, no. 11, pp. 2025–2074, 2019.
- [17] R. P. Rao, “A sensory–motor theory of the neocortex,” *Nature neuroscience*, vol. 27, no. 7, pp. 1221–1235, 2024.
- [18] L. Annabi, A. Pitti, and M. Quoy, “Bidirectional interaction between visual and motor generative models using predictive coding and active inference,” *Neural Networks*, vol. 143, pp. 638–656, 2021.
- [19] —, “Continual sequence modeling with predictive coding,” *Frontiers in Neurobotics*, vol. 16, p. 845955, 2022.
- [20] A. Pitti, P. Gaussier, and M. Quoy, “Iterative free-energy optimization for recurrent neural networks (inferno),” *PLoS One*, vol. 12, no. 3, p. e0173684, 2017.
- [21] A. Pitti, M. Quoy, C. Lavandier, and S. Boucenna, “Gated spiking neural network using iterative free-energy optimization and rank-order coding for structure learning in memory sequences (inferno gate),” *Neural Networks*, vol. 121, pp. 242–258, 2020.